



UNIVERSITY OF
MILANO - BICOCCA
Department of Material Science



ITALIAN INSTITUTE OF
TECHNOLOGY
Atomistic simulations

PhD program in Materials Science and Nanotechnology
· XXXVII Cycle ·

Machine Learning and Molecular Dynamics *in operando* simulations of catalytic reactions

by

Simone Perego

Mat. 813901

A dissertation submitted for the degree of
Dottore di Ricerca Ph.D.

Advisor:

Prof. Michele Parrinello

Supervisor:

Dr. Luigi Bonati

Tutor:

Prof. Marco Bernasconi

ACADEMIC YEAR 2023-2024

Abstract

The urgent need to transition to a green economy has generated significant interest in ammonia as an effective hydrogen carrier due to its high hydrogen content and ease of storage. To make this possible, efficient routes for extracting hydrogen via ammonia decomposition must be found. Ammonia thermal decomposition on iron-based catalysts has received much interest in the past decades, as iron offers an attractive balance between catalytic efficiency and sustainability. However, an atomic-level understanding of the reaction mechanisms, particularly under high-temperature *operando* conditions, has not been fully elucidated. Traditional models, often based on static representations of catalytic surfaces, fail to capture the dynamic nature of surface reactions, which is crucial for explaining catalytic activity and long-term stability. In this Thesis, we present a comprehensive computational investigation of ammonia decomposition on iron and iron-cobalt alloy catalysts, addressing these challenges via molecular dynamics simulations powered by machine learning and enhanced sampling techniques.

We first investigated ammonia decomposition on Fe(110) and Fe(111) surfaces, representing the two limiting cases of the low miller index surfaces, characterized by differing atomic densities and reactivities. Our simulations revealed that dynamic effects at *operando* temperatures critically influence all reaction stages. We demonstrated how the mobility of surface atoms influences adsorption, diffusion, and dehydrogenation processes, leading to an ensemble of reactive pathways rather than a well-defined static one. Nonetheless, we found a unified description of the reaction mechanism through charge transfer analysis along these pathways, providing a comprehensive understanding of the catalytic behavior of iron surfaces.

After the dehydrogenation steps in the low surface coverage limit, we studied the fate and influence of adsorbed nitrogen (N^*) on the surface morphology and activity.

For Fe(110), we investigated the nitridation mechanisms during ammonia decomposition in collaboration with experimentalists, who observed the formation and decomposition of iron nitrides during catalyst activation and ammonia decomposition. Our molecular dynamics simulations provide insights into the early stages of nitridation by showing that nitrogen migration into the iron bulk catalyst is favored over recombinative desorption, affecting nitrogen storage and release dynamics. In particular, we showed that when finite coverage effects are considered, the N^* dissolution process, which leads to bulk nitride formation, is enhanced. Similarly, we studied the influence of a finite N^* coverage on ammonia synthesis, particularly on the molecular nitrogen cleavage step for the Fe(111) surface. We found that N^* atoms stabilize triangular surface structures, a sign of a frustrated phase

transition at the surface. We discovered that N^* lateral interactions reduce, but not eliminate, iron atom mobility. Here, the effects of dynamics are even more relevant and counter-intuitive, as the reduction in the number of active sites is counterbalanced by an increase in their lifetime, resulting in the molecular nitrogen dissociation barrier remaining essentially unchanged across the studied coverages.

Leveraging the knowledge gained from iron catalysts, we then investigated iron-cobalt alloys, particularly FeCo, which have recently been proposed as an alternative to iron because of the higher activity shown in ammonia decomposition. Due to its complexity and prohibitive computational cost, this alloy has never been investigated via atomistic simulations. Thus, we developed a novel data-efficient protocol for constructing reactive machine learning potentials, focusing on dataset construction by targeting all relevant, particularly reactive, configurations. This is achieved by combining enhanced sampling methods, Gaussian processes, and graph neural networks to generate and select critical configurations based on local environment uncertainty. As a result, we could characterize the mechanisms of several reactions involved at *operando* temperatures also accounting for lateral interactions, with only a fraction ($\sim 1/20$) of the usual computational effort. Our simulations showed substantial similarity between the FeCo(1 1 0) and Fe(1 1 0) catalytic mechanisms in the elementary steps of the reaction. We ascertained a double promotional effect of cobalt, which not only reduces the reaction’s rate-determining step but also increases the barrier for nitrogen migration into the bulk, preventing catalyst nitridation. This microscopic interpretation is confirmed by experimental observations that no bulk nitrides are formed.

Our findings indicate that it is the delicate balance between different surface processes that explains catalytic activity, such as the formation of iron surface- or bulk-nitrides under certain conditions and how iron-cobalt prevents it. This sheds light on the microscopic promotional effect of cobalt, driven by the synergistic effects of structural and electronic promotion. In conclusion, our results underscore the complexity of heterogeneous catalysis and its intrinsic dynamic nature, paving the way for even more realistic simulations.

Contents

List of abbreviations	vi
Outline	vii
1 Introduction	1
1.1 Ammonia as hydrogen vector	2
1.2 Thermocatalytic decomposition	3
1.2.1 Iron catalysts	3
1.2.2 Bimetallic alloy catalysts	5
1.3 Dynamical modelling of heterogeneous catalysis	6
2 Computational Methods	7
2.1 Molecular Dynamics	7
2.1.1 Interatomic Force Description	10
2.2 Machine Learning Interatomic Potentials	13
2.2.1 Dataset	15
2.2.2 Descriptors	16
2.2.3 Architectures	18
2.2.4 Uncertainty Quantification in MLPs	22
2.3 Building machine learning potentials for chemical reactions	23
2.3.1 Active learning via <i>query-by-committee</i>	24
2.4 Free energies and enhanced sampling	27
2.4.1 Probabilities, free energies, and manipulating it	28
2.4.2 On-the-fly Probability Enhanced Sampling (OPES)	31
2.4.3 Collective variables	32
2.5 Machine learning for partial charges	35
2.6 Other analysis tools	36
2.6.1 Transition states characterization	36
2.6.2 Surface analysis	37
2.7 Computational details	40
2.7.1 NH ₃ decomposition on iron	40
2.7.2 Effect of N^* on Fe(111)	46
2.7.3 The fate of nitrogen on Fe(110)	49
2.7.4 Ammonia decomposition on iron-cobalt	51

3	Ammonia adsorption and decomposition on iron surfaces	56
3.1	Fe surface behavior at high temperature	57
3.2	Interaction between NH_x and iron surfaces	59
3.2.1	Fe(110)	59
3.2.2	Fe(111)	62
3.3	Dehydrogenation steps	64
3.3.1	Reaction mechanism analysis	64
3.3.2	Transition state ensembles	66
3.3.3	Dehydrogenation barriers	68
4	The fate of nitrogen on Fe(110)	73
4.1	Interaction of nitrogen and hydrogen with Fe(110)	74
4.2	Nitrogen recombination	79
4.3	Nitrogen migration inside the bulk	81
4.4	Lateral interaction effects	83
4.5	Discussion and experimental comparison	85
5	Fe(111) under high N^* coverage	88
5.1	Active site for N_2 dissociation	89
5.2	Effect of N^* coverage on surface morphology	90
5.2.1	Triangular motives validation with DFT	92
5.2.2	$N^* - N^*$ pair interaction	93
5.2.3	Effective N^* coverage	97
5.2.4	Dynamical properties	98
5.3	Effect of N^* coverage on N_2 adsorption and decomposition	100
5.3.1	Reaction pathway analysis	104
6	A Data-Efficient protocol for modeling catalytic reactivity	107
6.1	A data-efficient protocol	108
6.1.1	Preliminary construction of reactant potentials	110
6.1.2	Reactive pathways discovery via uncertainty-aware flooding simulations	111
6.1.3	Uniform accuracy along reactive paths through GNNs and Data-Efficient Active Learning (DEAL)	115
6.2	Validation of the workflow	121
6.2.1	Data-efficiency	125
6.3	DEAL: general features and applicability	127
7	Ammonia cracking on iron-cobalt	131
7.1	Unraveling the free energy landscape and reaction mechanism	131
7.1.1	Dehydrogenation steps	132
7.1.2	Nitrogen adsorption and recombination	137
7.1.3	Comparison of elementary step with Fe	139
7.2	Nitrogen migration inside the bulk	141
7.3	Lateral interactions	143
7.3.1	Efficiently refining the ML potential with DEAL	143

7.3.2 Nitrogen recombination and dissolution	147
7.4 Experimental observations of no nitride formation	148
Conclusions and Perspectives	151
List of Publications	153
Bibliography	154
Acknowledgement	170

List of abbreviations

ACE	Atomic Cluster Expansion
ACSF	Atom-Centered Symmetry Functions
AIMD	<i>Ab-Initio</i> Molecular Dynamics
BCC	Body-Centered Cubic
CV	Collective Variables
DFT	Density Functional Theory
EAM	Embedded Atom Model
FCC	Face-Centered Cubic
FES	Free Energy Surface
GGA	Generalized Gradient Approximation
GNN	Graph Neural Network
GP	Gaussian Process
LDA	Local Density Approximation
MD	Molecular Dynamics
ML	Machine Learning
MLP	Machine Learning interatomic Potentials
NN	Neural Network
OPES	On-the-fly Probability Enhanced Sampling
PBE	Perdew–Burke–Ernzerhof functional
PES	Potential Energy Surface
QM	Quantum Mechanics
TAP	Temporal Analysis of Products
TS	Transition State
UHV	Ultra-High Vacuum
XAS	X-ray Adsorption Spectroscopy
XPS	X-ray Photoelectron Spectroscopy
XRD	X-Ray Diffraction spectroscopy

Outline

The global push towards a greener economy has amplified the urgency to identify sustainable solutions for energy storage and transportation. One of the promising approaches to enable this transition is through the use of hydrogen as a clean fuel. However, the challenge lies in finding safe and efficient ways to store and transport hydrogen. Ammonia (NH_3) has emerged as a potential carrier due to its high hydrogen density and ease of liquefaction at moderate conditions. Nevertheless, the efficient decomposition of ammonia to release hydrogen, especially under industrially relevant conditions, remains an area requiring further research.

Ammonia decomposition on iron-based catalysts has been a topic of considerable interest for several decades, as iron offers an attractive balance between cost and catalytic efficiency. However, the atomic-level mechanisms governing the ammonia decomposition process, particularly under high-temperature *operando* conditions, have not been fully elucidated. Traditional models, which rely on static representations of catalytic surfaces, fail to capture the dynamic nature of surface reactions, which are instead crucial in explaining catalyst long-term stability.

This thesis aims to address these challenges by using advanced computational techniques, including molecular dynamics (MD) simulations, machine learning interatomic potentials (MLPs), and enhanced sampling methods. These techniques allow for a dynamic and atomistic view of the catalytic processes occurring on iron (Fe) and iron-cobalt alloy (FeCo) surfaces during ammonia decomposition. An important focus of this work is to uncover the fate of nitrogen on surfaces, a critical aspect of both catalyst performance and durability.

The findings from this research offer not only a deeper understanding of ammonia decomposition but also insights into the design of more efficient catalysts for ammonia cracking. More importantly, these studies paved the way for the practical application of MD simulations to study heterogeneous catalysis under operative conditions.

In the following chapters, the motivation of this study (Chapter 1), theoretical background, and computational methods (Chapter 2) employed will be presented.

Then, in Chapter 3, we will give a detailed analysis of the adsorption and de-

composition of an NH_3 molecule over iron surfaces. Here, the role of dynamics in the reaction is demonstrated, thus emphasizing the need for an *operando* temperature description. In the following Chapters, the description of the ammonia cracking reaction will be expanded, and the remaining reaction steps will be described, as well as the accounting for lateral interactions. Moreover, we will focus on nitrogen's competitive processes (Chapter 4), which lead to the nitridation of iron catalysts. Chapter 5 also focuses on nitrogen effects by showing the impact of high nitrogen coverage and the early tendency of the Fe(111) surface to form nitrides. This results in a beneficial effect on N_2 dissociative chemisorption, a critical step in the Haber-Bosch synthesis.

Finally, in the last chapters, we will discuss the iron-cobalt alloy catalyst. Chapter 6 is devoted to presenting the methodological advancements required for the study of this complex system. In particular, we will present here a new protocol that allows for the building of reliable MLPs for *in operando* catalytic process in a data-efficient manner. Then, in Chapter 7, we will exploit the new method developed and the knowledge acquired on iron to simulate, for the first time, all the elementary steps of ammonia decomposition on FeCo catalysts. Moreover, nitrogen migration inside the bulk will be investigated, explaining the higher catalytic activity that has been experimentally observed.

Chapter 1

Introduction

Heterogeneous catalysis plays a vital role in many industrial processes, including the production of fuels, chemicals, and fertilizers. In this form of catalysis, reactions occur at the interface between different phases, typically a solid catalyst and gaseous or liquid reactants. The properties of the catalyst surface largely determine the efficiency and selectivity of these reactions.

However, studying heterogeneous catalysis under realistic conditions presents significant challenges. Industrial reactions often occur under extreme thermodynamic conditions, such as high temperatures and pressures, leading to complex and dynamic environments on the catalyst surface. These surfaces can undergo phase transitions and structural changes and exhibit dynamic rearrangements of atoms during the reaction. This makes it difficult to capture the true nature of catalysis, as its activity is influenced by the continuous interaction between the surface and the reactants [1, 2]. The presence of multiple phases and fluctuating surface configurations further complicates both experimental and computational studies.

While recent experimental advancements have made it possible to study catalytic processes under *operando* conditions, where the catalyst is observed in real-time during the reaction, computational studies have often relied on static models. These models assume a frozen surface, neglecting the dynamic rearrangements that are crucial to understanding the reaction mechanisms in detail. This limitation creates a gap between theoretical predictions and the real behavior of catalysts under operating conditions.

In this Thesis, we aim to move beyond the static surface picture and develop simulations that incorporate the dynamics of catalysis in realistic environments. Using molecular dynamics (MD) simulations, we study the ammonia decomposition reaction, a key process in hydrogen production, under *operando*-like conditions. This chapter provides a brief overview of the fundamental properties of ammonia, its role

in the hydrogen economy, and the catalytic processes involved in its decomposition. In the final section, we will discuss the importance of incorporating dynamics into computational models of heterogeneous catalysis, explaining the significant efforts made to push the frontiers of computational investigation toward a fully dynamic description.

1.1 Ammonia as hydrogen vector

Haber and Bosch’s discoveries more than a century ago paved the way for the industrial synthesis of ammonia, which has been instrumental in sustaining population growth through large-scale production of fertilizer [3]. In the Haber–Bosch process, ammonia (NH_3) is produced by molecular nitrogen (N_2) and hydrogen (H_2) using a iron metal catalyst:



This reaction is exothermic ($\Delta H = -46.1 \text{ kJ/mol}$) [4] and disfavored in terms of entropy because four equivalents of reactant gases are converted into two equivalents of product gas. As a result, high pressures (140–320 atm) and relatively low temperatures favor ammonia synthesis at equilibrium conditions. Still, temperatures need to be high enough (350–550 °C) to overcome the relatively high activation energy of the elementary steps, namely the cleavage of the triple bond of the nitrogen molecule followed by three hydrogenation steps [5]. The first is believed to be the rate-limiting one [6, 7]. Currently, more than 180 million tons of ammonia per year are produced via the Haber-Bosch process [8]. It has been estimated that this accounts for roughly 2% of fossil fuel use worldwide and releases more than 450 million tons of CO_2 , representing close to the 1% of total annual global emissions of the greenhouse gas [8–10]. For this reason, commercial ammonia is not called *green* as it encompasses the steam methane (CH_4) reforming process for hydrogen production followed by the Haber-Bosch [10]. The former process accounts for 90% of the CO_2 emission [10] and is the process widely adopted for H_2 production [11]. According to the road map for the ammonia economy, in the foreseeable future, the Haber–Bosch process is still the dominant ammonia synthesis technology. However, other emerging technologies such as electrochemical and photocatalytic synthesis of ammonia are also promising for the production of so-called *green* ammonia [8, 12, 13].

At the same time, ammonia has high potential roles in a hydrogen economy [14]. With the impending arrival of the energy transition, hydrogen technologies are gaining importance, and the generation of hydrogen from ammonia decomposition is an attractive pathway to produce highly pure and carbon-free hydrogen on-site

and on-demand [15]. Ammonia has a significant weight fraction of hydrogen (17.65% of its mass) and offers a significant advantage over molecular hydrogen since, at atmospheric pressure, its liquefaction temperature is as high as $-33\text{ }^{\circ}\text{C}$, making it much easier to store and transport [16, 17]. However, the development of efficient methodologies for extracting hydrogen from ammonia is needed for this process to be economically viable.

1.2 Thermocatalytic decomposition

Among the different techniques for generating hydrogen from ammonia, here we focus on the most employed one, namely thermal decomposition (or catalytic cracking). Thermocatalytic decomposition of ammonia is the reverse of the Haber-Bosch ammonia synthesis [18]. Compared to the synthesis process, the decomposition reaction is favored by higher temperatures and lower pressures, e.g., the operational temperatures for catalytic cracking on iron at 1 atm lie in the range of $400\text{-}750\text{ }^{\circ}\text{C}$ [15].

A variety of catalysts have been tested in this context, and, having in mind industrial applications, most of the attention has been focused on non-noble metals like Fe, Co, and Ni and their alloys [18]. On metals, ammonia decomposition is believed to proceed through a series of dehydrogenation steps followed by recombination of nitrogen and hydrogen [19]. On Fe, Co, and Ni, the last step is believed to be the rate-limiting one, while dehydrogenation becomes limiting over other metals such as Ru, Ir, Pd, Pt, Cu [20].

Even if Ru shows the best catalytic activity among the one component catalysts [21], non-noble metals, mainly Fe, Co, and Ni, are more sustainable and thus are better candidates as industrial catalysts [18, 22]

1.2.1 Iron catalysts

Iron is a prototypical thermocatalyst in both the Haber-Bosch process and its inverse. The catalyst can be derived from the iron oxides magnetite (Fe_3O_4), wustite (Fe_{1-x}O), or hematite (Fe_2O_3). Significant advancements have been achieved since the pioneering work of A. Mittasch [23], who developed the original Fe-based NH_3 synthesis catalyst from Swedish iron ore containing magnetite and several impurities. Recent progress has been summarized by Humphreys *et al.* [3], which mainly concerns multiply promoted wustite-based catalysts. During the synthesis of NH_3 , the formation of surface nitrides plays an important role, which was observed after the exposure of Fe single crystals to N_2 in ultra-high vacuum (UHV) experiments

in the seminal studies by G. Ertl and co-workers [24], which obtained activation energies of N₂ desorption for Fe(100) and Fe(111) of 243 kJ/mol and 213 kJ/mol, respectively. Later, they also reported the activation energy for N₂ desorption from the Fe(110) surface amounting to 238 kJ/mol [25]. In a recent work [26], B. Lu *et al.* reviewed recent experiments on the role of nitride phase during NH₃ decomposition, highlighting that the reported active phases are quite divergent with even a not necessarily negative role of catalyst nitridation. However, a poisoning effect due to bulk nitridation of the catalyst is well-established [18, 27]

Still, in spite of more than a century of efforts, a comprehensive characterization of the Haber-Bosch and its reverse process remains elusive due to extreme *operando* conditions, which make both experiments and simulations challenging. Indeed, the operational temperatures for catalytic cracking on iron at 1 atm lie in the range of 400-750 °C [15]. Experimental information about the reaction intermediates has thus been obtained in conditions far from the *operando* regimes [28, 29], or extrapolated through indirect temporal analysis of products (TAP) measurements [30]. These reactions are known to be highly structure-sensitive, with Fe(110) and Fe(111) representing two limiting cases among the low-index surfaces. In fact, the former is characterized by high density and stability but low catalytic activity, while the latter has an open structure and high activity [24, 25, 28, 29, 31, 32]. This sensitivity has been reported both for the decomposition reaction [28, 29], similarly to what was observed in the synthesis process [32, 33]. Only very recently has direct *operando* probing of synthesis reaction intermediates been performed [34], confirming the structure sensitivity of the reactions and showing non-trivial temperature dependency in the rate-limiting step.

From a computational perspective, there is a large number of theoretical studies on ammonia decomposition based on T=0 K density functional theory (DFT) calculations [26, 30, 35–43]. However, the results are obtained in conditions far from the operative ones, and the extrapolated activation energies are highly dependent on the frozen model property and the different types of correction applied.

In a recent work [44] on N₂ dissociative chemisorption on Fe(111), L. Bonati *et al.* demonstrate the need for a dynamic description of the reaction, showing that the behavior of active sites cannot be inferred from the static calculation. In this Thesis, we followed these results and conducted an extensive study on ammonia decomposition with a full dynamic description. Moreover, we will extend Bonati's work [44] in order to account for lateral interactions and coverage effects.

1.2.2 Bimetallic alloy catalysts

In recent years, bimetallic alloys, especially non-noble ones, have emerged as promising candidates for ammonia decomposition [30, 45–51]. These alloys are suggested to offer optimal binding energies for the reaction intermediates, enhancing their catalytic performance [20]. However, a simple explanation in terms of volcano shape [19, 52] is not always possible. In fact, Nørskov *et al.* proposed an activity volcano for ammonia decomposition based on dissociative N_2 adsorption energy [21]. They suggested that the optimal binding energy (the Ru one) can be obtained by combining weak and strong binding elements, and they proposed CoMo as an optimal combination. However, during ammonia decomposition, the alloy forms bulk nitrides ($\text{Co}_3\text{Mo}_3\text{N}$) [53].

Here, we focus on alloying iron with cobalt, which has been reported to have a superior catalytic activity for ammonia decomposition compared to the monometallic iron activity [30, 49, 51]. However, an atomistic understanding is still lacking. In Ref. [30], while the experiments were performed on Fe-Co alloys, the computational analysis was limited to pure Fe and Co systems. The authors have attributed the improved activity of FeCo alloy to the presence of Co and suggested that the reaction mechanism on FeCo would likely resemble an ideal average between the monometallic catalysts. In particular, they suggest that the free energy barriers would likely resemble the Co ones, but with a reduction of the hydrogen formation barrier, considered limiting on monometallic Co. On the other side, S. Chen *et al.* [51] reported different Fe-Co alloys supported on MgO, with the optimal catalytic activity observed for a Fe-Co ratio of 1:1. Based on the experimental data (XRD and *operando* XAS measurements), the authors argued that the active phase of the catalyst is the metallic Fe-Co surface, and particularly the (110) surface. These results suggest FeCo alloys remain stable during ammonia decomposition and do not lead to nitride formation, contrary to what is reported for iron. The higher activity of the alloy was attributed again to the different binding energies of nitrogen to FeCo compared to pure Fe, calculated at the DFT level. However, they do not report any information on the mechanism. In summary, there is experimental evidence of improved catalytic activity of FeCo alloy for ammonia decomposition [30, 51]. Still, no theoretical modeling of the reactions has been attempted because of the difficulties involved.

The lack of microscopic interpretation motivated us to study the operation of FeCo alloys under operating conditions. Because of the high computational cost of metallic and ferromagnetic calculations, we developed a new data-efficient scheme (see Chapter 6) and used it to simulate the mechanism and reconstruct the free energy profiles of the NH_3 decomposition (Chapter 7), unrevealing the difference

with iron catalyst and thus rationalizing the promotional effect of Co.

1.3 Dynamical modelling of heterogeneous catalysis

Traditional studies of heterogeneous catalysis often focus on static pictures of catalyst surfaces and reaction intermediates. However, recent advancements in both experimental and computational techniques have highlighted the importance of dynamics in catalytic processes. To be more precise, recent technological advancements have led to the spread of recognition of its role, but the idea that dynamics is crucial comes from the past. More than 40 years ago, Spencer offered a picture of heterogeneous catalysis that is far from static and requires establishing a dynamic steady state [54]. More recently, Schlögl has emphasized the dynamic nature of catalytic surfaces under operando conditions, where surface atoms and adsorbates are constantly in motion, and redefined heterogeneous catalyst as “a functional material that continually creates active sites with its reactants under reaction conditions” [1].

This crucial role is clearly emerging in recent simulations [44, 55–65], thanks to advancements in computational tools. In particular, the effects of dynamics under operating conditions can manifest themselves in many different ways, for example, surface dynamics and reaction-induced changes can be found in the literature, with notable examples being the activation of the copper surface by carbon monoxide [56] and the destabilization of lithium imide [64, 66] and barium hydride [65] surfaces under ammonia decomposition. Focusing on ammonia over metal surfaces, Gerrits *et al.* [59] have highlighted the role of dynamics on the first dehydrogenation step Ru(0001) surface through *ab initio* molecular dynamics (AIMD) simulations [59]. Moreover, Bonati *et al.* revealed the dynamic change of the Fe(111) surface, with the active sites for the N₂ dissociation acquiring a finite lifetime, getting continuously formed and broken during surface dynamics [44].

In this Thesis, we propose a fully dynamic description of ammonia decomposition on Fe and FeCo. To achieve this, we perform large-scale MD simulations on long timescales. These simulations are made possible by several methodological advancements, particularly the development of machine learning potentials (MLPs) and enhanced sampling techniques. Furthermore, we have developed new strategies to construct accurate MLPs in a data-efficient manner, which is crucial for studying the bimetallic alloy. To study the highly dynamic environment, we employ a range of advanced analysis tools, such as ML-based partial charge predictions and transition state (TS) analysis. These methods will be reviewed in detail in the following chapters.

Chapter 2

Computational Methods

The driving idea of this work is to move beyond a static description of heterogeneous catalysis using molecular dynamics and advanced simulation techniques. The possibility of using molecular dynamics to study complex catalytic processes such as ammonia decomposition or the Haber-Bosch process lies in the recent progress in the field, particularly by extensive exploitation of machine learning, supported by the increasing availability of computational resources. This extended chapter will present the computational methods employed to simulate catalytic processes at operating temperatures. The first sections are dedicated to molecular dynamics, machine learning-based interatomic potentials, and advanced sampling techniques. Then, we will introduce the analysis techniques for retrieving information within the complex simulation environment. Of particular importance here is the ability to track atomic charges during large-scale simulations. At the end of the chapter, we provide all the computational details necessary to reproduce the results of this thesis.

2.1 Molecular Dynamics

Molecular dynamics (MD) is a computational technique that allows us to simulate the temporal evolution of a system of atoms. It is the atomistic simulation technique that most closely resembles experimental techniques, as it can act as a microscope with both atomic spatial resolution and potentially infinite temporal resolution. The equilibrium and dynamic properties of the system can be calculated as ensemble averages once the initial conditions and the interatomic potential between the particles are known.

In classical MD, the positions and velocities of each atom at every moment in time after the initial time t_0 are calculated by numerically integrating the classical equations of motion. A fundamental approximation, considered valid in most molecular

dynamics schemes, is the adiabatic Born-Oppenheimer approximation. This approximation allows the decoupling of the motion of nuclei and electrons, *i.e.*, separating the variables corresponding to nuclear motion and the electronic coordinates in the Schrödinger equation associated with the system's Hamiltonian. The adiabatic Born-Oppenheimer approximation is based on the fact that typical electronic velocities are orders of magnitude higher than nuclear velocities. It also assumes that the lighter electrons adjust adiabatically to the motion of the heavier nuclei, remaining at any time in their instantaneous ground state. This approximation has been demonstrated to be accurate even for metallic systems [67–69].

The reliability of the simulations is primarily related to the accuracy of the calculation of forces acting between atoms. The different methodologies for calculating these forces will be discussed later.

Once the acting forces are defined, the temporal evolution of each atom is determined by Newton's equation of motion:

$$M_i \ddot{\mathbf{r}}_i = \mathbf{F}_i \quad (2.1)$$

where M_i , $\ddot{\mathbf{r}}_i$, and \mathbf{F}_i are the mass, acceleration, and force acting on the i -th nucleus, respectively. The solution of equation 2.1 is obtained through a finite difference numerical algorithm. Time is discretized in time steps Δt of the order of femtoseconds, and the system moves from a certain configuration at time t to a new configuration at $t + \Delta t$. The integration algorithm used in this thesis is the Velocity Verlet (VV) [70, 71]. In this scheme, the position \mathbf{r}_i of the i -th atom at time $t + \Delta t$ is calculated from the position \mathbf{r}_i , velocity $\mathbf{v}_i = \dot{\mathbf{r}}_i$, and acceleration $\mathbf{a}_i = \ddot{\mathbf{r}}_i = \mathbf{F}_i(\{\mathbf{r}_i\})/M_i$ at time t . Once the new ionic configuration $\{\mathbf{r}_i\}$ at time $t + \Delta t$ is defined, the forces, and thus the acceleration, can be obtained at this timeframe. The new velocities at time $t + \Delta t$ are calculated by considering the acceleration at both times t and $t + \Delta t$. The VV algorithm ensures an error less than Δt^4 for the positions \mathbf{r}_i and less than Δt^3 for the velocities \mathbf{v}_i .

$$\begin{aligned} \mathbf{r}_i(t + \Delta t) &= \mathbf{r}_i(t) + \mathbf{v}_i(t) \Delta t + \frac{1}{2} \mathbf{a}_i(t) \Delta t^2 + \mathcal{O}(\Delta t^4), \\ \mathbf{v}_i(t + \Delta t) &= \mathbf{v}_i(t) + \frac{\mathbf{a}_i(t) + \mathbf{a}_i(t + \Delta t)}{2} \Delta t + \mathcal{O}(\Delta t^3). \end{aligned} \quad (2.2)$$

The algorithm is time-reversible and generates trajectories in the microcanonical ensemble (NVE), where the total energy (kinetic energy plus potential energy) is conserved within numerical integration errors. The configurations visited during the dynamics, while deviating from the real physical trajectory, form a collection of states (q, p) consistent with the microcanonical probability distribution $\rho(p, q)$, where q

and p represent positions and momenta following the standard notation of statistical mechanics. Under this assumption, known as the ergodic hypothesis, which is valid at dynamic equilibrium, the ergodic theorem holds:

$$A_{obs} = \langle A \rangle = \int dqdp \rho(q, p) A(q, p) = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \int_0^\tau dt A(\{\mathbf{r}_i(t)\}, \{\mathbf{v}_i(t)\}) \quad (2.3)$$

which ensures that the value of a generic observable A_{obs} which is an ensemble average $\langle A \rangle$ is equivalent to the time average over the trajectories. In practice, we can approximate the time integral average to a discrete average over samples:

$$A_{obs} \approx \frac{1}{n_k} \sum_{k=1}^{n_k} A_k \quad (2.4)$$

where n_k is the number of samples of the trajectory, and A_k is the value of the variable A at the frame k [72]. Therefore, an average over MD trajectories provides a measure of the experimentally observable macroscopic quantity A_{obs} .

To control the system temperature, it can be coupled to a thermal bath that ensures the average temperature is close to the desired temperature, T_{ext} . When this is done, the equations of motion are modified, and the system no longer samples the microcanonical ensemble. At equilibrium, the trajectories instead sample the canonical ensemble (NVT). In this way, MD simulations can further mimic experimental conditions.

Fixing the temperature alone does not ensure that a simulation samples the intended canonical ensemble. For example, rescaling the particle velocities at every step results in a constant kinetic temperature, T_{ext} , as calculated by the equipartition theorem. However, this contradicts the expected fluctuations in T_{ext} predicted by the Maxwell-Boltzmann distribution. A good thermostat should not only control the temperature but also conserve certain quantities to confirm that the correct ensemble is being sampled and to ensure ergodic behavior. In the '80, Andersen [73] introduced the first thermostat, based on a stochastic process where particles randomly collide with a thermal bath, altering their velocities. Later, Berendsen et al. [74] proposed an algorithm to rescale the velocities to bring the system temperature to T_{ext} . However, this thermostat suppresses kinetic energy fluctuations and, therefore, does not guarantee proper canonical ensemble sampling. Although both methods are simple and relatively stable, they have drawbacks and should be avoided for production runs.

More advanced thermostats have been proposed to correctly generate trajectories consistent with a canonical ensemble, *e.g.*, the widely known Nosé–Hoover one [75,

76] and its advancement [77]. Here, we adopt the Bussi–Donadio–Parrinello thermostat [78] that could be considered an extension of the Berendsen one. In this scheme, the temperature is controlled through velocity rescaling but in a stochastic approach, which allows the kinetic energy to fluctuate naturally according to the Maxwell-Boltzmann distribution. This ensures both accurate temperature control and proper sampling of the ensemble without artificially fixing the kinetic energy, thus fulfilling the ergodic condition. The algorithm is also computationally efficient and stable over long simulations. The details of MD simulations are reported in Section 2.7

2.1.1 Interatomic Force Description

As previously mentioned, in classical MD, a key aspect is the description of the forces between atoms, which are described by the so-called interatomic potentials. In this section, we will explore various methods for calculating these forces, starting from the most accurate, *ab initio* molecular dynamics (AIMD), and progressively introducing more approximate models.

2.1.1.1 Ab Initio Molecular Dynamics (AIMD)

In AIMD, the term “*ab initio*” indicates that the forces are computed from first principles, without relying on empirical or fitted potentials. This is a powerful simulation technique that directly incorporates quantum mechanical (QM) calculations of electronic structure to describe atomic interactions. However, solving the Schrödinger equation for many-body interactive electronics is impossible, and approximations are required. In modern AIMD, the electronic structure method most commonly used is the Kohn–Sham formulation of density functional theory (DFT).

Density Functional Theory (DFT) DFT is the cornerstone of AIMD, as it allows for the calculation of electronic energies and forces with a favorable balance between accuracy and computational cost. The fundamental idea behind DFT is that the ground-state properties of a many-electron system can be described by a functional of the electron density, $\rho(\mathbf{r})$, instead of the much more complex many-body wavefunction. This simplification drastically reduces the complexity of the quantum mechanical problem. The central equation of DFT is the Kohn–Sham equation, which maps the many-electron problem onto an auxiliary system of non-interacting electrons moving in an effective potential, $V_{\text{eff}}(\mathbf{r})$. This potential includes contributions from the external potential (*e.g.*, nuclei), electron–electron Coulomb repulsion, and the exchange–correlation functional, $E_{\text{xc}}[\rho]$, which encapsulates all the many-body

quantum effects:

$$\left[-\frac{\hbar^2}{2m}\nabla^2 + V_{\text{eff}}(\mathbf{r}) \right] \psi_i(\mathbf{r}) = \epsilon_i \psi_i(\mathbf{r}) \quad (2.5)$$

Here, $\psi_i(\mathbf{r})$ are the Kohn-Sham orbitals, and ϵ_i are the corresponding eigenvalues. The electron density is then obtained as:

$$\rho(\mathbf{r}) = \sum_i |\psi_i(\mathbf{r})|^2 \quad (2.6)$$

where i runs over the occupied states. The total energy of the system is expressed as:

$$E_{\text{tot}}[\rho] = T_s[\rho] + E_{\text{ext}}[\rho] + E_{\text{H}}[\rho] + E_{\text{xc}}[\rho] \quad (2.7)$$

where:

- $T_s[\rho]$ is the kinetic energy of the non-interacting electrons,
- $E_{\text{ext}}[\rho]$ is the external potential energy (nuclei-electron interaction),
- $E_{\text{H}}[\rho]$ is the Hartree energy, representing the electron-electron Coulomb repulsion,
- $E_{\text{xc}}[\rho]$ is the exchange-correlation energy, which contains the many-body effects of exchange and correlation.

Solving the Kohn-Sham equations self-consistently yields the ground-state electron density and the corresponding forces on the nuclei. These forces are then used to update the positions of the nuclei during the molecular dynamics simulation, following Newton's equations of motion (Eq. 2.1). The accuracy of DFT crucially depends on the choice of the exchange-correlation functional, $E_{\text{xc}}[\rho]$. Since the exact form of this functional is unknown, several approximations are commonly used. The simpler approximation is the Local Density Approximation (LDA), where $E_{\text{xc}}[\rho]$ is approximated by assuming that the electron density locally behaves like a uniform electron gas. LDA works reasonably well for systems with slowly varying densities, such as simple metals, but can be less accurate for systems with rapidly varying densities. Another widely used approximation is the Generalized Gradient Approximation (GGA), which improves upon LDA by including not only the local value of the density but also its gradient, $\nabla\rho(\mathbf{r})$. This allows for better treatment of inhomogeneous systems, such as molecules and surfaces. In this thesis, when QM calculations are required to be performed, we will adopt this level of theory, in particular, using the Perdew-Burke-Ernzerhof (PBE) [79] functional. All the details are reported in Section 2.7. These two approximations are widely used, but they have limitations, particularly in describing systems with strong electronic correlations or excited states. To address these, more advanced methods have been developed, such as hybrid function-

als (e.g., B3LYP [80] or PBE0 [81]) or meta-GGAs (e.g., SCAN [82]). Furthermore, other corrections such as DFT+U (for strongly localized electrons) [83–85], many-body perturbation theory, time-dependent DFT [86] (to deal with excited states), or DFT-D [87, 88] (to include dispersion correction) have been proposed. However, the applicability criteria of these corrections are highly debated and not suitable for ferromagnetic metallic surfaces [89, 90]. Moreover, when going up to the so-called “*Jacob’s ladder*” of *density functional approximations*[91], the computational cost of these methods increases, becoming prohibitive for the systems we studied.

AIMD is particularly powerful for simulating systems where the electronic structure is critical to the dynamics, such as chemical reactions, phase transitions, catalytic processes, and materials under extreme conditions. Its ability to account for electronic effects on the fly makes it indispensable in scenarios where traditional force fields fail, particularly in cases involving bond-breaking and bond-forming events. However, due to its high computational cost, AIMD is typically limited to relatively small systems (up to a few hundred atoms) and short simulation timescales (hundreds of picoseconds).

2.1.1.2 Semi-Empirical Methods

Semi-empirical methods offer a faster and computationally cheaper alternative to fully *ab initio* approaches, making them an attractive choice when large systems or long timescales are involved. While these methods are still grounded in QM principles, they achieve significant speed-ups by simplifying the calculations. This is done by neglecting certain integrals that are computationally expensive and considered less critical for describing the system’s overall behavior. To compensate for these approximations, empirically derived parameters are introduced, which are optimized using experimental data or high-level *ab initio* results. This makes this method not universally applicable since they are highly dependent on how well the empirical parameters represent the system under study. Despite being slightly different from the majority of semi-empirical methods based on molecular orbit theory [92], the Tight Binding approach [93] can also be considered in this category. In summary, semi-empirical methods provide a practical middle ground between fully *ab initio* approaches and purely empirical models. Yet, they are generally inadequate for describing systems such as ferromagnetic metals and, at most, their interaction with small molecules, as in the case under study.

2.1.1.3 Empirical Force Fields

Moving further up the approximation scale, empirical force fields describe the forces between atoms using pre-defined analytical expressions. These models are based on fitting experimental data or high-level quantum calculations and do not involve solving the electronic structure problem directly. Instead, they represent the potential energy of the system as a function of atomic positions using relatively simple functional forms. The most simple form of force fields are the pair potentials, where the potential energy is described only in terms of the distances between couples of atoms. Some well-known examples of empirical force fields include the Lennard-Jones potential and the Morse potential. The complexity increases when angular (three-body) terms are included in the potential, as in the case of the Stillinger-Weber potential [94]. Further complexity arises with many-body potentials, which account for interactions involving multiple atoms. A notable example for metals (though not applicable to ferromagnetic ones) is the Embedded Atom Model (EAM) potentials [95]. Nowadays, there are hundreds of complex empirical force fields available, such as CHARMM [96], AMBER [97], and OPLS [98], to name just a few. These are widely used in biological and material science simulations. These methods allow simulations of very large systems (millions of atoms) and long timescales (microseconds or longer), but they often sacrifice chemical accuracy, especially in scenarios where bond formation or breaking occurs. To address such events, several modifications have been proposed, including Reactive Force Fields (reaxFF) [99]. However, the number of parameters in these models can grow excessively with the complexity of the systems, and in many cases, DFT remains a more practical solution. The main drawback of these empirical force fields is their poor transferability, which means that the reliability of simulations decreases significantly when applied under conditions different from those for which the potential was originally constructed. However, it is important to acknowledge that these potentials, being optimized to reproduce experimental data, can be even superior to DFT in capturing thermodynamic properties.

2.2 Machine Learning Interatomic Potentials

Machine Learning Potentials (MLPs) provide a computationally efficient means of representing the potential energy surface (PES). These potentials rely on machine learning models trained on data obtained from high-level QM calculations, allowing for the prediction of atomic interactions with near-quantum accuracy but at a fraction of the computational cost. MLPs are especially useful for simulating large systems or long timescales, where traditional *ab initio* methods would be prohibitively

expensive. In MLPs, the PES is represented by a machine learning model, such as a Neural Network (NN) or Gaussian process (GP), that maps atomic configurations to energies and forces. These models are trained on a dataset composed of atomic configurations (energies and forces), typically generated using DFT or other high-level QM methods. The goal is to create a model that can accurately interpolate the PES for new atomic configurations not included in the training set.

The idea of representing the PES through neural networks (NN) originated in 1995 with the work of Blank *et al.* [100], where feed-forward NNs were used to define relationships between the system’s structure and the desired properties. In these networks, the energy E is calculated by the NN from the atomic positions. However, this approach limits applications to systems with fixed dimensionality.

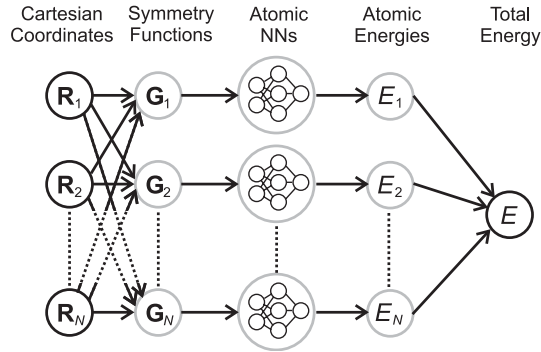


Figure 2.1: Schematic of a high-dimensional neural network according to the method proposed by Behler and Parrinello [101].

A method to generate NN potentials for high-dimensional systems containing thousands of atoms was proposed by Behler and Parrinello in 2007 [101]. In this approach, the conceptual limitations of traditional NN potentials are overcome by using a feed-forward NN for each atom in the system (Figure 2.1). Each “atomic NN” provides the contribution to the energy E_i of atom i as a function of its surrounding chemical environment, and the total energy E of the system is obtained as the sum over the total number of atoms N :

$$E = \sum_{i=1}^N E_i = \sum_{\mu=1}^{N_{elem}} \sum_{i=1}^{N_{\mu}} E_i^{\mu} \quad (2.8)$$

where N_{elem} is the number of chemical species, and N_{μ} is the number of atoms of the μ -th species. It must hold that $\sum_{\mu=1}^{N_{elem}} N_{\mu} = N$. For a given chemical element, the architecture of the atomic NN is fixed, meaning that the number of hidden layers, nodes, and their corresponding weights are all predetermined. The input vector \mathbf{G}_i of each atomic NN describes the local environment of the atom up to a certain cut-




Figure 2.2: Graphical representation of the three basic building blocks characterizing recent high-dimensional MLPs. The snapshot in the dataset block, as all the snapshots reported in the Thesis, was obtained with Ovito visualization software [102].

off distance. To ensure that the number of inputs (i.e., the dimension of the vector \mathbf{G}_i) is independent of the number of neighbors of each atom, the local environment is not described in Cartesian coordinates but through special many-body functions called “symmetry functions”. These functions contain all the information regarding the radial and angular distribution of the neighbors around the atom on which they are calculated. These functions, which depend on the Cartesian coordinates of each atom within the cut-off distance, are computed for every atom i and form the input vector \mathbf{G}_i of the atomic NN.

This description of the local environment through symmetry functions provides two fundamental advantages: first, it analytically introduces invariance of the NN potential with respect to particle exchange and roto-translational symmetry. Secondly, it eliminates the dependence of atomic NNs on the number of neighbors, thus allowing for the treatment of systems with arbitrary and non-fixed sizes.

This seminal work lays the foundation for all modern MLPs, particularly highlighting the three key ingredients needed to construct a machine learning potential: data, descriptors, and architectures (Figure 2.2).

2.2.1 Dataset

Data are the cornerstone of MLPs. The data constitute the training set, which must contain all the local environments that we want our potential to be capable of describing. The construction of the training set can follow very different paths. Generally, this is also the aspect that most interests a user of MLPs who, like in our case, does not develop the MLP but trains it to study a specific system.

The construction of the potential typically represents the most expensive aspect of a study, as all the configurations in the training set must be labeled (i.e., assigned energies, forces, and possibly other properties like stress, magnetic moment, or charge) using the highest possible level of theory (DFT or higher). Recently, the

dataset construction has been done progressively alongside the potential building process in the so-called active learning (see Section 2.3). However, most of the proposed methods are based on solid knowledge of the reaction under study or/and require too many QM calculations. Initially, we adopted a standard active learning scheme to collect the data for Fe reactions. The same scheme could not be applied to FeCo, where the cost of QM calculations turned out to be prohibitively expensive. Thus, we develop a new protocol for efficiently collecting structures for the training set, in particular with a novel data-efficient active learning (DEAL). We will discuss this aspect extensively in Chapter 6.

In this study, we constructed two training sets to investigate reactive chemical processes related to ammonia on Fe and FeCo surfaces. Both potentials were built using active learning processes. For Fe, we initially collected 110k configurations for the study of ammonia decomposition, then expanded the training set to about 220k configurations to account for lateral interactions. For FeCo, thanks to the new protocol developed here, we collected only 8K configurations to cover all reactions, including lateral interactions. The details of the training set composition are reported at the end of this chapter.

2.2.2 Descriptors

Once the dataset is assembled, the next critical step, which has a significant impact on the accuracy of the MLP, is the conversion of spatial configurations into an ML dataset through the use of descriptors. Descriptors are many-body functions of atomic coordinates $G=G(\mathbf{R}_i)$. As demonstrated with the Behler-Parrinello atom-centered symmetry functions (ACSF) [101], this approach allows for decoupling the dimensionality of the systems in the dataset from the target property we aim to predict. The choice of descriptors directly influences the highest level of accuracy the MLP can achieve, as machine learning algorithms are then employed to closely approximate the desired results (typically $E, \nabla E$). After defining the descriptors, essential tasks such as data cleaning, preprocessing, and normalization are required. As mentioned before, an effective descriptor includes the main physical property of the system, that is, spatial translation invariance (isometry of space), rotational invariance (isotropy of space), and permutation invariance (QM invariance). Furthermore, it must be continuous and unique to avoid degeneracy, and it is preferable to be parsimonious in order to minimize complexity and enhance computational efficiency.

With the expanding use of MLPs, dozens of types of descriptors have been proposed in recent years, along with their respective architectures. Here, we will briefly describe the two types of descriptors used in these studies, while directing interested

readers to specialized literature [103–106].

DeepPot-SE Descriptor The Deep Potential Smooth Edition (DeepPot-SE) descriptor [107] is a core component of the DeePMD framework, designed to map the atomic environment into a continuous and differentiable function that neural networks can use to predict the potential energy surface (PES). Rather than deciding a functional form of the descriptors, as in the case of Behler Parrinello ACSF, the descriptors themselves are learned using a neural network function. The atomic position $\mathbf{R} \in \mathbb{R}^{N \times 3}$ is mapped to an N atomic centered matrix of local coordinates $\mathbf{R}^i \in \mathbb{R}^{N_c \times 4}$ (where N_c is the maximum number of neighbors expected). The row j of \mathbf{R}^i can be written as:

$$\mathbf{R}_j^i = [s(r_{ij}), \hat{x}_{ij}, \hat{y}_{ij}, \hat{z}_{ij}] \quad (2.9)$$

where $\mathbf{r}_{ij} = \mathbf{r}_j - \mathbf{r}_i = (x_{ij}, y_{ij}, z_{ij})$ is the relative coordinate, $\hat{x}_{ij} = s(r_{ij}x_{ij})/r_{ij}$, and $s(r) \sim r^{-1}$ is a scalar and differentiable weighting function which goes to 0 at $r = r_c$. Additionally, DeepPot-SE descriptor includes an embedding matrix $\mathbf{J}^i \in \mathbb{R}^{N_c \times M}$, which consists of M nodes from the output layer of an NN function \mathcal{N}_e :

$$\mathbf{J}_j^i = \mathcal{N}_e(s(r_{ij})) \quad (2.10)$$

Then the embedding matrix \mathbf{J}^i , is decomposed in two matrix $\mathbf{J}^{i1} \in \mathbb{R}^{N_c \times M_1}$ and $\mathbf{J}^{i2} \in \mathbb{R}^{N_c \times M - M_1}$ by taking the first and last columns of \mathbf{J}^i respectively. The atomic descriptor $\mathbf{G}^i \in \mathbb{R}^{M_1 \times M_2}$ is constructed by taking a combination of \mathbf{J}^i and \mathbf{R}^i in order to maintain invariances by using symmetry-preserving transformations:

$$\mathbf{G}^i = (\mathbf{J}^{i1})^T \mathbf{R}^i (\mathbf{R}^i)^T \mathbf{J}^{i2}. \quad (2.11)$$

The use of the pairwise embedding neural network enables automatic learning of descriptors, providing flexibility and efficiency. By combining these properties, the DeepPot-SE descriptor achieves high accuracy in modeling complex systems, while also being computationally efficient, making it suitable for large-scale MD simulations.

ACE Descriptors The Atomic Cluster Expansion (ACE) [108] is a framework designed to efficiently and systematically describe the potential energy surface (PES) of atomic systems. Unlike other descriptors that rely on predefined functional forms, ACE provides a general and systematic way to represent atomic environments by considering many-body interactions through an expansion in terms of local basis functions. This approach is grounded in the idea that the PES can be expressed as

a sum of contributions from atomic clusters of increasing order, starting from single atoms, to pairs, triplets, and higher-order clusters.

Mathematically, the ACE descriptor represents the local environment of an atom i through a set of local basis functions $\phi_\alpha(\sigma_i, \sigma_j)$ that depend on the relative positions of neighboring atoms j , where σ_i is the vector of properties of atoms i , that is position, chemical species, and other learning parameters, and α a set of coefficients. ϕ_α can be generally written as the product of a radial and angular (as the spherical harmonics) and general function that encodes other properties like the chemical species. The first element of the expansion of atom i can be written as:

$$\mathcal{A}_{i,\alpha} = \sum_{j \in \mathcal{N}(i)} \phi_\alpha(\sigma_i, \sigma_j) \quad (2.12)$$

Where $\mathcal{N}(i)$ are the neighbors of atom i . The body ordered expansion of the site energy $E_i = V^{(1)}(\mathbf{r}_i) + \frac{1}{2!} \sum_j V^{(2)}(\mathbf{r}_i, \mathbf{r}_j) + \dots$ can be rewritten as:

$$E_i = E_i^{(0)} + \sum_{\alpha} \tilde{c}_\alpha^{(1)} \mathcal{A}_{i,\alpha} + \sum_{\alpha_1, \alpha_2} \tilde{c}_{\alpha_1 \alpha_1}^{(2)} \mathcal{A}_{i,\alpha_1} \mathcal{A}_{i,\alpha_2} + \dots \quad (2.13)$$

This representation needs to be further adjusted to include the rotational symmetry; we refer the reader to the original paper [108] for the complete mathematical description. Since the basis functions capture both radial and angular dependencies, this allows ACE to account for complex many-body interactions in a computationally efficient manner. The strength of ACE lies in its systematic nature, which allows for a controlled increase in accuracy by including higher-order terms in the expansion. Furthermore, ACE has shown excellent transferability, allowing it to generalize well across different atomic configurations, making it an optimal descriptor to be used in MLP.

2.2.3 Architectures

The architecture, also called the ML model or regressor, forms the backbone of an MLP. Its role is to represent the PES starting from the input features, which are the descriptors extracted from the dataset of atomic configurations. The effectiveness of the architecture in mapping these descriptors to potential energies and forces directly impacts the performance of the MLP. Various MLP models have been proposed for this purpose. Among the most prominent and widely used architectures in modern MLPs are artificial neural networks (NNs), graph neural networks (GNNs), and kernel-based methods such as Gaussian Processes (GPs).

2.2.3.1 Neural Networks

As previously mentioned, NNs are one of the most widely adopted machine learning architectures in the development of MLPs. These models consist of layers of interconnected nodes (neurons) that process the input features, or descriptors, and map them to the target properties, such as potential energy and atomic forces. The fundamental concept behind NNs is to approximate any nonlinear function by adjusting the weights and biases of the neurons through a process called training. They are indeed universal interpolators, that is, NNs are able to approximate any multi-input/output function

The output of a single neuron is given by applying an activation function f to a weighted sum of its inputs, \mathbf{x} :

$$y = f\left(\sum_i w_i x_i + b\right)$$

where w_i are the weights, x_i are the inputs, and b is a bias term. The higher the number of nodes (that are distributed in the hidden layers), the higher the flexibility of the network and its cost. Since f is a simple analytical nonlinear function, the propagation of the message across layers makes the NN capable of capturing arbitrary complex relationships in the data. The learning process of an NN is driven by minimizing a loss function, which measures the difference between the network's predicted outputs and the true values (e.g., energies or forces). One common loss function for MLPs is the Mean Squared Error (MSE), which, for a dataset with M training samples, is given by:

$$\mathcal{L} = \frac{1}{M} \sum_{i=1}^M \left(E_i^{pred} - E_i^{target}\right)^2$$

, where E_i^{pred} are the predicted values from the NN, and E_i^{target} are the target values in the training set one, typically from QM calculations. To minimize the loss function, the network adjusts its weights and biases using an optimization algorithm such as gradient descent. The gradients of the loss function with respect to the weights are computed via backpropagation, allowing efficient updates of the network parameters:

$$w_{i,j}^{new} = w_{i,j}^{old} - \eta \frac{\partial \mathcal{L}}{\partial w_{i,j}}$$

where η is the learning rate, controlling the step size during the optimization process.

A notable example of how NNs are used in MLPs is the DeepMD framework [109,

110], which utilizes neural networks in combination with the DeepPot-SE descriptors to efficiently represent the PES of large-scale systems. Today, DeePMD is one of the most popular MLPs. We have employed it to train the MLP for iron. The details are reported in Section 2.7.

2.2.3.2 Graph Neural Networks

Graph Neural Networks (GNNs) extend the flexibility and expressiveness of traditional neural networks by utilizing graph-based structures to represent atomic environments. Unlike conventional NNs, which process fixed-size input vectors, GNNs handle varying numbers of atoms and neighbors naturally by treating the system as a graph, where atoms are represented as nodes and interatomic interactions as edges.

GNNs capture local environments by passing messages along the edges, thus propagating information across neighboring atoms. The message-passing framework enhances the ability of GNNs to model atomic interactions by increasing the receptive field of the network, allowing it to incorporate both local and non-local interactions.

In this work, we employed MACE [111], one of the state-of-the-art GNN architectures, to study the ammonia decomposition process on FeCo surfaces. MACE excels in both in-domain and out-of-domain predictions, demonstrating remarkable data efficiency [112]. This architecture is built on a higher-order equivariant message-passing scheme related to the Atomic Cluster Expansion (ACE), which enables MACE to accurately model complex atomic interactions while maintaining computational efficiency. The specific parameters for the MACE network used in our simulations are detailed in Section 2.7.4.3.

2.2.3.3 Kernel-Based Methods

Kernel methods are widely used in machine learning for regression tasks, especially when the dataset is small but the problem requires flexibility in modeling complex relationships. The core idea of kernel methods is to project the input data into a higher-dimensional feature space, where the relationships between points become more linear. This is achieved through a kernel function, which measures the similarity between pairs of data points.

In kernel-based regression, we aim to predict the output y based on the similarity between the input \mathbf{x} and the training data points $\{\mathbf{x}_i\}$:

$$y(\mathbf{x}) = \sum_{i=1}^N \alpha_i K(\mathbf{x}, \mathbf{x}_i)$$

where $K(\mathbf{x}, \mathbf{x}_i)$ is the kernel function, and α_i are the learned coefficients from the

training process. A common choice for the kernel function K is the Gaussian Kernel:

$$K(\mathbf{x}, \mathbf{x}_i) = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}_i\|^2}{2\sigma^2}\right)$$

This kernel measures similarity based on the Euclidean distance between data points and is widely used due to its flexibility. Another common choice used is the polynomial kernel, which takes the dot product between the input vectors and raises it to the power p :

$$K(\mathbf{x}, \mathbf{x}_i) = (\mathbf{x} \cdot \mathbf{x}_i)^p$$

with the special case being the linear one with $p = 1$. Once the kernel function is chosen, kernel-based methods like Gaussian Process Regression (GPR) or Kernel Ridge Regression (KRR) fit the model by finding the optimal coefficients α_i that minimize a loss function, usually based on the difference between predicted and true outputs.

Starting with the Gaussian Approximation Potential (GAP) method [113], GPs have been successfully used for representing the potential energy surface of extended systems. In this work, we used the sparse GP implementation of FLARE (Fast Learning of Atomistic Rare Events) [114, 115], to study ammonia decomposition on FeCo surfaces. Similarly to the Behler-Parrinello high-dimensional NN, the energy is modeled as a sum over atomic contributions that depend on all atom-centered local environments of the system: $E = \sum_i \epsilon_i$. These environments are characterized by a set of invariant descriptors obtained as the product of equivariant ones [115] based on the atomic cluster expansion (ACE) [116]. What distinguishes GPs from other ML methods is their ability to provide confidence intervals on the predictions. Because of the energy decomposition, we can also measure the variance associated with each atomic contribution ϵ_i , obtaining a measure of the uncertainty of local environments, i.e., how different they are with respect to the training ones. Another key feature of GPs, is that they allow for the gradual inclusion of new atomic configurations into the training set, improving the model's accuracy without the need to retrain hyperparameters at every step. This makes kernel-based methods ideal for on-the-fly learning, particularly in situations where new environments are continuously encountered. However, a notable drawback is that the computational cost of making predictions increases as the size of the training set grows, which can become a limitation for large datasets. The specific parameters for the FLARE network used in our simulations are detailed in Section 2.7.4.2.

2.2.4 Uncertainty Quantification in MLPs

In machine learning potentials (MLPs), it is essential to estimate the uncertainty of predictions to ensure reliable simulations, particularly when exploring new atomic configurations not covered in the training set. In this section, we present the approaches we used to quantify uncertainty in both neural networks (NNs) and graph neural networks (GNNs), as well as Gaussian Processes (GPs). The ability to quantify uncertainty enables active learning schemes, where uncertain predictions can be used to query new data points from high-fidelity calculations, as we will discuss in Section 2.3 and Chapter 6.

While NNs (including GNNs) are highly expressive, they do not inherently provide uncertainty estimates. To address this, a common approach is to use an ensemble of models (referred to as a committee) trained on different partitions or permutations of the training dataset. Each model in the committee makes independent predictions, and the variability among these predictions serves as a proxy for the uncertainty of the model. Specifically, in this work we calculate the standard deviation of the predicted forces across the committee. For a force component α on atom i , the uncertainty $\sigma_{i,\alpha}$ is computed as:

$$\sigma_{i,\alpha} = \sqrt{\frac{1}{N_{\text{NN}}} \sum_{k=1}^{N_{\text{NN}}} \left(f_{i,\alpha}^{(k)} - \bar{f}_{i,\alpha} \right)^2}$$

where $f_{i,\alpha}^{(k)}$ is the force predicted by the k -th model in the committee, $\bar{f}_{i,\alpha}$ is the average force across all models, and N_{NN} is the number of models in the committee. To monitor the overall uncertainty of a configuration, we track the maximum uncertainty for each configuration:

$$\sigma_{\text{max}}^{(\text{NN})} = \max_{i,\alpha} \sigma_{i,\alpha} \quad (2.14)$$

This metric is used to identify configurations that should be re-labeled at the high level of theory, in a process known as query-by-committee selection [117]. In this work, we used this approach for both NNs and GNNs, to ensure the reliability of the simulations.

On the other hand, GPs naturally provide uncertainty estimates as part of their predictions. In GPs, the predicted variance for each atomic contribution reflects the model's confidence in its prediction. The energy of a system is expressed as a sum over atomic contributions, and for each atomic environment, the GP provides not only the predicted energy ϵ_i but also the variance $\tilde{\sigma}_{\epsilon_i}$, which represents the uncertainty of

the prediction. For each configuration, we monitor the maximum local uncertainty:

$$\sigma_{\max}^{(\text{GP})} = \max_i \tilde{\sigma}_{\epsilon_i} \quad (2.15)$$

In FLARE $\tilde{\sigma}_{\epsilon_i}$ is rescaled by the kernel hyperparameters to obtain a dimensionless quantity that ranges between 0 and 1 [115].

This feature allows GPs to perform uncertainty-aware MD. In this approach, the GP model continuously monitors its uncertainty during MD simulations. If the uncertainty of a prediction is low, the MD step is done using the GP forces otherwise, if it exceeds a predefined threshold, the system automatically switches to a higher-level QM calculation (e.g., DFT) to ensure accurate results. The GP model is then updated with the new data, improving its predictions as the simulation progresses. This active learning loop, where the model refines itself based on the predicted uncertainty, allows for efficient on-the-fly learning.[114].

In this thesis, we will adopt both the two schemes, but we will also combine them in a new protocol to build MLP in a highly data-efficient way (see Chapter 6)

2.3 Building machine learning potentials for chemical reactions

Simulating catalytic reactivity under operative conditions poses a significant challenge due to the dynamic nature of the catalysts and the high computational cost of electronic structure calculations. As we discussed in the previous Sections, MLPs [103] have emerged in recent years as promising tools to address the accuracy-efficiency trade-off. They are optimized to reproduce energies and forces from a dataset of reference calculations, typically performed at the DFT level. Hence, their effectiveness depends on the quality of the training dataset, which must include not only the equilibrium structures but also the relevant high-energy ones. In fact, as discussed in Section 2.2.1, the dataset is the basic building block of our MLP, and even with perfect descriptors and architecture, the quality of the ML prediction relies on the goodness of the dataset. This is especially crucial for transition state geometries, whose energies are connected to reaction rates through an exponential relationship (see Eq. 2.20). However, the identification of these structures in a complex and dynamic environment remains elusive, especially under operating conditions where an ensemble of TS configurations often exists [58, 118], as we will across this thesis results.

The training set construction typically involves active learning procedures [117, 119, 120], in which an ML model is trained on an initial dataset and used to generate

new structures (e.g., via molecular dynamics). A subset of these configurations is then labeled with single-point DFT calculations and added to the training set, proceeding iteratively until convergence. In the field of computational catalysis, such schemes have also been employed in combination with nudged elastic band [121] and minima hopping [122] to accelerate the calculation of energetic barriers and adsorbate geometries. However, we cannot rely on these static approaches to simulate the dynamics under operating conditions (e.g., $T=700\text{--}1050\text{ K}$ for ammonia decomposition [15]), since the mechanism and the relevant environments may deviate significantly from the calculations performed at $T=0\text{ K}$. In fact, the effects of dynamics under operating conditions can manifest themselves in many different ways [55, 58, 60], including the dynamic change of the surface (e.g., see Fig. 3.1) and the emergence of a broad ensemble of reactive pathways (e.g., see Figs. 3.9 and 7.4), which would be impossible to capture by sampling the potential energy landscape with static calculations. As we will see in Section 2.4, enhanced sampling techniques are better suited to sample the reactive landscape at finite temperatures and to collect a diverse set of atomic environments [123, 124]. By integrating these sampling techniques with active learning strategies, it has been possible to construct ML potentials for a wide variety of rare events, from phase transitions [125–127] to chemical reactions [128–137], and catalytic processes [44, 63, 64]. This is the strategy we adopted in the case of Fe, as we will discuss in the next section.

Recently, these techniques have been also used to enhance the model uncertainty to explore high-uncertainty configurations [138–141], but with the complexity of the system under study, we did not explore this direction.

2.3.1 Active learning via *query-by-committee*

In Figure 2.3, we report the basic flowchart of active learning strategy.

In the active learning procedure, we want to use the MLP to generate configurations that could potentially be added to the dataset. However, an initial set of configurations is still required to train a first model that provides reasonable prediction. It is common practice to use AIMD simulations to generate these configurations. Due to the limited timescales accessible with AIMD, enhanced sampling techniques, such as metadynamics or umbrella sampling, can be combined to explore the relevant regions of phase space more efficiently. However, AIMD configurations are highly correlated due to the temporal proximity between successive snapshots. One solution is extracting configurations less frequently, typically every 3 to 5 ps, significantly reducing redundancy. This implies discarding at least 70% of the generated configurations to avoid overrepresenting similar structures in the dataset, a substantial

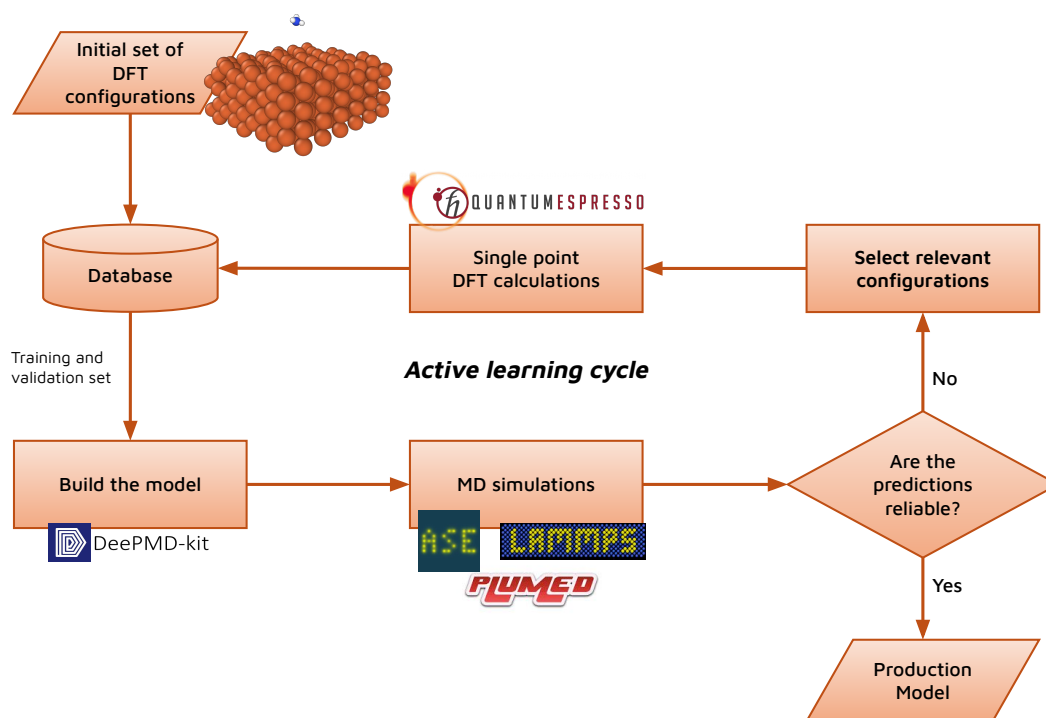


Figure 2.3: Flowchart of reactive MLP construction workflow via active learning. The icons of the principal software used during the active learning cycles in the case of Fe studies [110, 142–145]

computational waste. Another common approach is to generate initial configurations by randomly displacing atoms within the system. While this method allows for generating sparse and more varied configurations, it may lead to unphysical structures, and care must be taken when this approach is employed.

After collecting the first set of configurations, we enter the active learning cycle. In each iteration, a new MLP is trained and then used to sample new configurations through MD simulations. Then, a subset of these structures is selected, and single-point DFT calculations are performed. These data are then incorporated into the training set to optimize a new potential, which will be used to sample new configurations.

There are two crucial aspects of this procedure. The first is the sampling done through MD simulation, which should cover as much as possible of the phase space as possible relevant to the process under study. The use of advanced sampling methods, which we employed later for studying catalytic reactivity, can significantly enhance this active learning procedure. Indeed, they allow us to harvest uncorrelated structures, especially reactive ones. Generally, this is done by iteratively expanding the explored region. In the case of a reaction, in the first cycles, only metastable states

are sampled, and then high energy configuration and reactive paths should be sampled. In this stage, simulations are performed to only sample new configurations to be added to the training set and not for recovering observable expectation values. So, using the enhanced sampling techniques that ensure the quasi-static regime, such as OPES, is not strictly required, while Umbrella sampling and Metadynamics can also be used (see Section 2.4).

The second key element of this active learning cycle is the selection criteria. We would like to add to the training set only configurations that have not already been described well by the MLP. One way to do this is via the *query-by-committee* approach [117]. As we discussed in Section 2.2.4, from the prediction of a committee of models we can define an uncertainty $\sigma = \sigma_{\max}^{(\text{NN})}$ (see Eq. 2.14) and, thus, using it as a proxy for the prediction error. In Figure 2.4a, we report a possible selection

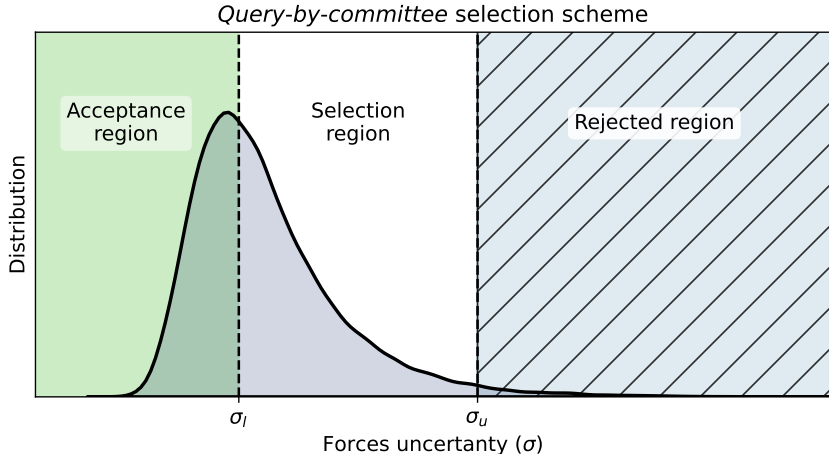


Figure 2.4: **Query-by-committee selection.** Typical distribution of the uncertainty of NN forces predictions $\sigma = \sigma_{\max}^{(\text{NN})}$ using a committee of models (see Eq. 2.14). A common choice of the different regions used to query the configuration based on σ is reported. See text for details.

criterion based on the uncertainty σ [63]. Starting from the distribution of σ , a lower (σ_l) and an upper σ_u limit are defined. The lower bound σ_l is typically set to a value slightly higher than the peak of the distribution and defines the acceptance region. Instead, the higher threshold σ_u is defined in order to avoid the selection of nonphysical configurations in which atoms are too close or correspond to improbable chemistry. Typically is defined as $\sigma_l + [0.2 \sim 0.3] \text{ eV } \text{\AA}^{-1}$. Then, a random fraction of the configuration with $\sigma \in [\sigma_l, \sigma_u]$ is selected to be relabeled at the QM level. This fraction depends on the total number of configurations in the distribution in order to limit the number of QM calculations to be performed.

This procedure is very effective for many cases, however, it has been shown that

it does not always adequately describe all reactive pathways, being unbalanced on reactants and products [129]. This is particularly true for very fast reactions such as dehydrogenation. For this reason, additional selection criteria based on collective variables that focus the selection on the TS region are adopted.

The active learning procedure (sampling, selection, QM relabelling, training) is iterated until a reliable potential for describing the full reactive pathways is obtained.

We applied this methodology to build the reactive potential for ammonia synthesis/decomposition on Fe. The total amount of configurations we collected was $\sim 110\text{k}$ (Table 2.1) for the studies on pristine surfaces, which became $\sim 250\text{k}$ (Table 2.4) when we included lateral interactions. To these configurations, we have to add the discarded configuration from AIMD. These are two times the ones included in the dataset, $\sim 120\text{k}$ (see Table 2.2), only for the case of pristine surfaces.

This strategy is effective but not efficient and cannot be applied to the computationally expensive FeCo catalyst. Thus, before studying this system, in Chapter 6 we will focus on the development of a method that allows us to train reliable MLP for chemical reactions in dynamic environments but requires a few QM calculations.

2.4 Free energies and enhanced sampling

Despite the acceleration provided by MLPs, numerous reactive processes, such as chemical reactions, unfold over timescales that surpass the limits of conventional simulations. Consequently, a spectrum of advanced sampling methods has been devised to facilitate the simulation of rare events, addressing the challenge posed by these extended timescales. We can categorize these methods into three main categories. The first family of techniques relies on the identification of collective variables (CVs) and introduces an external potential to enhance their fluctuations, thus reducing the free energy barriers along them, as in the case of Umbrella Sampling [146, 147], Metadynamics [148, 149] and Variationally Enhanced Sampling [150, 151]. In the second group, the equilibrium distribution is broadened in a more generic way, as in the case of replica exchange molecular dynamics [152] or generalized ensembles [153, 154]. Finally, the third category aims at sampling reactive paths by doing, for example, a Monte Carlo simulation in trajectory space [155]. We refer the interested reader to the review by Henin *et. al.* for a comprehensive overview of enhanced sampling methods [156].

In this section, we will briefly introduce the basic principle of the first category. In particular, we will focus the techniques that we used in this thesis, namely the CVs-based method On-the-fly Probability Enhanced Sampling (OPES) [157, 158], implemented in the open-source PLUMED plugin [144, 159]. At the end of the sec-

tion, we will describe the main CVs adopted across the studies.

2.4.1 Probabilities, free energies, and manipulating it

Probability and free energy profiles. Before talking about enhanced sampling, an introduction to the concept of probability and free energy profiles in MD simulations is required. In a canonical ensemble, this probability is related to the potential energy $U(\mathbf{x})$ of the system through the well-known Boltzmann distribution

$$p(\mathbf{x}) = \frac{e^{-\beta U(\mathbf{x})}}{Z} \quad \text{with} \quad Z = \int d\mathbf{x} e^{-\beta U(\mathbf{x})} \quad (2.16)$$

where $\beta = 1/k_B T$ is the inverse temperature (k_B the Boltzmann constant), and Z is a normalization constant called configurational partition function. The outcomes of an MD simulation conducted within the canonical ensemble are expected to be distributed accordingly to Eq. 2.16. Since this probability lives in a $3N$ dimensional space, where N is the number of atoms, it is convenient to understand the system in terms of a smaller set of degrees of freedom, the CVs. These CVs are, in general, functions of the atomic coordinates of the system $\mathbf{s} = \mathbf{s}(\mathbf{x})$, and, in the same spirit of reaction coordinates in chemistry or order parameters in physics, are intended to capture the progress of the molecular process under study. Defined one or a few CVs, one can define the probability of observing a given value of \mathbf{s} during the simulation as the marginal distribution of $p(\mathbf{x})$ along \mathbf{s}

$$p(\mathbf{s}) = \int d\mathbf{x} \delta(\mathbf{s} - \mathbf{s}(\mathbf{x})) p(\mathbf{x}). \quad (2.17)$$

The logarithm of this probability, which is called the *free energy surface* (FES) or free energy profile as a function of the CVs \mathbf{s} :

$$F(\mathbf{s}) = -\frac{1}{\beta} \log p(\mathbf{s}) \quad (2.18)$$

This quantity can provide numerous insights into the processes studied: the free energy minima correspond to the metastable states, and the free energy barriers correspond to the transition states between them. To get out of a metastable state, the system must overcome free energy barriers where the probability of being in that state is exponentially suppressed.

Furthermore, the relative stability between two metastable states, A and B, can

be measured by the free energy difference

$$\Delta F = F_B - F_A = \frac{1}{\beta} \ln \frac{p_A(\mathbf{s})}{p_B(\mathbf{s})} = \frac{1}{\beta} \ln \left(\frac{\int_A d\mathbf{s} e^{-\beta F(\mathbf{s})}}{\int_B d\mathbf{s} e^{-\beta F(\mathbf{s})}} \right) \quad (2.19)$$

where the integration domains correspond to the regions in CV space assigned to each metastable state. Note that the difference in free energy ΔF is hardly affected by the choice of CV, provided it can distinguish between the two minima [160]. However, more care should be taken to extract the free energy barrier from FES, as the profiles depend on the choice of CV [161]. A general rule of thumb when choosing a CV among a set of ones is the higher the barrier, the better the CV.

Free energy barrier and chemical reactions. Free energy can be related not only to ensemble quantities but also to experimentally observable quantities. In particular, the rate constant k of a chemical reaction ($A \rightarrow B$) can be calculated from transition state theory using the Eyring-Polanyi equation:

$$k = \kappa \frac{k_B T}{h} e^{-\frac{\Delta G^\ddagger}{RT}} \quad (2.20)$$

where κ is transmission coefficient, h is Planck's constant, ΔG^\ddagger is the Gibbs free energy of activation, and R is the universal gas constant. From the reaction rate, we can calculate the half-life $t_{1/2}$ of the reaction (assuming non reversible reactions):

$$t_{1/2} = \ln(2)k^{-1}, \quad (2.21)$$

which is the time it takes for half of the reactants to be converted to products. This equation connects the rate of a chemical reaction to the free energy barrier between the reactants and the transition state ΔG^\ddagger . Note that, while the equations are often presented in terms of the Gibbs free energy, in cases where pressure and volume are constant, the Helmholtz free energy can be used as an equivalent substitute.

In general, we will estimate reaction rates and not even present the results in these terms. However, it's good to remember that even without knowing all the parameters, such as the transmission coefficient, these equations provide an estimate of the times at which reactions occur, allowing us to assess whether a reaction is fast or slow under specific conditions. In particular, it is useful for comparing different reactions and evaluating how changing the conditions of the system alters the probability of observing an event.

Alter the probability through enhanced sampling. In CVs-based methods, an external potential that depends on the atomic coordinates through the CVs, i.e., $V(\mathbf{x}) = V(\mathbf{s}(\mathbf{x}))$, is added to the system to decrease the free energy barriers along the CVs and increase their fluctuations [162]. Notable examples of algorithms that utilize this principle include Umbrella Sampling [146] and Metadynamics [148]. From the probability perspective, the addition of the bias potential determines a new probability distribution which is sampled during the simulation, which is called *biased distribution*:

$$p_V(\mathbf{x}) = \frac{e^{-\beta(U(\mathbf{x})+V(\mathbf{s}(\mathbf{x})))}}{Z_V} \quad \text{with} \quad Z_V = \int d\mathbf{x} e^{-\beta(U(\mathbf{x})+V(\mathbf{s}(\mathbf{x})))} \quad (2.22)$$

Reweighting An essential part of the enhanced sampling technique is the requirement to evaluate the properties of the original system (e.g., its free energy surface) from the samples obtained via the modified distribution p_V . This necessitates implementing a correction method known as *reweighting*. We have already seen that under the ergodicity hypothesis, a generic observable's ensemble average is equivalent to its time average over trajectories (see Equation 2.3). This hypothesis holds in the case that the bias potential V is static (or can be approximated as such), which means that it is known and independent from time. With some math, we can compute the ensemble average $\langle O(\mathbf{x}) \rangle$ under the original Boltzmann distribution using samples that are distributed according to $p_V(\mathbf{x})$:

$$\langle O(\mathbf{x}) \rangle = \int d\mathbf{x} O(\mathbf{x})p(\mathbf{x}) = \int d\mathbf{x} \left(O(\mathbf{x}) \frac{p(\mathbf{x})}{p_V(\mathbf{x})} \right) p_V(\mathbf{x}) = \frac{\langle O(\mathbf{x})e^{\beta V(\mathbf{s}(\mathbf{x}))} \rangle_V}{\langle e^{\beta V(\mathbf{s}(\mathbf{x}))} \rangle_V} \quad (2.23)$$

where $\langle \cdot \rangle_V = \int d\mathbf{x} \cdot p_V(\mathbf{x})$ denotes the average under the biased distribution. For a sufficiently long trajectory, we can use Equation 2.4 to rewrite this quantity as a weighted average of the observable along the trajectory:

$$\langle O(\mathbf{x}) \rangle \approx \sum_{k=1}^{n_k} O(\mathbf{x}_k)w(\mathbf{x}_k) \quad \text{with} \quad w(\mathbf{x}_k) = \frac{e^{\beta V(\mathbf{s}(\mathbf{x}_k))}}{Z_w} \quad (2.24)$$

where \mathbf{x}_k are the samples obtained in the biased MD, and $w(\mathbf{x}_k)$ are weights of the k -th configuration of the trajectory which are normalized for $Z_w = \sum_k e^{\beta V(\mathbf{s}(\mathbf{x}_k))}$. The reweighting procedure then corresponds to correcting the weight of each sample to match that of the original distribution. In particular, we can use Eq. 2.24 to estimate the unbiased probability distribution $p(\mathbf{s})$, and consequently the free energy profile (Eq. 2.18). To do this, we note that the definition of the marginal of the density along \mathbf{s} (Eq. 2.17) can be rewritten as the ensemble average of a delta function, i.e.,

$p(\mathbf{s}) = \langle \delta(\mathbf{s} - \mathbf{s}(\mathbf{x})) \rangle$. Therefore, we obtain the following expression for the unbiased probability:

$$p(\mathbf{s}) \approx \sum_k^{n_k} \delta(\mathbf{s} - \mathbf{s}(\mathbf{x}_k)) w(\mathbf{x}_k) \quad (2.25)$$

2.4.2 On-the-fly Probability Enhanced Sampling (OPES)

Now, we describe the On-the-fly Probability Enhanced Sampling (OPES) method that we primarily adopted for biasing simulation in this thesis. OPES was developed not only to accelerate simulations but also to do so in a way that makes it easy to retrieve unbiased statistical information about the system. To do this, OPES adds a bias potential $V(\mathbf{x})$ to target a distribution $p^{tg}(\mathbf{x})$, which still retains a good overlap with the original distribution $p(\mathbf{x})$ in configuration space. Thus, one can efficiently perform reweighting and obtain statistical information about $p(\mathbf{x})$. Since the distribution $p(\mathbf{x})$ is unknown at the beginning of the simulation, the bias is optimized in a self-consistent way, at convergence hold:

$$V(\mathbf{x}) = \frac{1}{\beta} \log \frac{p(\mathbf{x})}{p^{tg}(\mathbf{x})} \quad (2.26)$$

As in Metadynamics, the bias potential is time-dependent, constantly updating the bias as new samples are collected. The iterative scheme aims at quickly reaching a regime of quasi-static bias. This way, most of the MD time can be used for statistical analysis. Many different variants of OPES have been proposed, depending on bias construction and target probability. Here, we focus on OPES-Metad [163] which is the original OPES formulation. Here, the target p^{tg} is chosen as the well-tempered distribution in CV space $p^{WT}(\mathbf{s})$ [149] that aims to broaden the equilibrium probability $p(\mathbf{s})$:

$$p^{WT}(\mathbf{s}) \propto [p(\mathbf{s})]^{1/\gamma} \quad (2.27)$$

where the parameter γ (called bias factor) regulates the magnitude of the smoothing. This $p^{WT}(\mathbf{s})$ is the same distribution that is sampled asymptotically in Well-Tempered Metadynamics [149]. In terms of free energy, the effective barriers in this distribution are decreased by a factor γ . The intuitive idea is to stay as close as possible to the original distribution, $p(\mathbf{s})$, while reducing the barriers enough to obtain frequent transitions. The bias is iteratively optimized after the estimation of the unbiased probability $p(\mathbf{s})$ through reweighting. A trick that makes OPES particularly efficient is employing a weighted kernel density estimation (KDE) to make $p(\mathbf{s})$ continuous and differentiable. In practice, the main parameters that we have to set in simulations are the pace, that is the stride of the bias, the barrier parameter

$\Delta E \sim \gamma\beta^{-1}$, which represents the expected height of the highest free barrier, and the sigma, that is the initial width of the deposited kernels (if not specified, it is estimated adaptively by the OPES algorithm in the initial $10 \times$ pace steps). The specific parameters adopted in simulations are reported in Section 2.7.

OPES-Flooding In our simulations, we will also use a recent variant of OPES, named OPES-Flooding [164, 165]. This variant was developed for efficiently collecting unbiased transition and studying kinetic rates by combining the OPES biasing scheme and the ideas of conformational flooding [166, 167] and hyperdynamics [168]. In this work, we do not attempt to use the flooding scheme to statistically recover the unbiased kinetics of the process [165, 169]. Instead, we used it to generate unbiased reactive trajectories escaping from a known metastable state and characterize transition states. In practice, OPES-flooding is a modification of OPES-Metad that realizes straightforwardly the requirement of not perturbing the transition state due to its flexibility. The OPES-Flooding idea is to build a bias $V(\mathbf{s})$ along a CV \mathbf{s} that is aimed at only partially filling one of the metastable states, increasing the probability of escaping from the minima but leaving TS region bias-free. The introduction of an additional parameter ensures this last condition, the so-called *excluded region* \mathbf{s}_{exc} , and imposing that no bias is deposited for $\mathbf{s} \in \mathbf{s}_{exc}$.

It is important to remember that in addition to the external bias aimed at increasing the fluctuation along specific CVs, one often has to add bias to confine the system to a certain state. This is done by applying a so-called constraint or wall. In this thesis, we have always utilized harmonic constraints, which act by applying an external potential of harmonic type when the variable s exceeds a certain threshold s_t :

$$V_{\text{wall}}(s) = \begin{cases} 0 & \text{if } s \leq s_t \\ k\left(\frac{s-s_t}{s_i}\right)^2 & \text{if } s > s_t \end{cases} \quad (2.28)$$

where k is the spring constant and s_i a regularization parameter. A similar definition can be used for lower bounds. Multiple bias walls can be defined, and they should be included in the reweighting procedure. Generally, it is only used to confine the system's dynamic during MD and the configuration where the wall act is discarded.

2.4.3 Collective variables

A key element in enhanced sampling methods is the choice of collective variables (CVs). The efficiency of enhanced sampling relies heavily on high-quality CVs; poorly chosen variables can significantly reduce the method's effectiveness. Properly chosen

CVs can accelerate sampling exponentially by targeting all relevant degrees of freedom.

CVs must satisfy two important requirements: continuity and differentiability across the configurational space, and respect for the symmetries of the system (such as rotational and translational invariance, and occasionally permutational symmetry). The biasing potential is applied based on the force $f = -\nabla_{\mathbf{x}}V(\mathbf{s}(\mathbf{x}))$, which depends on both the derivative of the bias with respect to the CVs, \mathbf{s} , and the derivative of the CVs with respect to atomic positions, \mathbf{x} .

Moreover, CVs should facilitate dimensionality reduction. Most biasing methods work best with a small number of CVs, and these CVs need to distinguish between different metastable states and encode slow modes that hinder sampling. For instance, CVs must be able to differentiate transition states from metastable ones.

Typically, CVs are chosen based on physical or chemical intuition, often relying on descriptors such as bond distances. While effective for simple systems, this method fails in more complex environments, such as dynamic catalytic reactions where multiple degrees of freedom are involved. To address these challenges, data-driven approaches have emerged, leveraging machine learning and statistical techniques to identify CVs directly from high-dimensional data.

It must be remembered that CVs are useful also besides adding bias to the system. Indeed, they are a powerful analysis tool, and even if bias is applied to a CV or several CVs, free energy need not be projected only onto those CVs. In the case CVs are not implied for biasing, the differentiability constraints can be relaxed. In this study, we used only traditional variables. In fact, the system's geometry is quite simple since the reactions of interest are related to NH_x molecules. However, we want to remember that while the description may be geometrically simple, a few light atoms on a metal substrate represent a drop in a bucket, making the study of these reactions challenging. In addition, the free energy barriers we're going to encounter are of the orders of several $k_B T$, making sampling even more complex. Now we will give an overview of the CVs that we have used during this work; the specific parameters are reported in Section 2.7

The distance between two atoms is the most elementary, but also powerful CV that we use in this work. It is very useful for the definitions of walls, but also for accelerating bond breaking. It can also be decomposed in its vectorial components to define, for example, a planar distance.

Other simple CVs are the vectorial components of the center of mass of a group of atoms (generally a molecule), which can be applied to enhance and study the diffusion of a molecule on the surface.

An important CV used for both enhancing sampling and analysis is the coor-

dination number between one or a group of atoms X atoms and another group of atoms Y , computed in differentiable form as:

$$C_{X,Y} = \sum_{i \in X} \sum_{j \in Y} s(r_{ij}), \quad s(r_i) = \frac{1 - \left(\frac{r_i - d_0}{r_0}\right)^n}{1 - \left(\frac{r_i - d_0}{r_0}\right)^m} \quad (2.29)$$

where r_{ij} is the distance between the i -th atom of X specie and the j -th atom of Y specie, r_0 , d_0 , n , and m , are tunable parameters. This is one of the key CVs that we will use in almost all of our studies.

To study the diffusion of atomic/molecular species over (110) surfaces, we construct CVs that embed the cristal. In particular, we define the position on the plane of the center of mass of the atom/molecule expressed in terms of the lattice vector direction:

$$\begin{aligned} v_{1\bar{1}1} &= \frac{1}{a\sqrt{3} \cos \alpha \sin \alpha} (x \sin \alpha + y \cos \alpha) \\ v_{1\bar{1}\bar{1}} &= \frac{1}{a\sqrt{3} \cos \alpha \sin \alpha} (-x \sin \alpha + y \cos \alpha) \end{aligned} \quad (2.30)$$

where x and y are the coordinates of the center of mass (or the position along the $[001]$ and $[1\bar{1}0]$ directions, α is the angle between $[1\bar{1}1]$ and $[1\bar{1}0]$ directions, and a is the lattice parameter.

Analogously, we define another transformation of the coordinates of nitrogen on the (110) surfaces to describe the migration inside direction. Here we have to consider also the z position:

$$\begin{aligned} \tilde{v}_{1\bar{1}1} &= \frac{1}{a} \left(x + \frac{1}{\sqrt{2}} y + \frac{1}{\sqrt{2}} z \right) \\ \tilde{v}_{1\bar{1}\bar{1}} &= \frac{1}{a} \left(-x + \frac{1}{\sqrt{2}} y + \frac{1}{\sqrt{2}} z \right) \\ v_{0\bar{1}0} &= \frac{\sqrt{2}}{a} (-y + z) \end{aligned} \quad (2.31)$$

where $[0\bar{1}0]$ is the migration inside direction. Here, we also redefine the \tilde{v}_{hkl} to include also the layer dependence through z . This set of CVs is able to efficiently describe the diffusion "channel" the N follows inside the Fe/FeCo (considering a slab exposing the (110) plane)

To simulate the first dehydrogenation step we use a set of 3 CVs: the $C_{N,Fe}$, the $C_{N,H}$, and a geometrical CV based on the angular distribution function defined as

follows:

$$\Theta = \int_{\alpha}^{\beta} \sum_{ij \in H} s(r_i)s(r_j) \frac{1}{\sqrt{2\pi}w} \exp\left(-\frac{(\theta - \theta_{ij})^2}{2w^2}\right) \quad (2.32)$$

where θ_{ij} are the HNH angles, r are the NH distances, $\alpha = 0.55\pi$, $\beta = 0.65\pi$, and $w = 0.5(\alpha - \beta)$. The latter parameters are used to define the angular windows in which the reaction occurs effectively. This is important to effectively simulate the back-and-forth of the reaction.

To conclude this section we want to anticipate that we will use an additional set of CVs to analyze the reactions, no longer based on geometric but rather chemical considerations, as we will discuss in the next section.

In this work, we will often project the free energy surface in two dimensions, find the minimum energy paths, and compute the free energy profile along these. For finding these paths, we have used the Minimum Energy Path Surface Analysis (MEPSA) code [170].

2.5 Machine learning for partial charges

In complex chemical environments, relying solely on CVs derived from geometrical configurations can be problematic due to the high variability and complexity of atomic arrangements. Traditional machine learning potentials MLPs focus on predicting only energies and forces, often overlooking crucial electronic information, which is essential for an accurate description of chemical processes. Moreover, the structures predicted by MLPs are typically too large to directly solve the electronic structure. In these cases, direct electronic calculations using QM methods (e.g. DFT) would be prohibitively expensive, making it impractical to perform *post hoc* charge analysis after simulations.

To address this issue, Bonati et al. [44] developed a surrogate MLP that retains the same descriptors and architecture used in energy and force predictions but is trained to predict atomic charges instead of energy or forces. Since partial charges are a *per-atom* property, the neural network can be trained on smaller systems and subsequently applied to predict charges for arbitrarily large systems.

Following the methodology developed in [44], we trained a second NN to predict partial charges based on atomic positions and chemical species. Reference values were obtained from a subset of the Fe dataset (approximately 75,000 structures) using the Bader charge decomposition scheme [171, 172].

We utilized the deep tensor neural network architecture SchNet [173], a GNN, with the same parameters as in [44]: 5 interaction layers, pairwise distances expanded

on 30 Gaussians, 64 atom-wise features, 3.5 Å cosine cutoff, and an atom-wise module with 2 hidden layers and [64, 64] nodes per layer. The loss function was the root mean square error (RMSE) between predicted charges $\{q_i\}$ and reference charges $\{q_i^{DFT}\}$:

$$L = \frac{1}{N_{at}} \sum_i (q_i - q_i^{DFT})^2 \quad (2.33)$$

The dataset was split into training and validation subsets with an 80:20 ratio, and the model was optimized using the ADAM optimizer [174] with a learning rate of 0.001 and an early stopping criterion. This approach achieved an RMSE of 10^{-4} e on the validation set.

The charges transferred between atoms were expressed as deviations from their nominal valence charge q^V (8e, 5e, 1e for Fe, N, and H, respectively). The total charge transferred from the metallic substrate to the N_xH_y species was calculated as:

$$q_{N_xH_y} = \sum_{i \in N_xH_y} (q_i - q_i^V). \quad (2.34)$$

This quantity has been used as a CVs in our studies.

At this stage, we do not report results on charge transfer for FeCo, as those studies are still ongoing. In these ongoing studies, we are extending the approach to compute charges alongside energy and forces using the same descriptors, integrating charges into the MACE framework. As a future perspective, this methodology would provide access to analytical derivatives of the charges with respect to atomic coordinates. Although this is not a physical observable, the output from the NN preserves the bias-CVs requirement, and this could be used to bias simulations directly, turning atomic charges into an active simulation tool rather than just an analysis feature.

2.6 Other analysis tools

2.6.1 Transition states characterization

To obtain an unbiased characterization of the transition states (TS), no bias potential should be added in this region [175, 176]. This has been done for studying the dehydrogenation steps on both Fe and FeCo. To this aim, we adopt the OPES-flooding scheme described in Section 2.4.2. Briefly, in this method, the OPES barrier parameter ΔE is properly chosen to a value lower than the effective barrier, and the transition state region is excluded from any bias potential modifications. In such a way, dehydrogenation events are observed to be bias-free. These simulations enable us to collect statistics on the reactive pathways and characterize the TS ensemble

without any prior assumptions about the location of the transition and final states.

The same CVs of the standard OPES simulations are applied to describe each dehydrogenation step, but with lower barrier parameters and no harmonic constraints applied. For Fe(110), the OPES barrier parameters ΔE are set equal to 50(60), 20(65), and 65(50) kJ mol⁻¹ for each step respectively, while no bias is added in the region where $C_{N,H}$ is lower than 2.82, 1.85, 0.9. For Fe(111), the OPES barrier parameters ΔE are set equal to 60, 65, and 50 kJ mol⁻¹.

Analogously, for the FeCo(110), we exclude from bias deposition the same region, and we use the coordination $C_{N,H}$ as bias-CV with ΔE equal to 70, 30, and 80 kJ mol⁻¹ for the three hydrogenation steps. All simulations were stopped once the reaction occurred, terminating the simulation when the maximum NH distance exceeded 1.9 Å.

One hundred simulations of dehydrogenation events were performed for each step and surface. We then defined the TS as the configuration in which the bonding distance between the N atom and the reactive H atom, $d_{N,H}$, is closest to the maximum free energy of the dehydrogenation barrier. In the study of Fe, to characterize the TS ensemble, we apply the k-medoids cluster analysis as implemented in the `kmedoids` Python library [177] to identify at least two classes of TS geometry for each dehydrogenation step on both surfaces by using the coordination coordinates $C_{N,Fe}$ and $C_{H,Fe}$.

2.6.2 Surface analysis

Here, we present the analysis tools that we used for the study of the morphology and dynamics of Fe(111) under N^* coverage.

To examine the structure and movement of the iron surface, we initially employ the Alpha-Shape technique [178], which is implemented in OVITO [102], to reconstruct the surface based on the atomic positions. This method utilizes a virtual sphere with a radius of 2 Å to create a three-dimensional surface mesh that distinguishes the accessible volume (void) from the inaccessible volume (slab). By doing so, it becomes feasible to identify the atoms associated with the surface at each time step and focus the subsequent analysis solely on these atoms, despite their temporal variations resulting from their high mobility at operando conditions ($T = 700\text{ K}$).

2.6.2.1 Environment Similarity

To assess the similarity of the environment surrounding an atom with a reference, we use the Environment Similarity measure introduced by Piaggi and Parrinello [179]. This can be considered as a variation of the widely used SOAP (Smooth Overlap of

Atomic Positions) kernel [180], but without rotational invariance. Initially, we define a smooth local density around the central atom by applying Gaussian fitting to the position of each neighboring atom i :

$$\rho_x(r) = \sum_{i \in \chi} \exp\left(-\frac{|r_i - r|^2}{2\sigma^2}\right) \quad (2.35)$$

where σ is a broadening parameter and r_i is the position of atom i with respect to the central atom. The environment similarity between χ and a reference χ_{ref} is subsequently defined as follows:

$$S(\chi, \chi_{ref}) = \int dr \rho_x(r) \rho_{x_{ref}}(r) = \frac{1}{n} \sum_{i \in \chi} \sum_{j \in \chi_{ref}} \exp\left(-\frac{|r_i - r_j|^2}{4\sigma^2}\right) \quad (2.36)$$

where n represents the number of atoms in the environment χ_{ref} . After performing integration, we normalize the kernel such that $S(\chi_{ref}, \chi_{ref}) = 1$.




Figure 2.5: Representative of χ_7 and $\chi_{Fe_3^*}$ local environments. Iron atoms are colored based on their z coordinate, as shown in the scale bar.

The reference environments chosen here are the active sites on (111) surface (χ_7) and the triangles forming on top of the vacancy ($\chi_{Fe_3^*}$), see Figure 2.5. The χ_7 is the environment of a surface atom of the third layer corresponding to the cavity in which the N_2 is adsorbed in the α' state, surrounded by 7-coordinated Fe atoms. While for $\chi_{Fe_3^*}$, the reference environment is chosen with one of the three iron atoms as the central atom. We also considered the locally deformed χ_7 by the presence of 1-3 N^* in its vicinity as the references for the cases with finite N^* concentration and the best match was taken. Regarding the computational parameters, we considered all neighboring Fe atoms within a distance of 3.5 \AA to construct the local density. To mitigate the effects of thermal fluctuations, we applied a moving average to the atomic positions using a window size of 2.5 ps. We used a broadening parameter of

$\sigma = 0.185$ at $T = 700 K$, identical to the one used by Bonati et al. [44].

After obtaining the environment similarity values for each surface atom, we identify the sites C_7 and Fe_3^* as those with a similarity score, $\kappa(\chi, \chi_7)$ and $\kappa(\chi, \chi_{Fe_3^*})$, respectively, equal to or greater than a threshold. The selection of these threshold values is determined based on the minima of the distributions of $\kappa(\chi, \chi_7)$ and $\kappa(\chi, \chi_{Fe_3^*})$ for each N^* concentration, reported in Figure 2.6.

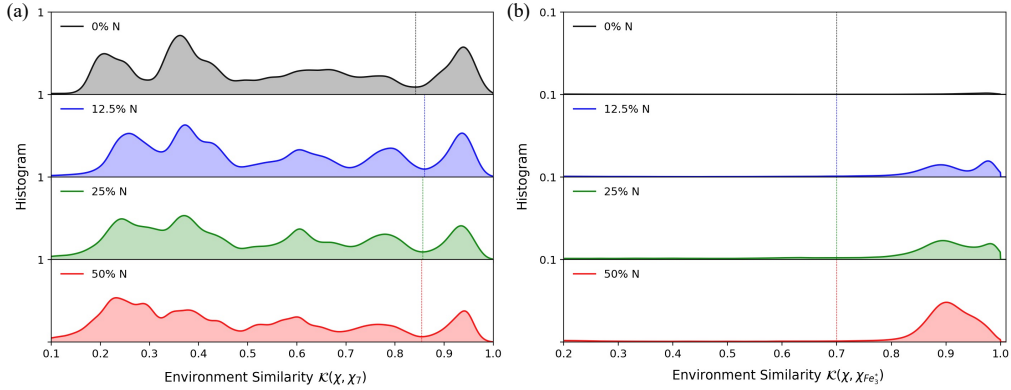


Figure 2.6: The distribution of the environment similarity for the surface iron atoms of various N^* covered surfaces with (a) χ_7 and (b) $\chi_{Fe_3^*}$ at 700 K. The threshold value to determine the density of C_7 is chosen based on the minima of the histogram in a) as denoted by dashed vertical lines. For Fe_3^* , the threshold value of 0.7 is chosen as indicated by vertical dashed lines in b).

2.6.2.2 Mean Square Displacement

We calculate the mean square displacement (MSD) of Fe(111) surface atoms and N^* atoms separately using the following Einstein relation [181] :

$$\text{MSD}(t) = \left\langle \frac{1}{N} \sum_{i=1}^N |r_d(t) - r_d(t + \tau)|^2 \right\rangle_{\tau} \quad (2.37)$$

where N is the number of equivalent particles the MSD is calculated over, r are their coordinates, d is the desired dimensionality of the MSD and τ represents the lag time. MSD is averaged over all possible lag-times $\tau \leq \tau_{max}$, where τ_{max} is the length of the trajectory. To calculate the MSD, we utilized the MDAnalysis package [182–184].

The first 20 ns of the trajectory were discarded as considered equilibration time. The diffusion coefficients (D) for the self-diffusion of Fe(111) surface atoms and N^* atoms on the xy-plane are determined by fitting a straight line between 20 ns and

40 ns of MSD_{xy} (MSD in xy-plane) vs t and using the following correlation:

$$D = \frac{MSD_{xy}}{4t}. \quad (2.38)$$

2.7 Computational details

In this section, we report all the details needed to reproduce the results in our simulations that have not yet been reported. To avoid possible misunderstandings, we decided to separate this section according to the different studies conducted in this thesis work. First we will present the details of the studies of ammonia decomposition on iron. Then those employed for the study of lateral interactions and nitridation. Finally those for the study of FeCo-related reactions. The order we present here does not reflect the order in which the results will be presented. For the sake of completeness in reporting all details, some repetition may be present.

2.7.1 NH_3 decomposition on iron

2.7.1.1 MD simulations

Classical molecular dynamics simulations were performed using the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) software [142], supplemented by DeepMD-kit 2.1 [109] and PLUMED [144]. The time step was set to 0.5 fs, and the temperature was controlled with a stochastic velocity scaling thermostat [78] with a coupling constant of 100 fs.

Small surfaces containing 40 to 108 Fe atoms and up to 10 adsorbates (N/H) were simulated during the active learning phase.

After having optimized the potential, longer simulations of at least 20 ns were performed using an $8 \times 10 \times 10$ Fe(110) slab (800 atoms) along with an NH_x ($x=1-3$) molecule to simulate adsorption and dissociation. Analogously, an $8 \times 8 \times 12$ Fe(111) slab (768 atoms) along with an NH_x ($x=1-3$) molecule were simulated. For the (110) and (111) slabs, two and three layers of the slab were fixed to create a boundary condition that imitates that of a semi-infinite slab. Periodic boundary conditions were utilized in the x- and y-directions. At the same time, a reflecting wall was implemented over the surface in the z-direction, in order to prevent the NH_3 molecule from flying away. The distance between the wall and the top layer was set such that the volume accessible to the ammonia molecule when detached from the surface is similar to the average volume of NH_3 in the gas phase at 1 bar. This is the pressure normally used in ammonia decomposition [15].

2.7.1.2 Data set compositions

The dataset was composed through an active learning scheme (see, Section 2.3.1). The resulting dataset comprises about 110k configurations. These include some of the calculations we collected for the ammonia synthesis studies on Fe [44, 185], along with many others made specifically for this work to ensure that all of the following physical scenarios of interest were considered:

- long-term dynamics of Fe surfaces
- surface with nitrogen hydrides (NH_x , where x ranges from 3 to 1)
- diffusing hydrides
- reactive configurations of all dehydrogenation steps.

The detailed composition is reported in Tables 2.1 and 2.2.

2.7.1.3 ML-based potential details

To fit the potential energy we employed the Deep Potential Molecular Dynamics (DeepMD) [107, 186] method. Here, we report the technical details. Three hidden layers with [30, 60, 120]/[240, 240, 240] nodes each were used respectively for the embedding and fitting network. A cutoff radius of 6.0 Å (with 0.3 Å switching function for continuity) was set for interacting atoms. We check that decreasing the cutoff up to 5.0 Å doesn't affect the structural property of iron. In the training phase, an exponential-type learning curve was applied to decrease the learning rate from 0.001 to 3.5×10^{-8} . The decay step was set to 2 epochs, and the models were trained for 200 decay steps. The performance was evaluated on energy and forces with a weighted root mean square error (RMSE) loss function. The weights of the energy and forces were adjusted during the optimization from 0.02 to 1, and from 1000 to 1, respectively. In the active learning procedure, an ensemble of 4 different models is optimized in parallel on permutations of the training and validation datasets to assign an uncertainty to the configurations. At the end of the active learning procedure, the RMSE achieved for energy and forces are 1meV/atom and 50meV/Å, respectively (see Figure 2.7)

2.7.1.4 Reliability of ML-based potential

Furthermore, to benchmark the reliability of the potential, we used it to calculate the potential energy barriers for all dehydrogenation steps. The dimer method was used to identify the transition states [187, 188]. A $4 \times 6 \times 5$ Fe(110) slab (120 atoms) and a $4 \times 4 \times 6$ Fe(111) slab (96 atoms) along with an NH_x ($x=1-3$) molecule were

Surface	Formula	Adsorbate	N. config.	Surface	Formula	Adsorbate	N. config.	
110	Fe ₄₀	-	353	111	Fe ₄₅ HN	NH	1067	
		N-H	1176		Fe ₄₅ N	N	2423	
	Fe ₄₀ H ₂ N	NH-H	2182		Fe ₄₅ N ₂	N	2399	
		NH ₂	1013		Fe ₄₅ N ₄	N ₂	6285	
	Fe ₄₀ H ₃ N	NH ₃	1054			N	2388	
	Fe ₄₀ H ₄	H ₂	1870			NH-H	1000	
	Fe ₄₀ H ₈	H ₂	1887		Fe ₅₄ H ₃ N	NH ₂ -H	998	
	Fe ₄₀ HN	NH	1010			NH ₃	1000	
	Fe ₄₀ N	N	2423			NH-H-NH ₃	500	
		N	2351		Fe ₅₄ H ₆ N ₂	NH ₂ -H-NH ₃	500	
	Fe ₄₀ N ₂	N ₂	3362			NH ₃	1000	
	Fe ₇₂ H ₂ N	NH ₂	1500		Fe ₇₂	-	630	
	Fe ₇₂ H ₃ N	NH ₃	1000		Fe ₇₂ H ₁₀ N ₂	N-H	100	
	Fe ₇₂ HN	NH	1500		Fe ₇₂ H ₁₂	H ₂	3741	
	Fe ₈₀ H ₁₂ N ₄	N-H-H ₂ -N ₂	444		Fe ₇₂ H ₁₂ N	N-H	286	
	Fe ₈₀ H ₁₈ N ₆	N-H-H ₂ -N ₂	109			N-H	100	
	Fe ₈₀ H ₇ N ₃	N-H-H ₂ -N ₂	1500		Fe ₇₂ H ₁₂ N ₂	N ₂ -H	1233	
	Fe ₈₀ H ₈	H ₂	1069		Fe ₇₂ H ₁₂ N ₄	N-H	542	
	Fe ₈₀ H ₈ N	NH	664		Fe ₇₂ H ₂₀ N ₄	N-H-N ₂	563	
		N-H	298		Fe ₇₂ H ₃₀ N ₆	N-H-N ₂ -H ₂	251	
	Fe ₈₀ H ₈ N ₂	N ₂ -H	1232		Fe ₇₂ H ₆ N ₂	N ₂ -H	420	
	Fe ₈₀ H ₈ N ₃	NH ₂	1500		Fe ₇₂ H ₇ N ₂	N ₂ H-H	584	
	Fe ₈₀ H ₈ N ₄	N-H	512		Fe ₇₂ H ₈ N ₂	N-H	598	
	Fe ₈₀ N	N	1611		Fe ₇₂ N	N	990	
	Fe ₈₀ N ₂	N	1292		Fe ₇₂ N ₂	N	836	
	Fe ₈₀ N ₄	N	72			N ₂	994	
	111	Fe ₄₅	-		4370	Fe ₇₂ N ₄	N	92
		Fe ₄₅ H ₁₂	H ₂		374	Fe ₉₆ HN	NH	497
Fe ₄₅ H ₁₂ N		N-H	1604	Fe ₉₆ H ₂ N	NH ₂	497		
Fe ₄₅ H ₁₂ N ₂		N-H	1550	Fe ₉₆ H ₃ N	NH ₃	497		
Fe ₄₅ H ₁₂ N ₄		N-H	1533	Fe ₁₀₈	-	2456		
Fe ₄₅ H ₂₀		H ₂	656	Fe ₁₀₈ H ₈ N ₂	N-H	600		
		N-H	1017	Fe ₁₀₈ N	N	859		
Fe ₄₅ H ₂ N		NH-H	1793	Fe ₁₀₈ N ₂	N	854		
		NH ₂	1036	Fe ₄₈	-	674		
		N-H	762	Fe ₄₈ HN	NH	1114		
Fe ₄₅ H ₂ N ₂		NH	720	Fe ₄₈ H ₂ N	NH ₂	1116		
		NH-N-H	1925	Fe ₄₈ H ₃ N	NH ₃	1060		
		NH ₂ -H	1463	Fe ₄₈ N ₂	N ₂	1939		
Fe ₄₅ H ₃ N		NH ₃	4586		NH-H	994		
Fe ₄₅ H ₄ N		N-H	2003	Fe ₇₂ H ₃ N	NH ₂ -H	1000		
Fe ₄₅ H ₄ N ₂		N-H	1774		NH ₃	992		
Fe ₄₅ H ₄ N ₄		N-H	1596		NH-H	1000		
Fe ₄₅ H ₈ N		N-H	1902	100	Fe ₅₄ H ₃ N	NH ₂ -H	1000	
Fe ₄₅ H ₈ N ₂		N-H	1875		NH ₃	996		
Fe ₄₅ H ₈ N ₄		N-H	1592					

Table 2.1: **Dataset composition.** Composition of the configuration in the dataset for training and validation of the potential divided by surface, chemical formula, and adsorbate type.

implied at this stage. Geometries were optimized using a force-based conjugate gradient method in all calculations. The lowest two and three Fe layers were kept fixed during optimization and MD calculations for (110) and (111) surfaces, respectively. Saddle points and minima were considered convergent when the maximum force in

Group subdivision					
Surface	N. config.	Adsorbate	N. config.	Driver	N. config.s
110	32984	-	8483	from AIMD	58601
111	71911	NH ₃	13648	active-learning	58179
211	8889	NH ₂	14135		
100	2996	NH	19286		
		N	18590		
		N ₂	12580		
		H ₂	9597		
		combinations	20461		

Table 2.2: **Dataset composition, summary.** Summary of the Composition of the configuration in the dataset for training and validation of the potential.

each degree of freedom was less than 0.01 eV\AA^{-1} . Then, we recalculated the energy difference between the structures obtained at the DFT level. Table 2.3 shows the results achieved. The activation energies are in excellent agreement, with a maximum discrepancy of 0.07 eV .

	$E_a^{DFT} - E_a^{NN} [\text{eV}]$	
	Fe(110)	Fe(111)
$\text{NH}_3^* \rightarrow \text{NH}_2^* + \text{H}^*$	0.01	-0.01
$\text{NH}_2^* \rightarrow \text{NH}^* + \text{H}^*$	-0.07	0.06
$\text{NH}^* \rightarrow \text{N}^* + \text{H}^*$	0.01	-0.07

Table S2.3: **Validation of NN potential.** Difference between the activation energy E_a for all the dehydrogenation steps computed with our NN potential and DFT-PBE, according to Section 2.7.1.4. The maximum discrepancy obtained is 0.78 meV/atom , consistently with the expected accuracy of the potential

2.7.1.5 DFT calculations

All the single-point DFT calculations used to build the reference database were performed using the PWscf code of Quantum ESPRESSO (QE) [145, 189, 190]. Exchange-correlation effects have been treated within the generalized-gradient approximation with the Perdew-Burke-Ernzerhof (PBE) functional [79]. We employed ultrasoft RKKJ pseudopotentials [191] along with plane-wave basis with a 640 and 80 Ry cutoff for wave-function kinetic energy and charge density, respectively. The Marzari-Vanderbilt cold smearing technique [192] with a Gaussian spreading of 0.04 Ry was used to treat the state occupations. Collinear-spin polarization was included.

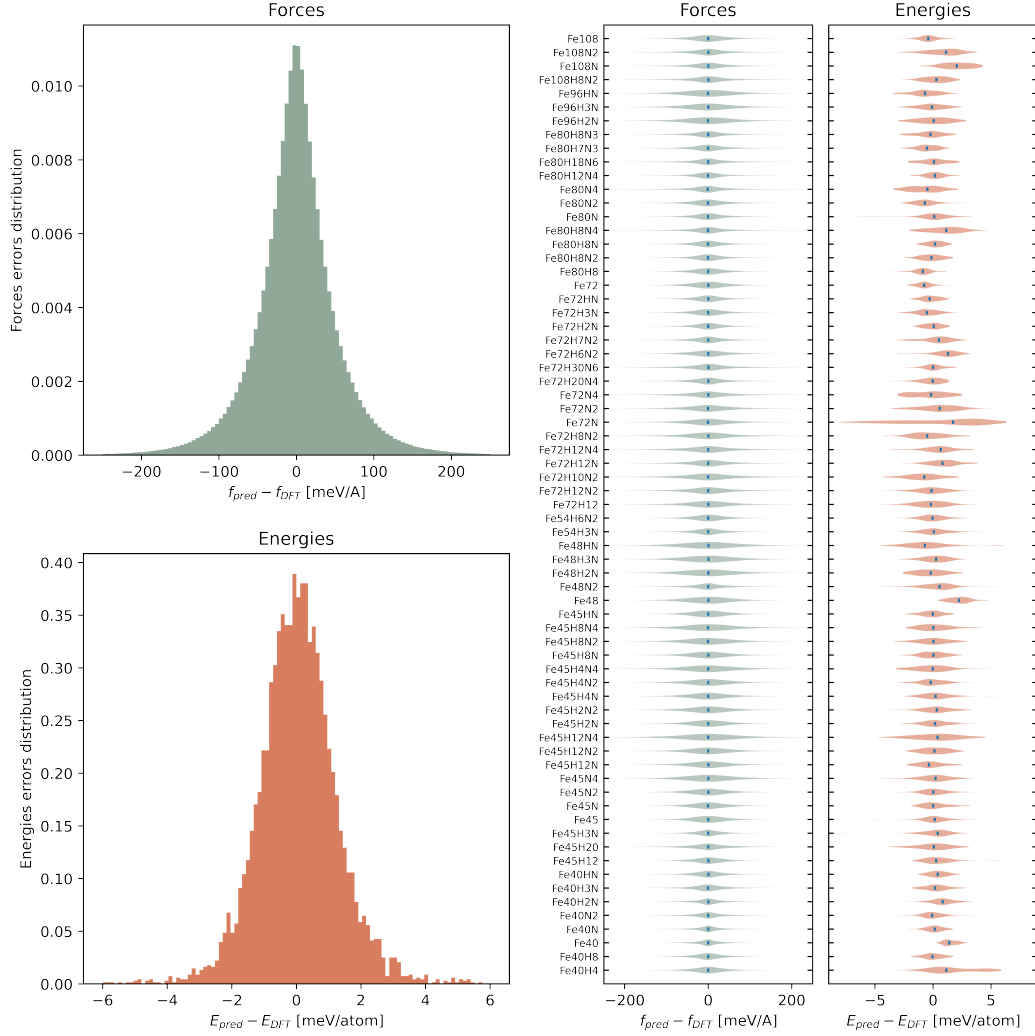


Figure 2.7: **Accuracy of energy and force predictions.** Left panel, the histogram of the errors in the forces (top) and energies (bottom) in the test set. Right panel, we decomposed the error distributions according to the chemical formula. The violin plot represents the error distributions on the forces and energies of the configurations grouped by chemical formula. The lines represent the mean value for each system. The test set contains 100 configurations for each formula. The average root mean square errors (RMSE) are $55 \text{ meV}/\text{\AA}$ and $1.3 \text{ meV}/\text{atom}$ for forces and energies, respectively. We choose to test on an unbalanced set with respect to the training one to assess the reliability of the potential on the different scenarios. The RMSEs on the balanced test on the studied system are $46 \text{ meV}/\text{\AA}$ and $1.0 \text{ meV}/\text{atom}$ for forces and energies, respectively.

In the QE program, the starting spin polarization of iron was set to 0.6 (site magnetization per valence electron). Thus, only ferromagnetic iron structures were considered. To mimic bulk-terminated structures, the lowest two Fe layers were kept fixed during optimization and MD calculations for (100) and (110) surfaces, and the lowest three layers for the (111) and (211) surfaces. A vacuum layer of at least 10\AA was included in the z-direction of all the slab models to prevent a self-interaction

effect. The Brillouin zone was sampled using a $2 \times 2 \times 1$ Monkhorst-Pack k-point grid[193].

The preliminary AIMD simulations were performed in a constant volume and temperature (NVT) ensemble using the stochastic velocity rescaling thermostat[78] with a time step of 1.0 fs at temperatures ranging from 600 to 800 K. In addition, enhanced sampling simulations (see Section 2.4) were performed to simulate the reactive events and to explore a broader configurational space.

2.7.1.6 CV and OPES specification

We report here the specific parameters of the CV and OPES parameters used in the study of NH_3 decomposition.

The set of CVs implied both as bias-CVs as analysis tool is the coordination number between the nitrogen atom and iron/hydrogen atoms, as defined in Eq. 2.29 for the definition). When used as a CV with OPES, the parameters are set as follows: $r_0 = 3.0$ (0.8) Å, $d_0 = 0.0$ (0.7) Å, $n = 6$ (5), and $m = 10$ (7) for $C_{N,Fe}$ ($C_{N,H}$). In the analysis context, the following set of parameters is adopted for the $C_{N,Fe}$: $r_0 = 1.5$ Å, $d_0 = 1.0$ Å, $n = 6$, and $m = 12$. This latter choice allows us to identify a sharper coordination number related to only the first neighborhood shell.

Analogously, for inspecting the dehydrogenation events and their transition state (see Section 2.6.1), we define the coordination between the reactive H and Fe atoms $C_{H,Fe}$ with the following parameter: $r_0 = 1.5$ Å, $d_0 = 0.5$ Å, $n = 6$, and $m = 12$.

To study the diffusion of NH over Fe(110), we use as CVs $v_{1\bar{1}1}$ and $v_{1\bar{1}\bar{1}}$. We set OPES barrier parameter ΔE equal to 40 kJ mol^{-1} and the OPES bias was updated every 500 steps. Harmonic restraints with an elastic constant of $10000 \text{ kJ mol}^{-1}$ were applied on both the two CVs when $|v| \geq 1.6$. In such a way, we restrict the number of accessible metastable states without reducing the size of the supercell. The same restraints were applied to study the diffusion of NH_3 and NH_2 , while no biases were applied due to the small diffusion barriers.

To simulate the first dehydrogenation step we use a set of 3 CVs: the $C_{N,Fe}$, the $C_{N,H}$, and Θ . The OPES barrier parameter was set at 60 (80) kJ mol^{-1} for the (110) ((111)) surface and the OPES bias was updated every 1000 steps. A harmonic restraint was applied when $C_{N,H} \leq 2.5$, utilizing an elastic constant of 2000 kJ mol^{-1} . This restraint, which acts immediately after the dehydrogenation occurs, is aimed at facilitating reversible sampling, eliminating the need to wait for hydrogenation to take place.

For the following dehydrogenation steps, we use as biased CVs the $C_{N,Fe}$, and the $C_{N,H}$ coordination numbers. The OPES barrier parameter was set at 40 (90) kJ mol^{-1} and

80 (50) kJ mol⁻¹ for the second and third dehydrogenation steps over the 110 (111) surface, respectively. The bias was updated every 1000 steps in all the simulations and a harmonic restraint with an elastic constant of 2000 kJ mol⁻¹ was applied when $C_{N,H} \leq 1.5$ or $C_{N,H} \leq 0.5$.

2.7.2 Effect of N^* on Fe(111)

2.7.2.1 Atomistic models

We modeled the (111) iron surface with a slab of a BCC oriented along the: $a = [\bar{1}\bar{1}0](x)$, $b = [11\bar{2}](y)$, and $c = [111](z)$. This model consisted of 768 Fe atoms distributed across 16 atomic layers in the x-direction, 24 atomic layers in the y-direction, and 12 atomic layers in the z-direction, see Figure 2.8(a). The bottom three layers were kept fixed at their equilibrium positions. We studied three different concentrations of N^* on the (111) surface: 12.5%, 25%, and 50% of a Fe monolayer. Figure 2.8 describe the initial structures, where the N^* were arranged in ordered patterns. In the case of 25% concentration, the N^* arrangement is the one suggested by Bozso et al. [24].

2.7.2.2 DFT calculations

Density functional theory (DFT) calculations were performed with the same parameters as the previous study, see Section 2.7.1.5.

2.7.2.3 ML interatomic potential details

To perform long-timescales and large system sizes simulation, we trained a machine learning potential on an extensive dataset comprising *ab initio* molecular dynamics and single-point DFT calculations. We started from the potential developed by Bonati et al. [44] and we expanded it via an active learning strategy to describe also high N^* and H_2 coverages, up to 100% of a monolayer on the Fe(111). For the fitting of the potential, we used the Deep Potential Molecular Dynamics Smooth Edition scheme [107, 186], as implemented in the DeePMD-kit software [110] using the same settings of Ref. [44]. To collect the relevant configurations to be labeled with DFT single points, we used an active learning scheme combined with enhanced sampling methods. We used a *query-by-committee* algorithm (Section 2.3.1), in which, at each iteration, five potentials were optimized on different permutations of the training and validation datasets. Then, MD simulations were performed with one of the potentials and the standard deviation of the predictions was used as a proxy for the uncertainty. This allows us only to select configurations for DFT labeling that are not already

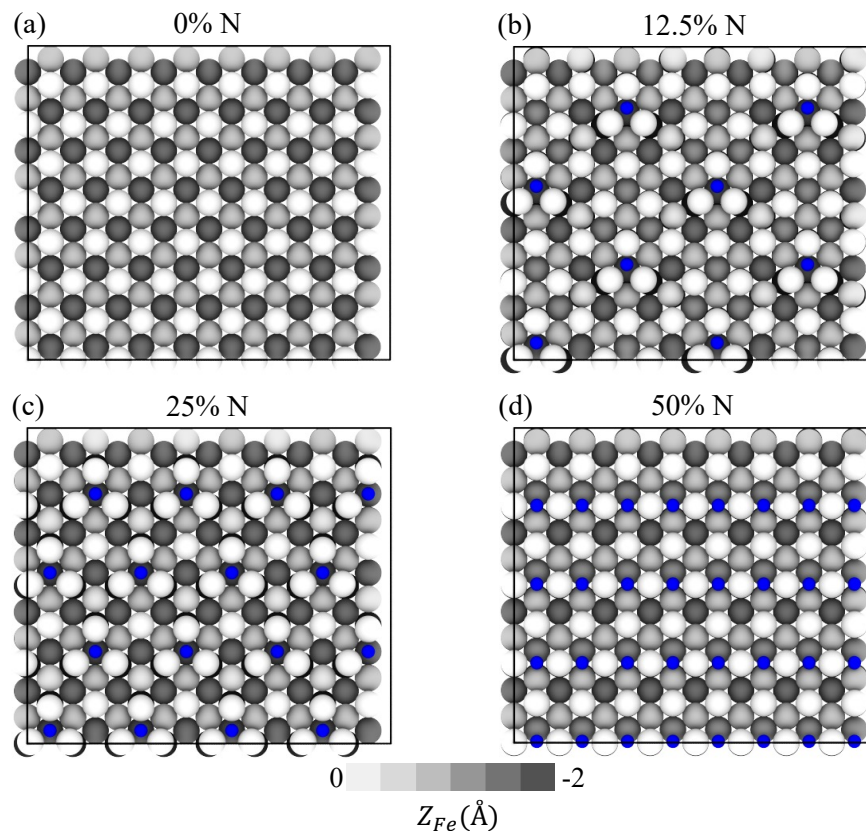


Figure 2.8: Snapshots from the top view of the $0K$ minimized Fe (111) structures with various N^* concentrations: a) 0 %, (b) 12.5 %, (c) 25 %, and (d) 50 % of the monolayer Fe concentration along the $[111]$ direction. Fe atoms are color-coded based on their height along the $[111]$ axis, as indicated by the color bar, while N atoms are depicted in blue.

well described by the potential. This strategy was used with both standard MD simulations and enhanced sampling ones. The usage of enhanced sampling simulations is crucial for this stage to ensure sufficient sampling of all the thermodynamically relevant configurations, including molecules adsorption/desorption, dissociation, and atomic species diffusion. It also enables to collect configurations not only from local minima but also along the entire reactive pathways [44, 125, 126, 129]. The active learning procedure was iterated eight times until we had a reliable potential. The final potential was trained on more than 200k configurations; the detailed composition is reported in Table 2.4. Figure 2.9 shows the distribution of errors in forces and energies on the validation sets of pure and nitrogenated Fe(111) without the N_2 ($Fe - N^*$) and with N_2 ($Fe - N^* - N_2$). The final potential achieves a root mean squared error (RMSE) of 1.3 meV/atom for the energy and 54 meV/Å for the forces

on the validation set of $Fe - N^*$. In the validation set of $Fe - N^* - N_2$, we obtained an RMSE of 1 meV/atom for energy and 50 meV/Å for forces.

System	Surface	Training	Validation
Fe	(111)	10600	3700
$Fe - 11\%N^*$	(111)	7800	2200
$Fe - 22\%N^*$	(111)	7200	2000
$Fe - 25\%N^*$	(111)	4000	1500
$Fe - 44\%N^*$	(111)	3600	1200
$Fe - 50\%N^*$	(111)	2000	500
$Fe - 75\%N^*$	(111)	3200	800
$Fe - 100\%N^*$	(111)	2900	700
$Fe - N_2$	(111)	5900	1500
$Fe - 75\%N^* - N_2$	(111)	2100	400
$Fe - 22\%N^* - H_2 - N_2$	(111)	4400	800
$Fe - H^*$	(111)	10000	2400
$Fe - N^* - H^*$	(111)	33200	11800
$Fe - H^* - N_2$	(111)	16300	400
$Fe - N^* - H^* - NH_x$	(111)	17800	5100
Fe	(110)	800	200
$Fe - N^*$	(110)	24300	7200
$Fe - H^*$	(110)	13800	3400
$Fe - N^* - H^*$	(110)	7200	1700
$Fe - N_2$	(110)	4400	1300
$Fe - N^* - N_2$	(110)	2600	600
$Fe - H^* - N_2$	(110)	6000	1500
$Fe - N^* - H^* - N_2$	(110)	2900	600
$Fe - N^* - H^* - NH_x$	(110)	11500	3500

Table 2.4: Composition of the dataset of *ab initio* calculations used to train the machine learning potential. % content of N^* refers to the percentage of monolayer Fe concentration along the surface direction. NH_x is equivalent to $NH_3 - NH_2 - NH$

2.7.2.4 MD simulations

Once we obtained a reliable interatomic potential, we utilized the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) [142] program, integrated with DeepMD-kit 2.1 [110] and PLUMED [144], to perform MD simulations. Simulations were performed using a constant volume and temperature (NVT) ensemble at 300 K and 700 K with a timestep of 0.5 fs. The stochastic velocity scaling thermostat [78] with a coupling constant of 100 fs was used to control the temperature. To investigate the effect of N^* on surface dynamics, simulations with a duration of 50 ns were performed. The bottom three layers were held in a fixed position to establish a




Figure 2.9: Histogram of the error in forces (left) and energies (right) on the validation set (20%) of $Fe - N^*$ and $Fe - N^* - N_2$. The histogram is normalized with respect to the total number of data points. For forces, each component is considered separately. Please refer to the Table 2.4 for the various N^* coverages used in the training and validation data sets.

boundary condition that simulates a semi-infinite slab. Periodic boundary conditions were applied in the x- and y-directions.

CV specification To build the OPES bias we used the N-N interatomic distance (d) and the coordination number between Fe and N atoms ($C_{N,Fe}$) as CVs, which were identical to those used by Bonati et al. [44] in their study on N_2 adsorption and dissociation on clean Fe(111) surface. The coordination number was calculated following Eq. 2.29. The chosen parameters were as follows: $r_0 = 2.5 \text{ \AA}$, $n = 6$, and $m = 12$. The update of the OPES bias was performed every 1000 steps, with an initial kernel width of 0.025 for d and 0.25 for $C_{N,Fe}$. The barrier parameter was set at 80 kJ/mol. A harmonic restraint was applied when $d \geq 2 \text{ \AA}$, utilizing an elastic constant of 2000 kJ/mol/ \AA^2 . This restraint aimed to facilitate the reversible sampling of adsorption states and the dissociation barrier, eliminating the need to wait for recombination to take place. MD simulations within OPES framework were performed for 25 ns and 75 ns at 300 K and 700 K, respectively, for all N^* concentrations.

2.7.3 The fate of nitrogen on Fe(110)

2.7.3.1 MD Simulations

MD simulations were performed using the MLP trained on the expanded database, see Section 2.7.2.3.

A $(5 \times 8 \times 10)$ slab of Fe, with Miller indexes (110) and 400 atoms was simulated, where the three bottom layers were kept fixed to simulate a bulk-terminated surface,

and periodic boundary conditions were applied on the x and y directions. We also repeated the simulations with a larger system of 800 atoms to rule out possible finite-size effects. Still, the initial simulation to study the interaction of nitrogen and hydrogen with the surface (4.1) was performed on the larger slab. MD simulations were performed using the LAMMPS MD software [142], using the canonical sampling through velocity rescaling thermostat with a target temperature of 700 K [78], an integration time step of 0.5 fs and a length of 10 ns.

2.7.3.2 CV and OPES specifications

To study the diffusion of N and H over Fe(110), we use as CVs $v_{1\bar{1}1}$ and $v_{1\bar{1}\bar{1}}$, as we have done for NH . For N , we set OPES barrier parameter ΔE equal to 45 kJ mol^{-1} , and the OPES bias was updated every 500 steps. Harmonic restraints with an elastic constant of 1000 kJ mol^{-1} were applied on both the two CVs when $|v| \geq 1.6$. In such a way, we restrict the number of accessible metastable states without reducing the size of the supercell. The same restraints were applied to study the diffusion of H , while no biases were applied due to the small diffusion barriers.

Similarly, the same OPES parameters were adopted to study the diffusion of N and H in the presence of another N^* . Here, for an effective sampling, we applied a reflective wall at $|v_{hkl}| < 2.55$, $hkl = 1\bar{1}1, 1\bar{1}\bar{1}$, displayed in Figure 4.5 with a black line. Similarly, to avoid the diffusion of the central nitrogen, we applied another wall at $|v_{hkl}| < 0.55$, $hkl = 1\bar{1}1, 1\bar{1}\bar{1}$.

To study the nitrogen recombination, we use the N-N distance as a bias CV. Here, some particular wall based on the lattice position was adopted to confine the N^* in the precursor state of the recombination. See Section 7.1.2 for the details. The OPES barrier parameter was set to 200 kJ mol^{-1} with a pace of 100 steps. The same parameters were adopted to perform the simulation at finite monolayer coverage except for the barrier parameter that was decreased to 195 kJ mol^{-1} and 190 kJ mol^{-1} for the coverage of 20% and 40%, respectively.

Finally, we exploited the CV defined in Equations 2.31 to study the nitrogen migration inside the bulk. Specifically, $v_{0\bar{1}0}$ was set as bias-CV, with an OPES barrier parameter of 170 kJ mol^{-1} and a pace of 100 steps. To constrain the simulation to a single “diffusion channel”, walls are placed when $\|\tilde{v}_{1\bar{1}1}\| > 0.55$ and $\|\tilde{v}_{1\bar{1}\bar{1}}\| > 0.55$, with the same harmonic parameter adopted for surface diffusion. Analogously to nitrogen recombination studies, OPES barrier parameters were decreased to 160 kJ mol^{-1} and 150 kJ mol^{-1} as coverage increased.

2.7.4 Ammonia decomposition on iron-cobalt

2.7.4.1 DFT calculations

All DFT calculations used to build the reference database were performed using the PWscf code of Quantum ESPRESSO [145, 189, 190]. Exchange-correlation effects were treated using the generalized-gradient approximation with the Perdew-Burke-Ernzerhof (PBE) functional [79]. We employed ultrasoft pseudopotentials selected from the SSSP PBE Precision v1.3.0 library [194] with 1, 5, 16, and 17 valence electrons for H, N, Fe, and Co, respectively. The wave function and the charge density cutoff were set at 90 and 1080 Ry. The Marzari-Vanderbilt cold smearing technique [192] with a Gaussian spreading of 0.04 Ry was used to treat the collinear spin-polarised electronic occupations. Initial spin polarizations were set to 0.6 for Fe|Co. To ensure the consistency of the dataset, spurious high-energy non-ferromagnetic configurations were removed from the dataset. A vacuum layer of at least 12 Å was included in the z-direction of all the slab models to prevent self-interaction effects. The Brillouin zone was sampled using a Monkhorst-Pack k-point grid [193] with a maximum k-spacing of 0.25 \AA^{-1} , and the SCF convergence threshold was set to 10^{-6} Ry.

2.7.4.2 GP specification (FLARE)

The local environments are described through the ACE B2 descriptors, using a basis expansion with $N_{max} = 8$, and $l_{max} = 3$. We used a cutoff equal to 5.5 Å for Fe and Co and 4 Å for the interaction with the other species. The kernel used to compare the environments was a normalized dot product, squared.

2.7.4.3 GNN specification (MACE)

All the GNN models in this work were constructed using MACE version 0.3.0 [195]. The models employed a cutoff radius of 6 Å for the atomic environments with 256 channels and $L=0$, finding a balance between accuracy and computational efficiency. The dataset was split into training/validation subsets with a ratio of 95:5, and the model was optimized with AMSGrad using a learning rate of 0.01 and a batch size of 4. The performance was evaluated on energy and forces with a weighted root mean square error (RMSE) loss function. The weights of the energy and forces were initially set to 1 and 100, respectively. In the last $\sim 25\%$ epochs, the weight of energy has been increased by a factor of 10 three times. The optimization was interrupted via an early stopping criterion with a patience of 200 epochs.

2.7.4.4 Datasets compositions

The dataset was composed first via incremental and then through an active learning scheme (DEAL) in a highly data-efficient way. The protocol adopted during the construction of the dataset employed to study the decomposition of NH_3 on FeCo(110) is detailed in Chapter 6. Then, the dataset was expanded to study lateral interactions using DEAL as described in Section 7.3.1. This results in a final dataset that comprises about 8200 configurations. The detailed composition is reported in Table 2.5, while its composition during each phase of the construction will be reported later in Tables 6.1 and 7.2.

2.7.4.5 Molecular dynamics

MD simulations were performed using the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) software [142], supplemented by FLARE [196], MACE v0.3.0 [195], and PLUMED v2.9 [144]. The canonical sampling through velocity rescaling thermostat [78] with a coupling constant of 50 fs was employed to control the temperature in all the simulations. During the construction of the preliminary reactants potentials, we simulated FeCo slabs cut along the (1 1 0) surface containing 60 (3x4x5) and 144 (4x6x6) atoms. A vacuum region of at least 12 Å, was added in all configurations to avoid self-interaction effects. The slab was simulated in both bulk-terminated (by fixing the two bottom layers) and double-interface (by fixing a central layer) conditions. In the latter, adsorbate molecules were placed on both sides. The surface dynamics was simulated at 700 K and then 800 K, while the interaction with the different reaction intermediates (NH_3 , NH_2 , NH , N/N_2 , H) was simulated at 700 K. At this stage, atomic interactions were described by FLARE, and the timestep was set to 1 fs. In the reaction discovery stage, as well as during the GNN active learning and the final simulations, we simulated 120 atoms (4x6x5) bulk-terminated (2-fixed layers) FeCo(110) slabs at $T = 700$ K. Periodic boundary conditions were applied in the x and y directions, while a reflecting wall was implemented over the surface in the z-direction. For the final simulations, we repeated the simulations with larger systems (up to 800 atoms), using the identical setups we used in the case of Fe for the sake of comparability.

2.7.4.6 CV and OPES specification

In the preliminary step, we used OPES to enhance a) the exploration of the different adsorption states by using as CV the coordination between N and Fe/Co atoms and b) the diffusion of the adsorbed molecule by using the x and y components of its center of mass.

FeCo original dataset composition

Formula	Adsorbate	N. config.	Formula	Adsorbate	N. config.
Co ₃₀ Fe ₃₀ H ₂	H	251	Co ₆₀ Fe ₆₀ HN	NH	618
Co ₃₀ Fe ₃₀ H ₂ N ₂	NH	301	Co ₆₀ Fe ₆₀ N	N	10
Co ₃₀ Fe ₃₀ H ₄ N ₂	NH ₂	173	Co ₆₀ Fe ₆₀ N ₂	2 N	84
Co ₃₀ Fe ₃₀ H ₆ N ₂	NH ₃	125	Co ₆₀ Fe ₆₀ N ₂	N	17
Co ₃₀ Fe ₃₀ N ₂	N	275	Co ₆₀ Fe ₆₀ N ₂	N ₂	86
Co ₃₀ Fe ₃₀ N ₄	2 N	157	Co ₆₀ Fe ₆₀ N ₄	2 N	53
Co ₃₀ Fe ₃₀ N ₄	N ₂	140	Co ₆₀ Fe ₆₀ N ₄	N ₂	9
Co ₆₀ Fe ₆₀ H ₂	H	13	Co ₇₂ Fe ₇₂	–	130
Co ₆₀ Fe ₆₀ H ₂ N	NH ₂	1236	Co ₇₂ Fe ₇₂ H ₄ N ₂	NH ₂	130
Co ₆₀ Fe ₆₀ H ₂ N ₂	NH	25	Co ₇₂ Fe ₇₂ H ₆ N ₂	NH ₃	129
Co ₆₀ Fe ₆₀ H ₃ N	NH ₃	886	Co ₇₂ Fe ₇₂ N	N	⁽¹⁾ 181
Co ₆₀ Fe ₆₀ H ₄ N ₂	NH ₂	8	Co ₇₂ Fe ₇₂ N ₂	2 N	130
Co ₆₀ Fe ₆₀ H ₆ N ₂	NH ₃	29			
Total					5193

Dataset expansion for later interaction study

Formula	N. config.	Formula	N. config.
Co ₆₀ Fe ₆₀ H ₁₂ N	29	Co ₆₀ Fe ₆₀ H ₈ N	28
Co ₆₀ Fe ₆₀ H ₁₂ N ₂	34	Co ₆₀ Fe ₆₀ H ₈ N ₂	14
Co ₆₀ Fe ₆₀ H ₂ N ₃	3	Co ₆₀ Fe ₆₀ H ₉ N ₄	45
Co ₆₀ Fe ₆₀ H ₂ N ₄	15	Co ₆₀ Fe ₆₀ H ₉ N ₅	568
Co ₆₀ Fe ₆₀ H ₃ N ₂	2	Co ₆₀ Fe ₆₀ N ₁₀	7
Co ₆₀ Fe ₆₀ H ₃ N ₃	26	Co ₆₀ Fe ₆₀ N ₁₃	7
Co ₆₀ Fe ₆₀ H ₃ N ₅	54	Co ₆₀ Fe ₆₀ N ₆	3
Co ₆₀ Fe ₆₀ H ₄ N	3	Co ₆₀ Fe ₆₀ N ₉	8
Co ₆₀ Fe ₆₀ H ₄ N ₂	4	Co ₇₂ Fe ₇₂ H ₂ N ₃	3
Co ₆₀ Fe ₆₀ H ₄ N ₅	9	Co ₇₂ Fe ₇₂ H ₃ N ₂	3
Co ₆₀ Fe ₆₀ H ₄ N ₆	170	Co ₇₂ Fe ₇₂ H ₃ N ₄	25
Co ₆₀ Fe ₆₀ H ₅ N ₇	301	Co ₇₂ Fe ₇₂ H ₄ N ₅	82
Co ₆₀ Fe ₆₀ H ₆ N ₃	14	Co ₇₂ Fe ₇₂ H ₅ N ₆	277
Co ₆₀ Fe ₆₀ H ₆ N ₄	103	Co ₇₂ Fe ₇₂ H ₆ N ₃	44
Co ₆₀ Fe ₆₀ H ₆ N ₇	83	Co ₇₂ Fe ₇₂ H ₆ N ₇	234
Co ₆₀ Fe ₆₀ H ₆ N ₈	309	Co ₇₂ Fe ₇₂ H ₇ N ₄	47
Co ₆₀ Fe ₆₀ H ₇ N ₅	99	Co ₇₂ Fe ₇₂ H ₉ N ₄	376
Total		Total	3029

Table 2.5: **FeCo dataset composition.** Composition of the configuration in the dataset for training and validation of the potential divided by chemical formula and adsorbate type. The first block refers to the original data collected following what is described in Section 6.1. The second part refers to the database expansion for the study of lateral interaction, see Section 7.3.1.

⁽¹⁾: Configuration collected to study N migration inside the bulk. These configurations are not included in Table 6.1.

In the exploratory stage, we iteratively performed OPES-flooding simulations to harvest an ensemble of reactive pathways. For N_2 dissociation, we used the distance between the two N atoms $d_{N,N}$ and the coordination $C_{N,Fe|Co}$ as CVs, while for the hydrogenation steps, we used the inverse of the planar distance between N and one H atom. The probability estimate was updated every 50 steps, and the barrier was chosen to be 50 kJ mol^{-1} , in order to apply the bias only in the reactant region. Simulations were interrupted after the occurrence of the reaction, namely when $d_{N,N}$ exceeded 2.5 \AA , or when the coordination $C_{N,H}$ was decreased by 1 with respect to the equilibrium value.

During the active learning stage with MACE and the final production simulations, we used OPES with the standard well-tempered target distribution. To study nitrogen recombination, we use the distance $d_{N,N}$ as CV. The barrier parameter was set to 200 kJ mol^{-1} , and the bias was updated every 200 steps. For the first dehydrogenation step ($\text{NH}_3 \rightarrow \text{NH}_2 + \text{H}$), we biased the simulation along three CVs: C_{N-H} , $C_{N,Fe|Co}$, and the angular variable Θ defined in Eq. 2.32. For the second reaction ($\text{NH}_2 \rightarrow \text{NH} + \text{H}$), we used $C_{N,H}$ and $C_{N,Fe|Co}$, while for the last step ($\text{NH} \rightarrow \text{N} + \text{H}$), we employed the distance $d_{N,H}$ and $C_{N,Fe|Co}$. The barrier parameters were set equal to 80, 60, and 100 kJ mol^{-1} , respectively, and the bias was updated every 200 steps. Harmonic restraints were applied in all reactions to facilitate reversible sampling. For the $2\text{N} \rightarrow \text{N}_2$ process, the restraint was applied if $d_{N,N} \leq 1.2$ or $d_{N,N} \geq 3.2 \text{ \AA}$, while for the dehydrogenation steps if $d_{N,H} \leq 2.5 \text{ \AA}$, with an elastic constant of $5000 \text{ kJ mol}^{-1} \text{ \AA}^{-2}$.

To study the reactive pathways in detail, we perform a set of 100 OPES flooding simulations for each dehydrogenation step. The CV used for all the steps was the coordination $C_{N,H}$. To avoid biasing the TS region, we exclude from bias deposition the regions where $C_{N,H}$ is lower than 2.82, 1.85, and 0.9 for the three hydrogenation steps, using an OPES barrier parameter equal to 70, 30, and 80 kJ mol^{-1} , respectively. All simulations were stopped once the reaction occurred, terminating the simulation when the maximum NH distance exceeded 1.9 \AA .

It's important to notice that the tunable parameters of the coordination numbers r_0 , d_0 , n , and m are set differently for the dehydrogenation studies and the others. In the former case, for the coordination between N and Fe, Co, and Fe|Co, we set the parameters $d_0 = 1 \text{ \AA}$, $r_0 = 1.5 \text{ \AA}$, $n = 6$, and $m = 12$, while for the coordination between H and Fe, Co, and Fe|Co, we used $d_0 = 0.8 \text{ \AA}$, $r_0 = 1.0 \text{ \AA}$, $n = 6$, and $m = 12$. Finally, for the coordination between N and H, we used $d_0 = 0.7 \text{ \AA}$, $r_0 = 0.8 \text{ \AA}$, $n = 5$, $m = 7$. These are the analogous of $C_{N,Fe}$, $C_{H,Fe}$, and $C_{N,H}$ used in the study of NH_3 dehydrogenation on Fe. In all the other simulations and analyses related to N, for the coordination between N and Fe, Co, and Fe|Co, we used $d_0 = 0 \text{ \AA}$, $r_0 = 2.5 \text{ \AA}$,

$n = 6$, and $m = 12$.

The final studies on nitrogen recombination and migration into bulk were conducted using the same CVs as well as the same constraints of the studies on Fe. For the former reaction, the OPES barrier parameter was set to 195 kJ mol^{-1} , and the bias was updated every 100 steps. In the case of nitrogen migration inside, we use 220 kJ mol^{-1} as the barrier parameter and the same pace parameter. In the case of finite coverage (20% and 40% monolayer coverage), the barrier parameters for nitrogen recombination were reduced to 190 kJ mol^{-1} and 185 kJ mol^{-1} ; similarly, for the dissolution/segregation process the barrier parameters were decreased to 210 kJ mol^{-1} and 200 kJ mol^{-1} .

Chapter 3

Ammonia adsorption and decomposition on iron surfaces

Thermal decomposition, or catalytic cracking, is the technique of choice for generating hydrogen from ammonia [15]. In the following, we focus on iron, a prototypical thermocatalyst in both the Haber-Bosch process and its inverse, which is expected to be a better catalyst than the traditional optimal choices at low ammonia concentrations [20]. These reactions are known to be highly structure-sensitive, with Fe(1 1 0) and Fe(1 1 1) representing two limiting cases among the low-index surfaces. In fact, the former is characterized by high density and stability but low catalytic activity, while the latter has an open structure and high activity [24, 25, 28, 29, 31, 32].

Unfortunately, in spite of more than a century of efforts, a comprehensive characterization of both processes remains elusive due to extreme *operando* conditions which make both experiments and simulations challenging. Indeed, the operational temperatures for catalytic cracking on iron at 1 atm lie in the range of 400-750 °C [15]. Experimental information about the reaction intermediates has thus been obtained in conditions far from the *operando* regimes [28, 29], or extrapolated through indirect temporal analysis of products (TAP) measurements [30]. Notably, a high surface sensitivity has been reported [28, 29], similarly to what was observed in the synthesis process [32, 33]. Only very recently has direct *operando* probing of synthesis reaction intermediates been performed [34], confirming the structure sensitivity of the reactions and showing non-trivial temperature dependency in the rate-limiting step.

From a computational perspective, there is a large number of theoretical studies on ammonia decomposition based on T=0 K density functional theory (DFT) calculations [26, 30, 35–43]. However, these studies do not fully consider the influence of dynamics, which we already argued to be of paramount importance in industrial

catalysis [1, 54]. This crucial role is clearly emerging in recent simulations [44, 55, 57–62]. In particular, in the case of ammonia decomposition, Gerrits *et al.* have highlighted the role of dynamics on the Ru(0001) surface through *ab initio* molecular dynamics (AIMD) simulations [59].

In this Chapter, we employ the methodologies previously described to investigate the adsorption, diffusion, and dehydrogenation of nitrogen hydrides (NH_x , $x=1-3$) at the *operando* temperature of $T=700$ K on both the Fe(111) and Fe(110) surfaces. Here, we provide a comprehensive statistical and chemical characterization of these processes. An indispensable tool in the characterization of these processes is our ability to monitor the electronic charge distribution at modest computational effort also in large systems (see Section 2.5). This ability allows the adsorption environments and the progress of the reaction under evolving surface conditions to be followed.

All the results reported in this Chapter are obtained from MD simulations with the MLP, which was trained over 110 K configuration with DeePMD-kit [107]. The specifications on the MLP, as well as the details on simulation techniques, are reported in Section 2.7.1.

Given the challenges of studying these systems with MD simulations, we focus here on the interaction of one single molecule with the hot iron surfaces. Thus, we do not include the effect of lateral interactions [24, 25, 28, 29, 185] and reagent-induced surface reconstructions [63, 65, 197, 198]. These two topics will be discussed in Chapter 4 and 5.

We find that, as in the case of the ammonia formations, the two surfaces behave rather differently. The low and high temperatures are similar in the more rigid Fe(110) case, where an Arrhenius-type behavior is observed. In contrast, the high mobility of surface atoms of the less dense Fe(111) leads to the formation of new adsorption sites and results in a non-Arrhenius behavior. However, despite the different behavior that manifests itself in a variety of transition states and reaction pathways, a unified description can still be obtained.

3.1 Fe surface behavior at high temperature

Before discussing the ammonia decomposition-related problems, we briefly describe the behavior of pristine Fe surfaces. To study the reported surface dependence [28, 29, 32, 33], we simulate both the dense Fe(110) and the more open Fe(111) at the *operando* temperature of 700 K.

As this temperature is studied above the Hüttig temperature (≈ 550 K)[199], the mobility of surface atoms is expected, hence the diffusion of adatoms and vacancies. However, the energy for forming surface defects on Fe(110) is too high to be observed

spontaneously. Thus, the surface atomic arrangement is stable, and we do not observe any atomic diffusion (see Figure 3.1, left panel). We can anticipate that adsorption sites remain the same as those found in the low-temperature regime, as shown in Figure 3.2a. That is why we do not report this surface's 300 K behavior, which can be easily deduced from the *operando* temperature results.

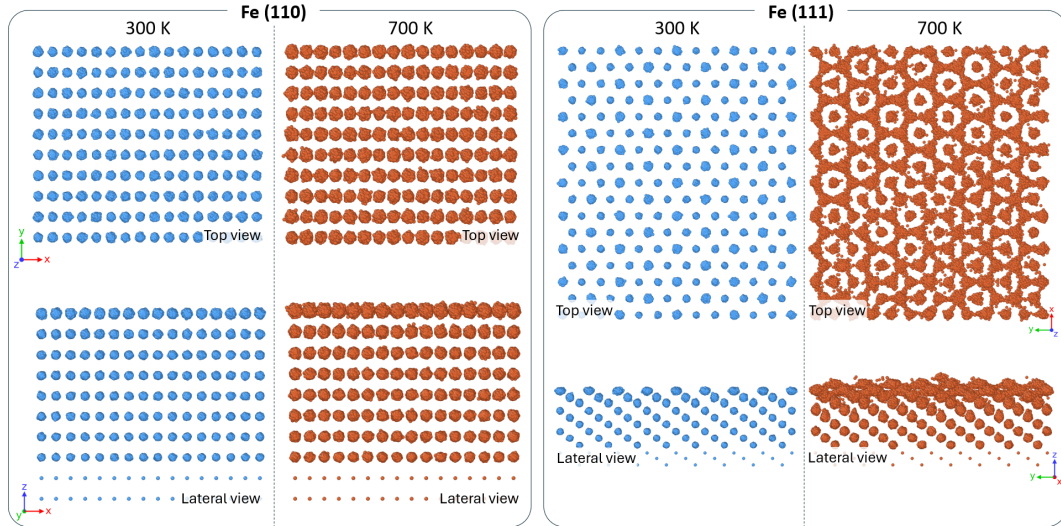


Figure 3.1: Scatter plot of atomic positions of Fe atoms for (1 1 0), on the left, and (1 1 1), on the right, surfaces at $T=300$ K (light blu) and 700 K (orange). The atom positions are recorded every 2.5 ps for 10 ns after 2 ns of equilibration. Both top and lateral views are reported.

In contrast, for Fe(1 1 1), the energy required for surface defect formation is lower, making it entropically favorable at a temperature of 700 K. As a result, the diffusion of surface atoms is observed, although the crystalline order of the surface is maintained, as discussed in Ref. [44]. This leads to the continuous formation and annihilation of surface defects, such as adatoms and vacancies. In Figure 3.1, we show the scattered positions of iron atoms over time, clearly illustrating the aforementioned phenomena. Furthermore, as we will examine it thoroughly in Chapter 5, the presence of vacancies on the surface allows the formation of triangular motifs, which are stabilized by adsorbed nitrogen (N^*). Thus, being the surface behavior rather different from the low temperature one, we also report the 300 K results for comparison. In particular, we find that in addition to the low-temperature adsorption sites (Figure 3.2b), several new ones are available at $T=700$ K, as shown in Figure 3.2c.

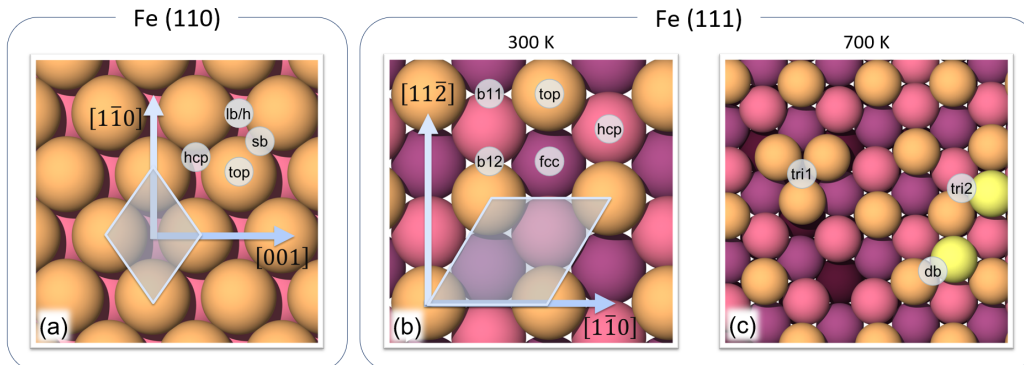


Figure 3.2: **Adsorption Sites.** Top view of Fe(110) (a) at 700 K, Fe(111) at 300 K (b) and 700 K (c). Frames are captured from unbiased ML-based simulations where the atomic positions are 10 ps averaged to suppress the thermal fluctuations. Iron atoms are colored by layer. Gray-shaded areas mark the primitive unit cell. White-gray labeled circles are used to denote high-symmetry adsorption sites. (a) For Fe(110): sb, short-bridge; lb, long-bridge; h, hollow. (b) For ideal Fe(111): b11, bridge between top-layer atoms; b12, bridge between first and second layer. (c) For high-temperature Fe(111), we highlight the formation of new adsorption sites due to the surface roughness: tri1, triangular reconstruction; tri2, high symmetry site between three consequential layers; db, defective b12. These last two sites exist only on the defective surface due to the absence of b11 site.

3.2 Interaction between NH_x and iron surfaces

We now discuss the results of simulations of the adsorption and diffusion of the intermediates of the reaction (NH_3 , NH_2 , and NH) on Fe(110) and Fe(111). In the dynamic approach adopted, adsorption and diffusion are intimately related and describe the catalyst's interaction with the adsorbed species. Two-dimensional free energy surfaces (FES) projections are used to explore the coordination numbers, charge transfers, and diffusion pathways.

3.2.1 Fe(110)

First, we present the results of the Fe(110). The most energetically stable adsorption configurations of nitrogen hydrides have already been investigated through DFT static calculations; top, short-bridge, and hollow sites have been computed as the minimum energy adsorption sites for NH_3 , NH_2 , and NH , respectively [30, 38]. At $T=0\text{ K}$, our results are in agreement with the previous theoretical work. In the following, we thus focus on the effects of temperature on this more rigid surface.

In Figure 3.3, we report the most probable adsorption state obtained from the minimum of free energy surfaces (FES) as a function of the nitrogen-iron coordination

number $C_{N,Fe}$, and the total charge transferred from the metallic substrate to the NH_x species q_{NH_x} . This allows the available adsorption sites to be characterized, from both the geometrical and chemical points of view. In Figure 3.4, the FES are projected along the two in-plane crystallographic directions $[001]$ and $[1\bar{1}0]$, from which we can analyze the diffusion mechanism and the related free energy barrier.

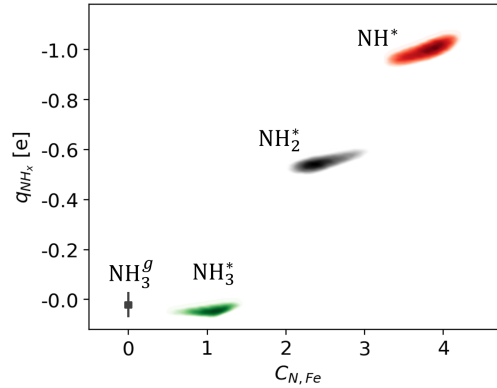


Figure 3.3: **Interaction of NH_x with $Fe(110)$ surface.** Most stable adsorption sites for the nitrogen hydrides on $Fe(110)$ surface at 700 K as a function of the nitrogen-iron coordination number $C_{N,Fe}$ and the total charge transferred from the metallic substrate to the NH_x species q_{NH_x} . Ammonia in both the gas and adsorbed phases is reported. For the adsorbed species, the profiles coincide with the minimum of their FES within the range $0-2 k_B T$ and are accordingly colored. For NH_3^g , the representation reduces to a point since $C_{N,Fe}=0$ in the gas phase.

NH_3 is preferentially adsorbed on top sites (1-fold coordination). No charge transfer from the substrate to the molecule is observed (Figure 3.3). However, the ammonia dipole induces a local redistribution of the iron charges, which in a first approximation can be described as a dipole induced-dipole interaction (Figure 3.4c). This weak bond allows the NH_3 to diffuse easily at 700 K through short bridge sites (Figure 3.4a-b). The computed free energy barrier for diffusing along this pathway is 0.15 eV. This diffusion mechanism competes with desorption and the subsequent physisorption process.

Contrary to NH_3 , NH_2 forms strong chemical bonds with the surface, preferring bridge to top sites. In these 2-fold coordination sites, the metal substrate transfers a charge of $-0.6 e$ to the molecule (Figure 3.3). Despite its strong chemisorption energy ($-3.17 eV$ [38]), NH_2 is the most mobile surface species. As shown in Figure 3.4a-b amide diffuses between short and long bridge sites through hpc sites (3-fold coordination) with an estimated free energy barrier of 0.05 eV, which is less than $1 k_B T$. Thus, we cannot speak of a single adsorption site due to NH_2 almost barrierless diffusion. The temperature effects play an important role even on these hard sur-

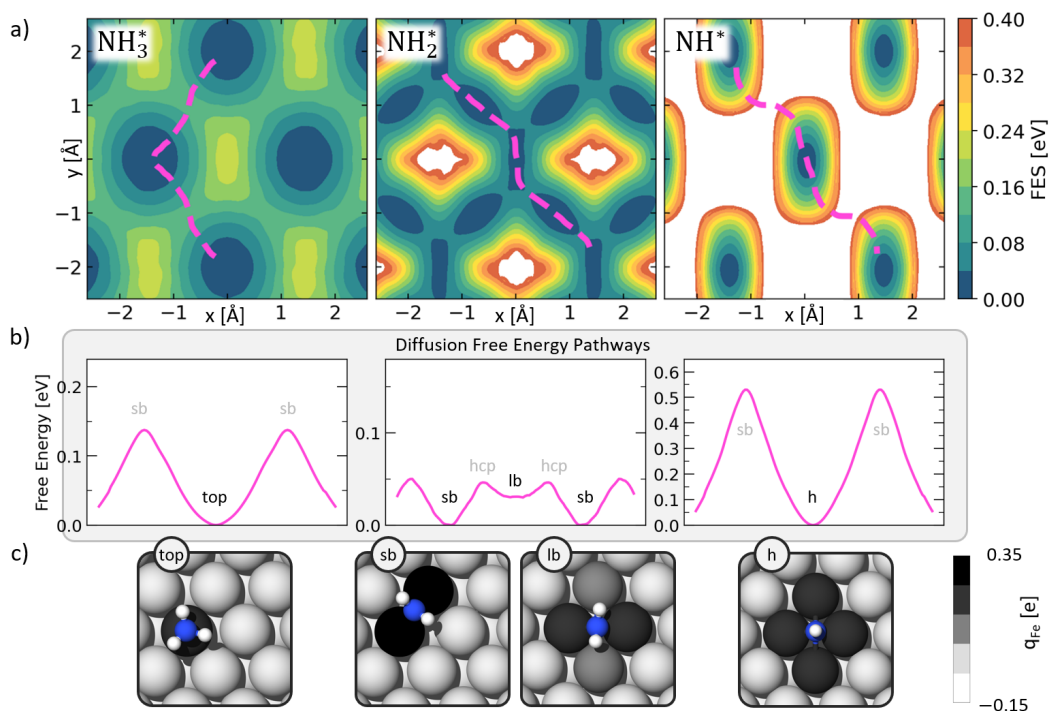


Figure 3.4: **Diffusion of NH_x^* on Fe(110) surface at 700 K.** (a) Free energy of NH_x^* interaction projected along the two crystallographic directions $[001]$ (x -axis) and $[1\bar{1}0]$ (y -axis). Local minima represent the metastable states, and fuchsia dashed lines denote the minimum free energy diffusion pathways. (b) Free energy calculated along the minimum free energy pathways. The high symmetry adsorption sites are labeled according to Figure 3.2. Different scales are used for the free energy to highlight the minima. (c) Snapshots of representative geometries of the minimum free energy adsorption states. Iron atoms are colored according to net charges predicted by the neural network model.

faces. We can anticipate that the absence of long-lived adsorption states implies a non-negligible increase in the free energy barrier for dehydrogenation.

Finally, NH lives on the hollow sites (4-fold coordination). In this configuration, as shown in Figure 3.3, there is a charge transfer of approximately $-1 e$ to the molecule. Among the nitrogen hydrides, NH is the least diffusive species. As shown in Figure 3.4a, imide diffuses between two adjacent hollow sites through a short bridge site. The computed free energy barrier for the diffusion is $0.54 eV$ (Figure 3.4b), making the diffusion of NH an activated process. For this species, we can thus speak a of long-lived ($\sim ns$) adsorption state.

Going from NH_3 to NH , there is an increase in interaction strength in terms of both the charge transfer and coordination (Figure 3.3). However, this is not correlated with the mobility of the intermediates (Figure 3.4).

3.2.2 Fe(111)

We now discuss the interaction of nitrogen hydrides with the Fe(111) surface, contrasting, as anticipated, the behavior at $T=300\text{ K}$ and $T=700\text{ K}$. In Figure 3.5, we show free energy surfaces as a function of the nitrogen-iron coordination number $C_{N,Fe}$, and the total charge transferred q_{NH_x} for each species in the two different temperatures. It is important to recall that at $T=700\text{ K}$, the surface iron atoms

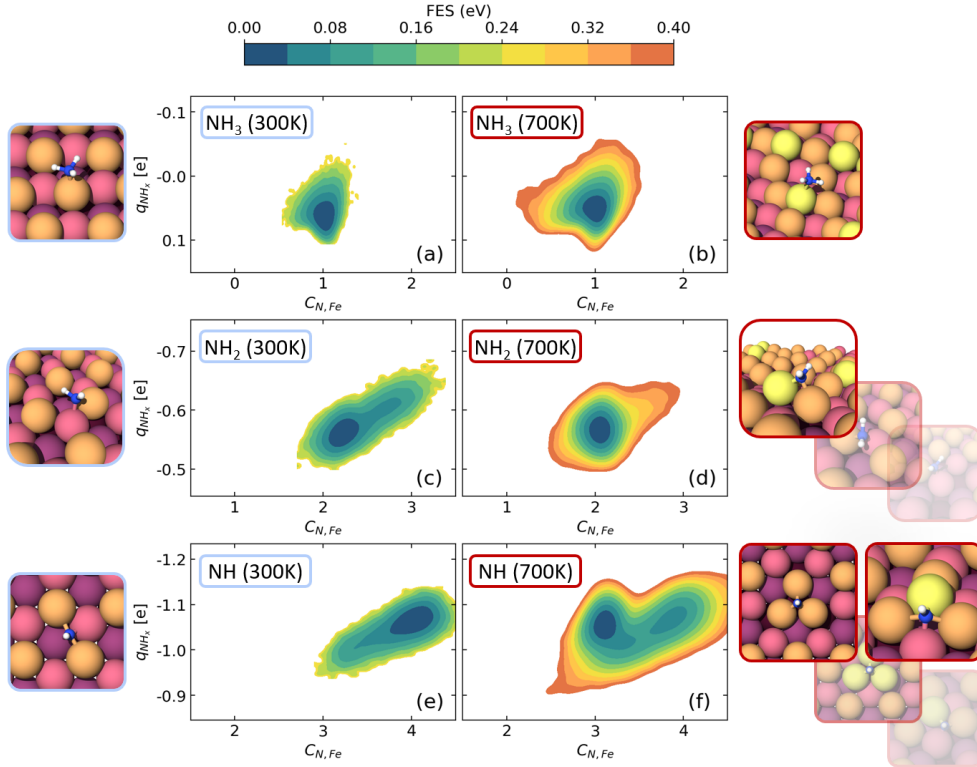


Figure 3.5: **Interaction of NH_x with Fe(111) surface.** Free energies as a function of the N-Fe coordination and the total charge transferred from the metallic substrate to the NH_x species at $T=300\text{ K}$ (a,c,e) and $T=700\text{ K}$ (b,d,f). In the snapshots, representative configurations belonging to the adsorption free energy minima are shown.

become mobile, forming additional adsorption sites (see Figure 3.2c). As shown in Figure 3.5, all the N exhibit a tetrahedral environment due to the drive of N to form sp^3 hybrid. As a result of surface atom mobility, sp^3 hybrid can be achieved in different ways in high-temperature regimes. In contrast, at low temperatures, such geometrical environments are not accessible. Such behavior is particularly important for the last two intermediates (NH_2 , NH) and has significant implications for dehydrogenation, as will be discussed below. As in Fe(110), NH_3 shows on-top adsorption at low and high temperatures, with no charge transfer associated. However, unlike

in Fe(110), NH_3 does not diffuse on the surface. This can be rationalized by considering that the shortest pathway between two on-top sites must visit the lower layers, passing via configurations less favorable than in the gas phase ones (see Figure 3.2).

As mentioned above, the adsorption of NH_2 is very different in the two temperature regimes. At $T=300\text{ K}$, the amide is bonded in bridge sites between the first and second layers (b12, Figure 3.2b). In this configuration, NH_2 can move almost barrierless ($\approx 0.04\text{ eV}$) between three equivalent b12 bridge sites via the top-layer bridge site (b11, Figure 3.2b). This behavior is illustrated in Figure 3.6 and explains the free energy ring minima between the two- and three-fold coordinated state of Figure 3.5c. At $T=700\text{ K}$, these low-temperature adsorption sites are no longer ob-

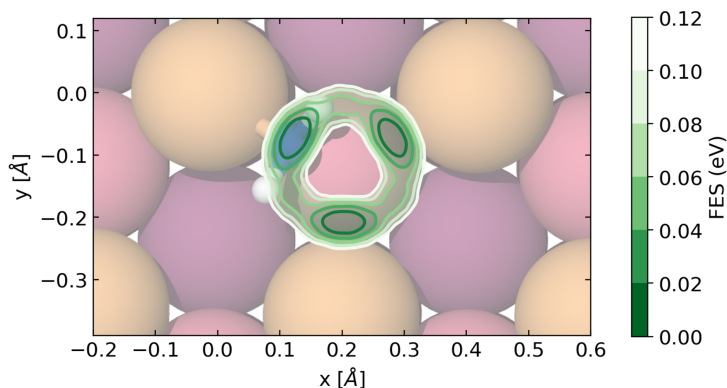


Figure 3.6: **Interaction of NH_2 with Fe(111), $T=300\text{ K}$.** Free energy(FES) of NH_2^* interaction with Fe(111) surface at $T=300\text{ K}$. FES is projected along the two crystallographic directions $[1\bar{1}0]$ (x -axis) and $[11\bar{2}]$ (y -axis). The three equivalent local minima (b12, see Figure 3.2b) are separated by free energy barriers of 0.04 eV . Thus, at low temperatures, NH_2 can diffuse between these three b12 sites across the b11, forming a unique “ring” shape adsorption site. At high temperatures, when the mobility of iron surface atoms breaks the surface symmetry forming defects, this adsorption pattern is no longer observed.

served, and NH_2 is placed only at defective bridge sites (db, Figure 3.2c). In these configurations, barrierless diffusion of NH_2 does not take place, and the adsorbate is stabilized in this two-fold coordinated state. The fluxionality of NH_2 at 700 K is significantly reduced with respect to 300 K . This counterintuitive effect can be explained by the availability of the isolated defective two-folded sites in which NH_2 is sp^3 hybrid.

Finally, we present the interaction of NH with Fe(111), which exhibits significant differences between the high- and low-temperature regimes. At 300 K , NH is adsorbed at bridge sites (b11, Figure 3.2b), while at 700 K , a variety of geometrical configurations is observed. The most stable sites for adsorption are the three-folded coordinated ones (tr1 and tri2, Figure 3.2c), but also four-fold coordination defective ones are still observed. Also on this surface, the iron substrate transfers a charge of

$\approx 1e$ to NH.

Although the mobility of iron surface atoms does not allow a quantitative analysis of the diffusion barriers similar to the one carried out for the (110) surface (Figure 3.4) to be performed, a detailed analysis of the free energy profiles can provide information on the dynamics of the reaction intermediates (Figure 3.5). The free energy surface reflects the probability that a certain state is occupied. In this low-dimensional representation, each state is projected to a point; and a point of the two-dimensional space is associated with one or more states. Thus, a broad free energy minimum reflects the fact that the system visits different environments. Using this tool, we can compare the free energy surface for NH_3 at the two temperatures (Figure 3.5a-b) which can be almost superimposed, reflecting a similar adsorption behavior. When we compare the two different temperatures, it is important to observe that at 700 K, the system can also visit higher free energy regions (yellow and orange areas of contours). In fact, the probability of visiting a state is exponentially related to its free energy and the inverse of the temperature (see Equation 2.18). If we now consider NH_2 , at 300 K (Figure 3.5c), the free energy minimum is placed in the coincidence of b12 sites (two-fold coordinated) with a pronounced low free energy tail that extends to high coordination region (in the occurrence of the b11 sites, Figure 3.2b). This is a result of the above-described adsorption behavior of NH_2 , which is illustrated in Figure 3.6. Instead, at 700 K (Figure 3.5d), the free energy minimum is well-defined in the two-fold coordinated region, with the other areas characterized by high free energy. Since diffusion between different two-fold coordination states necessarily requires going through low- or high-coordination states, this implies that diffusion proceeds via high free energy states. This is reflected in the fluxionality of NH_2 , which results in the paradoxical effect of being reduced at high temperatures, with more stable adsorption in the two-folded coordination states at the high temperature. A similar analysis can be performed for NH (Figure 3.5e-f). At 300 K, the adsorption free energy minimum presents a long tail in the low-coordination region, indicating a significant dynamism of this state. If we now consider the adsorption at 700 K, we find that NH can perform much longer excursions than NH_3 and NH_2 .

3.3 Dehydrogenation steps

3.3.1 Reaction mechanism analysis

As outlined in Section 2.6.1, one hundred simulations of dehydrogenation events were performed for each step and surface at *operando* temperature, using the OPES-flooding approach. The statistical sampling of reactive pathways was performed with-

out preconceived assumptions about the transition and final states.

A visual inspection reveals immediately a variety of reaction pathways, as we will elucidate in the next section. To get a unified perspective, we use, as in Ref. [44], the charge transfer to the nitrogen hydrides to describe the reaction processes. The evolution of the charges in the nitrogen moieties q_x during each dehydrogenation event occurring on both surfaces are plotted in Figure 3.7. We reasonably assume that the TS is closest to the point where the distance $d_{N,H}$ reaches the breaking point of 1.40 Å (see Section 2.6.1) and plot the charge variability on a segment of 40 fs centered around the TS.

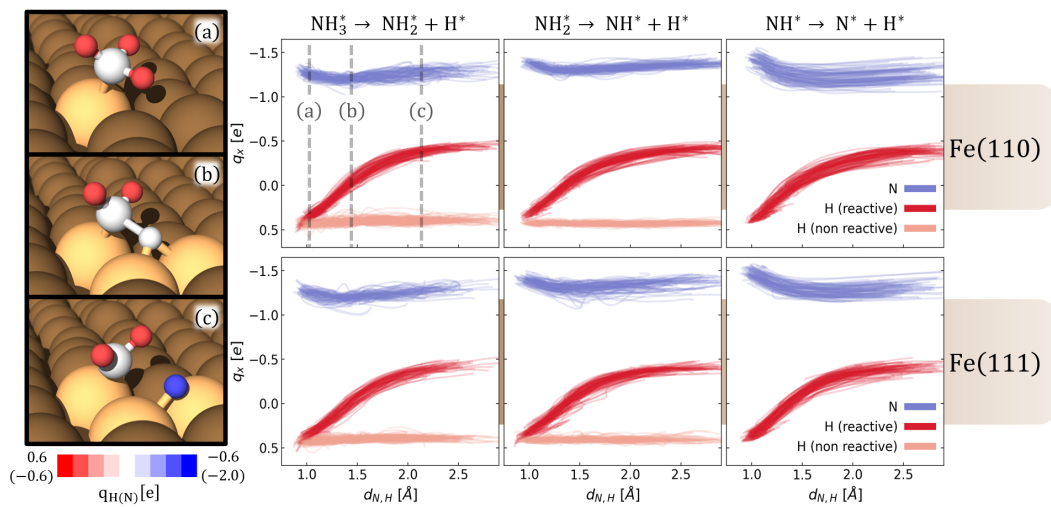


Figure 3.7: **Dehydrogenation charge variation.** Atoms in the molecule charge variation during reactions as a function of the distance between nitrogen and reactive hydrogen. Each line in the plot (one for each atom) corresponds to a reactive simulation (100 for each event). For each simulation, 40 fs were plotted around the transition state ($d_{N,H} = 1.40$ Å). On the left, we report three snapshots corresponding to the first dehydrogenation step on the Fe(110) surface.

It can be seen that throughout the reaction, a significant charge variation takes place only to the H atom that is being reduced. This distinctive behavior is exemplified in the left panel of Figure 3.7, which shows a representative first-step dehydrogenation event on the Fe(110) surface.

When the reactions are studied through these Bader charge lenses, a coherent underlying reaction mechanism becomes evident in spite of the fact that the reaction evolves following different geometrical trajectories. Strikingly, not only is this shared between the individual reactions but also across all the different steps and surfaces.

3.3.2 Transition state ensembles

We now analyze the observed variety in TS configurations. At the *operando* temperature, the adsorbed molecules and the surface atoms exhibit a pronounced mobility that gives each reactive event a distinctive geometric character. Therefore, it is not possible to identify a single transition state (TS); instead, it must be described by an ensemble of structures [118, 200, 201]. In Figure 3.8, we report the TS coordinations $C_{N,Fe}$ vs $C_{H,Fe}$ for all reaction steps on the Fe(110) and Fe(111) surfaces. These variables, which provide a coarse geometrical representation of the relevant reactive structure, exhibit a broad distribution.




Figure 3.8: **Transition state ensemble.** Distribution of the coordinations between nitrogen and iron, $C_{N,Fe}$, and between reactive hydrogen and iron, $C_{H,Fe}$, for the transition state configuration at 700 K. Each TS point is indicated by a cross.

Applying k-medoids analysis, we attempted to rationalize this complexity in a simple way and partition each TS ensemble into two clusters, as outlined in Section 2.6.1. The resulting clusters, along with their respective medoids, are depicted in Figure 3.9 for both Fe(110) and Fe(111) surfaces. In the top panels of Figure 3.9, we observe for Fe(110) at least two distinguishable classes of geometries, primarily differentiated by $C_{N,Fe}$ for the first and third dehydrogenation steps, and by the coordination number $C_{H,Fe}$ for the second dehydrogenation step. A parallel analysis is conducted for dehydrogenation steps on Fe(111). As depicted in the bottom panels of Figure 3.9, the first and second TS ensembles are classified based on $C_{N,Fe}$, while the third is distinguished by a combination of the two descriptors. It is interesting to note that NH_2 on Fe(110) is the only intermediate for almost all events that occur on a well-defined site (hpc, Figure 3.2a). The TS analysis identified two classes of reactive events based on how hydrogen is released on the surface. Dehydrogenation can occur on one of the three top atoms (characterized by a low $C_{H,Fe}$ coordination, blue cluster) or on one of the bridge sites (characterized by a high $C_{H,Fe}$ coordination, purple cluster). From an inspection of the transition states, we observe that

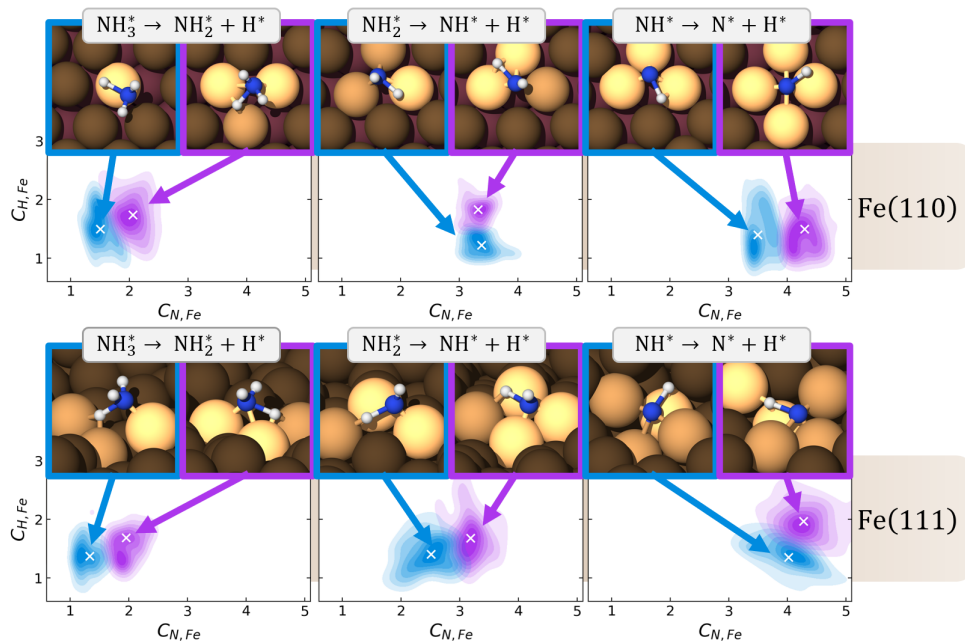


Figure 3.9: **Transition state ensemble analysis.** Distribution of the coordinations between nitrogen and iron, $C_{N,Fe}$, and between reactive hydrogen and iron, $C_{H,Fe}$, for the transition state configuration at 700 K. For each reaction, data are separated into two clusters according to the k -medoids analysis (see Section 2.6.1), in which color gradients represent density. The medoids are indicated with white crosses and shown in the snapshots. Raw data are reported in Figure 3.8. Iron atoms are colored to guide the eyes to the difference between clusters: $C_{H,Fe}$ for the second dehydrogenation step on the Fe(1 1 0) surface and the last one on the Fe(1 1 1) surface; $C_{N,Fe}$ for the others.

hydrogen binds to different bridge bonds, as on the different top atoms, regardless of the initial adsorption state (long or short bridge). It is noteworthy that the coordination number $C_{N,Fe}$ at the TS consistently exceeds the value assumed in the stable adsorption configuration of the reactant (refer to Figures 3.3 and 3.5 for Fe(1 1 0) and Fe(1 1 1), respectively). We have reported this behavior in Figure 3.10, showing a reaction segment of 40 fs (centered around the TS) in the two-dimensional space nitrogen-iron coordination $C_{N,Fe}$ total charge q_{NH_x} . As we will see below, since reactive events occur in the over-coordinated regime, the probability of each species reaching high-coordinated states is linked to the dehydrogenation barrier.

We want to emphasize once again that we only determined two clusters for each TS ensemble in order to obtain an economical description of the broad geometric complexity. This choice was meant to familiarize the reader with the great variety of TSs encountered and provide an overall description that encompasses all the reactions treated. A more detailed description of each reaction, such as looking at the electronic

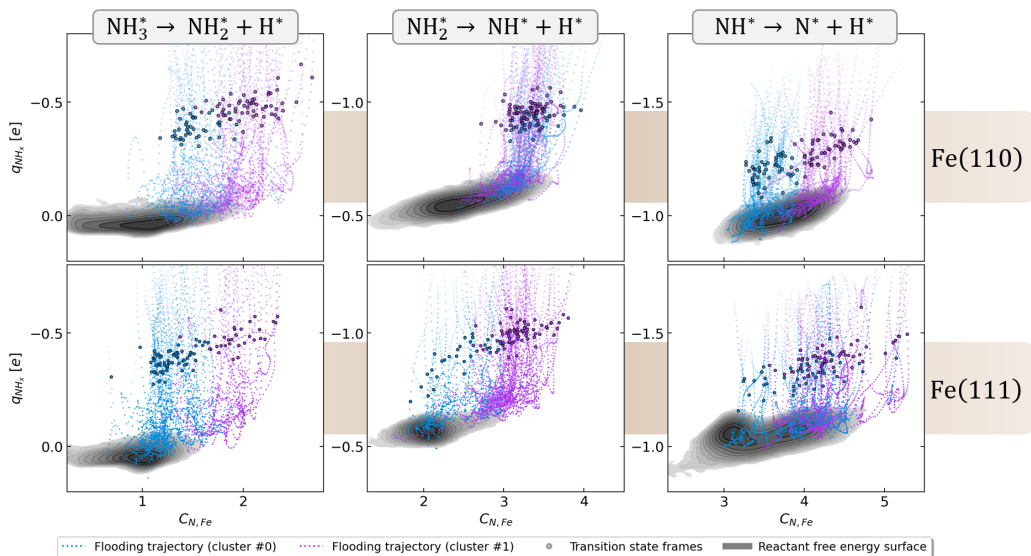


Figure 3.10: **Adsorption and dehydrogenation mechanism.** Reaction pathways for the three dehydrogenation steps for Fe(110), top panels, and Fe(111), bottom panels, at 700 K as a function of the nitrogen-iron coordination number $C_{N,Fe}$ and the total charge transferred from the metallic substrate to the NH_x species q_{NH_x} . Each line in the plot corresponds to a reactive simulation (100 for each event). For each simulation, 40 fs were plotted around the transition state ($d_{N,H} = 1.40 \text{ \AA}$), which are marked by dots. The trajectories are colored accordingly to clusters, following the k -medoids analysis reported Figure 3.8, with reducing intensity from the first to the last frame. To relate reactive trajectories with reactant adsorption, we include in the background the reactant free energy surface (see also, Figures 3.3 and 3.5), which is colored with grey gradients, from 0 up to $10 k_B T$ (0.6 eV) in 10 levels. Adsorption free energies are recovered from unbiased simulations, while reactive trajectories are extracted from OPES-flooding simulations.

occupation levels or orbital orientation, is beyond the scope of this work.

3.3.3 Dehydrogenation barriers

In Table 3.1 we present the free energy barriers associated with the dehydrogenation steps on both surfaces at the *operando* temperature. The values are obtained from the corresponding free energy profiles, depicted in Figure 3.11.

Since the reaction mechanism analysis shows that the underlying chemical process is the same for all reactive events, we would have expected similar reaction barriers for the various intermediates and surfaces. Instead, differences, even significant ones, are observed between the different nitrogen hydrides and the two different surfaces. We can attempt to rationalize these differences in the light of the various aspects of the dynamics described so far.

Firstly, we have observed how the different reaction intermediates have very different mobilities (see Section 3.2). In particular, for the Fe(110), we were able to

Free energy barriers ΔG^\ddagger [eV] at T=700 K		
	Fe(110)	Fe(111)
$\text{NH}_3^* \longrightarrow \text{NH}_2^* + \text{H}^*$	1.01	1.14
$\text{NH}_2^* \longrightarrow \text{NH}^* + \text{H}^*$	0.79	1.22
$\text{NH}^* \longrightarrow \text{N}^* + \text{H}^*$	1.25	1.10

Table 3.1: **Free energy barrier.** Estimated free energy barrier for each dehydrogenation step on Fe(110) and Fe(111) surface at T=700 K. The values are derived from the associated free energy profiles, illustrated in Figure 3.11, respectively. Considering the accuracy of the ML potential and the sampling uncertainty of the free energy calculations, we estimated an uncertainty within the range [0.05-0.1] eV.

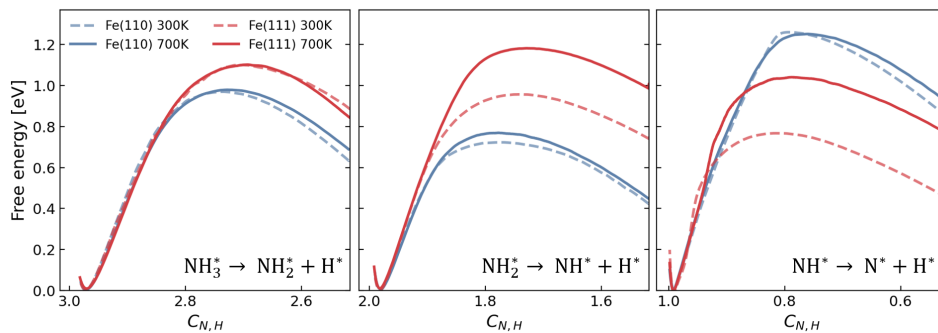


Figure 3.11: **Dehydrogenation free energy barrier.** Free energy profiles of dehydrogenation steps on Fe(110) and Fe(111) surface at T=300 K and T=700 K. The free energy is projected along the coordination number between nitrogen and hydrogen atoms $C_{N,H}$. The free energy of each initial state is set to zero.

calculate the diffusion barriers (Figure 3.4b), which can be directly related to the dehydrogenation barriers. We have reported this relationship in Figure 3.12. This phenomenon can be rationalized if we consider that transition states (see Figure 3.10) have generally higher coordination between N and Fe than the more stable adsorption states. Intermediates must, therefore, reach these states before they can react. Moreover, mobility is a fundamental ingredient for the reaction since the removal of one hydrogen atom is greatly influenced by the kinetics of NH_x on the surface (see Section 3.3.1).

In the case of Fe(111) surface, calculating the diffusion barrier at *operando* temperature is more complex due to the mobility of surface atoms, and we have not attempted to do so. However, in Section 3.2.2, we have described how the shape of the diffusion free energy surfaces, combined with a visual inspection, provides indirect information about the species' stability, with broader minima indicating greater dynamism. In particular, at 700 K, NH, which has been argued to have the highest

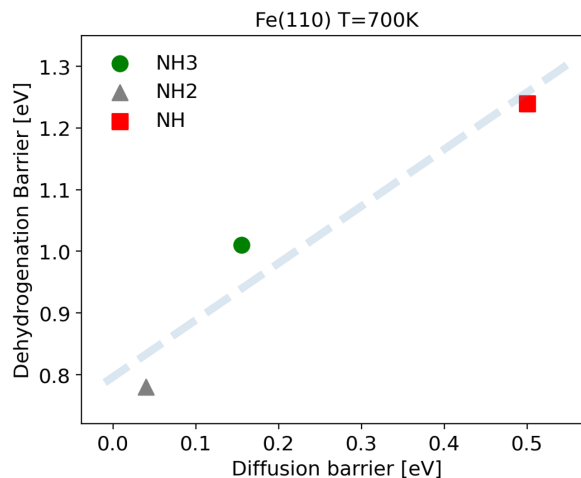


Figure 3.12: **Correlation between dehydrogenation and diffusion barriers on Fe(110)** Free energy barrier of dehydrogenation steps as a function of the diffusion barrier on Fe(110) surface. All the calculations are performed at the *operando* temperature of 700 K. The gray dashed lines follow the linear trend as a guide to the eye.

mobility, has the lowest dehydrogenation free energy barrier. Thus, also in this case, albeit indirectly and qualitatively, a correlation between dehydrogenation barriers and species mobility can be observed.

Even more illuminating is the comparison between the T=300 K and T=700 K behavior. For the more stable Fe(110), no substantial difference is observed between the two temperatures. However, it is interesting to observe a small discrepancy for NH₂, with respect to static calculation (see Figure 3.13). We can rationalize this difference if we consider that static methods probably overestimate the free energy of reactants due to a failure to consider all accessible states. This discrepancy is, in fact, more significant in the case of the highly diffusive reagents.

Significant changes in the two different temperature regimes are instead observed in the case of the Fe(111) surface. For the first dehydrogenation step, the free energy profiles at the two temperatures can be almost exactly superimposed. On the other hand, in the case of the latter two reactions, the catalyst is deactivated at higher temperatures. An explanation for this phenomenon can be suggested by an analysis of the behavior of the reaction moieties at the surface. In fact, the adsorption sites for NH₃ are similar, but for NH₂ and NH, the Fe(111) surface atoms mobility forms new, more stable adsorption sites in which the two nitrogen hydrides can be accommodated such that the N atom can have a tetrahedral geometry. This discourages these two species from visiting states with high coordination where the reaction occurs (see Figure 3.12), leading to an increase in the dehydrogenation barriers. As one might have anticipated, for this more complex surface, the static method fails in predicting

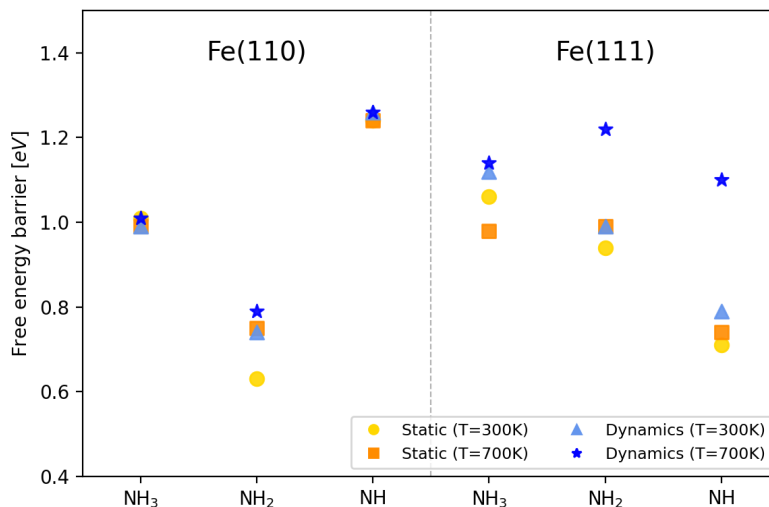


Figure 3.13: **Static versus dynamic methods.** Comparison between free energy barriers computed: as the difference in internal energy using the dimer method (see Section 2.7.1.4), dimer (NN) and corrected from finite temperature effects (see below); from the free energy barrier along the collective variable $C_{N,H}$ and obtained by reweighting the reactive MD simulations (see Section 2.4), FES. The free barriers are reported at 300 K for surface 111 and at 700 K for both surfaces. In static calculation, finite temperature correction was included using the harmonic approximation. The vibrational modes are calculated from a finite difference approximation of the Hessian matrix. We verified that all the minima, as well as the transition states, are true minima and sella-point of the potential energy surface, respectively. All calculations were obtained using the same level of theory (the MLP), and quantum mechanical effects were not included. We can observe that free energy computed with the static method is in good agreement with the dynamic method at low temperatures. In particular, the static method is also in agreement at 700 K with the close-packed Fe(110) surface, where harmonic approximations are expected to be valid. Nevertheless, a little discrepancy is observed for the highly diffusive NH_2 , for which these approximations are not valid. The static method instead misses the temperature effects for the Fe(111) surface by overlooking the mobility of iron atoms. It is important to note that these discrepancies are not equal for all the reactions due to the different effects of the dynamics on the different moieties (see Section 3.2).

free energy, see Figure 3.13.

Consistent with experimental evidence [28], our calculations predict NH as the most stable hydrogenated intermediate on the Fe(110) surface. In the case of Fe(111), we observe an interesting change in the rate-limiting step as a function of temperature. At 300 K, the most stable nitrogen hydride is NH_3 . Instead, at 700 K, the role of the most stable intermediate is played by NH_2 in encouraging agreement with the experiments [29]. Overall, the barrier for the limiting dehydrogenation step is comparable for both surfaces. Therefore, the different activities of the surfaces should be addressed to the recombinative desorption steps, which represent the limiting step for the entire catalytic cycle. This reaction will be described in the next Chapter.

To summarize, in this Chapter, we have investigated the decomposition of ammonia and described its adsorption, diffusion, and dehydrogenation on two clean iron surfaces through MD simulations at the *operando* temperature. All through these reaction steps, the effects of dynamics manifest themselves pervasively and in multiple ways. Even on the (1 1 0) surface, which is the most compact and stable, we saw how a dynamic description helps in rationalizing the link between the interaction of reaction moieties with the metal surface and dehydrogenation events. The effect is indeed dramatic for the open (1 1 1) surface, where in addition to the intermediate's mobility, one has to take into account the iron surface atom diffusivity that, by creating new adsorption sites, significantly alters the free energy profiles. This aspect will be further studied in Chapter 5 where, in addition to thermal effects, we also introduce lateral interactions.

Analyzing the transition states makes it clear that there is not a single, well-defined one, but rather one has to consider an ensemble of transition states. For each dehydrogenation step, we are able to identify at least two classes of reactive paths. This is the case not only for the more “dynamic” surface (1 1 1) but also for the more “static” (1 1 0).

It is important to note how these dynamical effects result in a non-trivial temperature dependence for the (1 1 1) surface, confirming the gap between the standard approach based on idealized assumption and what happens in the messier *operando* environment, even in the low adsorbate coverage limit. Molecular dynamics and enhanced sampling techniques, on the other hand, fully account for entropic effects, allowing these processes to be studied without any previous assumption. However, it is interesting to note that the same tools that have allowed us to simulate the complex behavior of the system also permit the unifying of all different dehydrogenation steps on different surfaces using the charge transfer as a descriptor.

Of course, we recognize that this is a simplified model compared to the actual industrial catalyst, which should require the entire catalytic cycle and other effects, such as lateral interactions with adsorbed intermediates and promoter activities, to be considered. However, it already demonstrates the need for a dynamical description of heterogeneous catalysis. A step toward this complexity will be taken in the following Chapters.

Chapter 4

The fate of nitrogen on Fe(110)

Once the ammonia dehydrogenation process is complete, relaxed nitrogen (N^*) and hydrogen (H^*) remain on the surface. These species must undergo recombination before being released into the gas phase. In this Chapter, we focus specifically on the fate of nitrogen, particularly examining its behavior on the Fe(110) surface. This surface is the most prevalent in experimental samples and thus holds the most significance for industrial applications. The behavior of the Fe(111) surface under high N^* coverage, which is of great interest for the Haber-Bosch process, will be discussed in the next Chapter.

To understand the possible processes nitrogen can undergo on the surface, we analyze the behavior of isolated N^* and H^* atoms on Fe(110). As expected from previous studies [38], hydrogen weakly interacts with the surface, which is not the case for N^* . For this reason, we focus on the recombination processes $2N^* \rightarrow N_2$ and explore the competing process of nitrogen migration into the bulk.

The latter process is the first elementary step toward bulk nitride formation. Therefore, comprehending the interplay between these two competing processes is essential for accurately reflecting the catalyst's condition during operation. In fact, as discussed in Section 1.2.1, there are experimental observations of nitride formation during ammonia decomposition.

Finally, to move towards a realistic surface model, we will study the recombination and migration processes in the presence of a finite surface coverage of nitrogen and hydrogen. This allows us to explore how surface dynamics and lateral interactions impact nitrogen behavior under more industrial-like catalytic conditions, allowing us to give answers at the microscopic level to the question “*which is the fate of nitrogen?*”




Figure 4.1: Graphical visualization of the Fe(110) $8 \times 10 \times 10$ slab used as a catalyst, in perspective and top view. The tripod of Cartesian directions (x, y, z) corresponds to $([001], [1\bar{1}0], [110])$. Additionally, three vectors representing the variables v_{hkl} , with $[h, k, l] = [1\bar{1}1], [1\bar{1}\bar{1}], [0\bar{1}0]$, defined in Equations 2.30, 2.31, are displayed. Note that their magnitude and the point of application do not agree with the slab depicted in the figure, while their directions are representative.

4.1 Interaction of nitrogen and hydrogen with Fe(110)

Before describing in detail the various interactions, to facilitate the reader, we present in Figure 4.1 an example of the studied system, where we graphically display the most relevant directions. In addition to the three Cartesian coordinates (x, y, z) , corresponding to the directions $([001], [1\bar{1}0], [110])$, we indicate the three variables v_{hkl} , with $[h, k, l] = [1\bar{1}1], [1\bar{1}\bar{1}], [0\bar{1}0]$, previously defined in Equations 2.30 and 2.31. These three variables, which capture the crystal geometry and thus its symmetries, will be extensively used during the analysis. Additionally, as described in Section 2.7.3.2, they are employed as collective variables in enhanced sampling simulations to sample all relevant processes. As we will see, all nitrogen-related processes that we will study here exhibit barriers ranging from ~ 10 to $\sim 40 k_B T$ ($T=700$ K). These processes occur over much longer time scales ($\sim 10^4$ s) than those accessible by ML potentials ($\sim 10^{-7}$ s). To simulate them, it is therefore necessary to apply advanced simulation techniques, as described in Section 2.4.

All the results reported in this Chapter are derived from MD simulations at the operando temperature of 700 K. Even though we do not focus directly on the non-trivial aspects introduced by finite-temperature simulations, we want to emphasize that the simulation environment is highly dynamic and deviates significantly from the static one, as mentioned in the previous Chapter. This note is made because

the figures presented here, static for their nature, may not fully convey the system's dynamics.

We begin by presenting the interaction of a single nitrogen and hydrogen atom adsorbed on the Fe(110) surface. The most energetically stable adsorption configurations of nitrogen hydrides have already been investigated through static DFT calculations; hollow sites have been computed as the minimum energy adsorption sites for N^* [38, 41, 43], while the hpc sites are the most stable for H^* [38, 197]. Regarding the adsorption site nomenclature, we refer to the one presented in Figure 3.2. At $T=0$ K, our results are in agreement with previous theoretical work. In the following, we focus on the effects of temperature on this relatively rigid surface to obtain a statistical description of the adsorption and related processes.

In Figures 4.2 and 4.3, we present the most probable adsorption states obtained from the minima of free energy surfaces (FES) projected along the two in-plane crystallographic directions $[001]$ and $[1\bar{1}0]$. From the FES, we can also analyze the diffusion mechanism and the related free energy barrier.

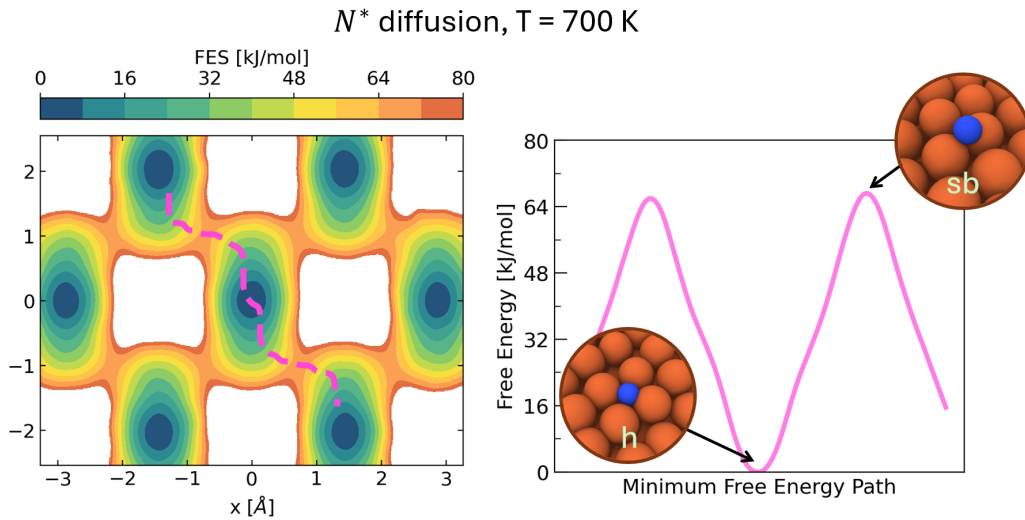


Figure 4.2: On the left panel, free energy of N^* projected along the two crystallographic directions $[001]$ (x -axis) and $[1\bar{1}0]$ (y -axis). Local minima represent metastable states, and fuchsia dashed lines denote the minimum free energy diffusion pathways. On the right panel, free energy is projected along the minimum free energy pathways (MEP). Some snapshots of representative geometries along the minimum free energy path are also shown. The high-symmetry adsorption sites are labeled according to Figure 3.2.

Atomic nitrogen is preferentially adsorbed at the hollow (h) sites. In this configuration, N^* is 4-fold coordinated with four adjacent Fe atoms on the surface or 5-fold coordinated with four adjacent Fe atoms on the surface and the Fe atom di-

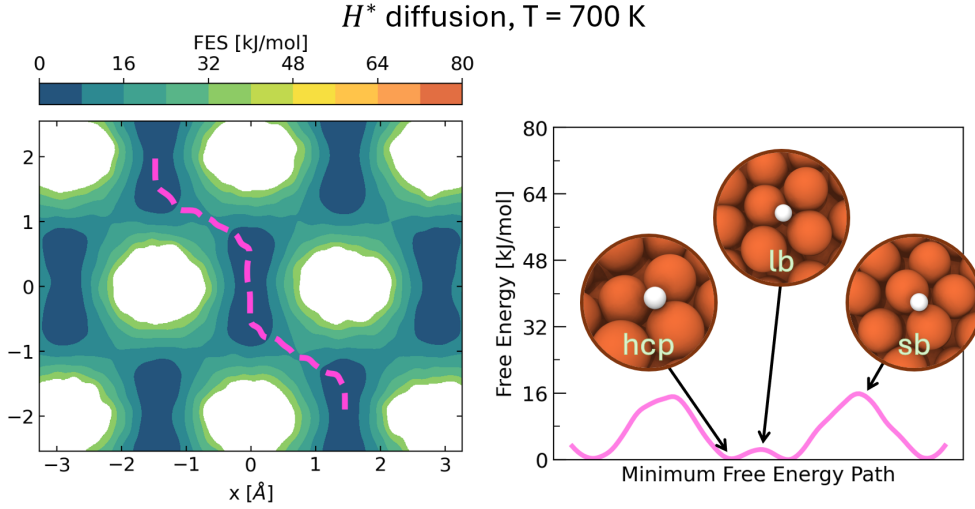


Figure 4.3: Analogous to Figure 4.2, for H^* interaction with the substrate.

rectly below in the second layer when deeply adsorbed, as we will discuss later. N^* is firmly bonded to the iron surface, compensating for its valence shell's lack of three electrons. To diffuse between two adjacent hollow (h) sites, N^* passes through a short-bridge (sb) site with a barrier of 67 kJ/mol (see Figure 4.2).

Atomic hydrogen, on the other hand, is highly mobile. In Figure 4.3, we show the free energy surface, where it can be observed that hydrogen quickly diffuses on the surface following a minimum energy pathway similar to that of nitrogen. Unlike N^* , the minimum energy adsorption site for hydrogen is the hcp site. Diffusion to adjacent hcp sites occurs either barrierlessly ($\Delta G < 1 k_B T$) via long-bridge (lb) sites or through short-bridge (sb) sites with a small barrier of 16 kJ/mol. In general, H^* can explore many surface sites. This can be explained by hydrogen being partially reduced on the metallic surface, which can effortlessly supply only one electron necessary for bonding in various configurations.

In Figure 4.4, we present the free energy of adsorption for N^* and H^* projected along the two variables $v_{1\bar{1}1}$ and z . Unlike the x and y variables, the variable $v_{1\bar{1}1}$ allows us to identify the most stable adsorption sites along the crystallographic direction $[1\bar{1}1]$ (which contains the minima and maxima of the MEP), as it is orthogonal in crystallographic space to the direction $[1\bar{1}\bar{1}]$ (see Equations 2.30). In this lateral projection, we can observe the vertical oscillations along the z variable, normal to the (110) surface. The zero along the z variable is set at the average z coordinate of the top-layer iron atoms. As can be seen from the free energy profile, N^* can make deep vertical excursions much more significant than H^* . In particular, considering the region most accessible to the adsorbed species ($\Delta G < 3 k_B T$), z can vary in the

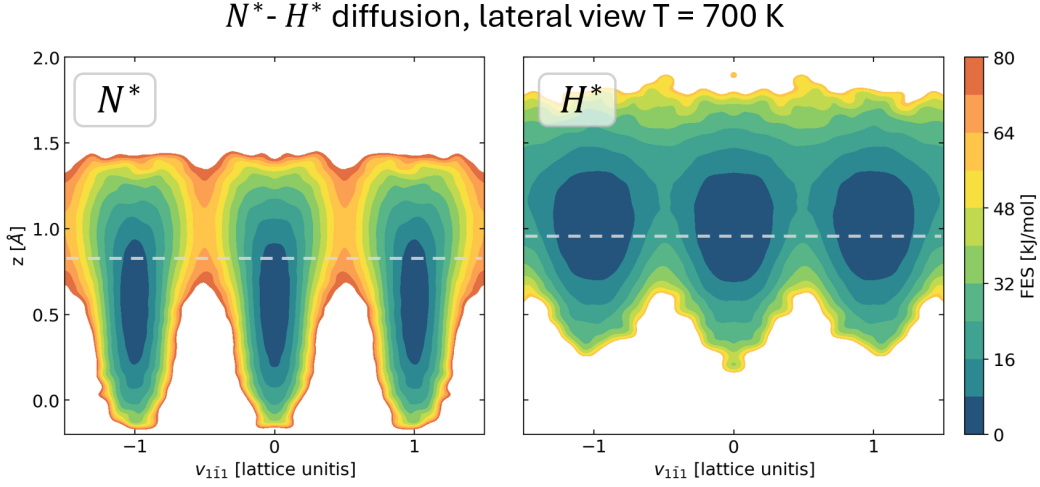


Figure 4.4: Free energy of adsorption for N^* and H^* projected along two variables, $v_{1\bar{1}1}$ and z , i.e., along the crystallographic direction $[1\bar{1}1]$ (orthogonal in crystallographic space to the direction $[1\bar{1}\bar{1}]$, see Equations 2.30) and $[110]$, at a temperature of 700 K. The zero along the z variable is set at the average z coordinate of the top-layer iron atoms. White dashed lines represent the z values for the lowest energy adsorption sites at $T = 0$ K [38].

range 0.1-1 Å for N^* , while it varies between 0.6-1.4 Å for H^* . At a temperature of 700 K, the accessible states for N^*/H^* significantly increase the gap from the mere 0.15 Å difference between the static adsorption sites, represented by the white dashed lines in Figure 4.4. These results show that H^* preferentially stays on the surface rather than within it, unlike what can be stated for N^* .

Before discussing the reactive processes, we want to show how the presence of another N^* alters the diffusion of N^*/H^* . In Figure 4.5, we report the FES projected onto the exact coordinates as in Figures 4.2 and 4.3. To obtain this free energy profile, we placed a nitrogen atom at the center of the cell ($x, y = (0, 0)$). We constrained it to stay in that adsorption site by applying a reflective wall at $v_{1\bar{1}1}$ and $v_{1\bar{1}\bar{1}}$ at ± 0.55 , corresponding to the maxima of the diffusion MEP. Additionally, to achieve more effective sampling, we restricted the region accessible to the N^*/H^* atom using another wall at $v_{1\bar{1}1}$ and $v_{1\bar{1}\bar{1}}$ at ± 2.55 , displayed in Figure 4.5. As can be seen, the presence of the adsorbed nitrogen atom in the center decreases the probability of finding another N^* in one of the adjacent adsorption sites (the four hollow sites in the $[1\bar{1}1]$ direction and equivalent), which present an energy approximately ~ 40 kJ/mol higher than the second-neighbor sites along the same direction. These latter sites represent the most stable adsorption sites for the second N^* , in agreement with low-energy electron diffraction (LEED) patterns [25] and previous theoretical calculation [202]. On the other hand, H^* is less affected by the presence of the nitrogen atom, which

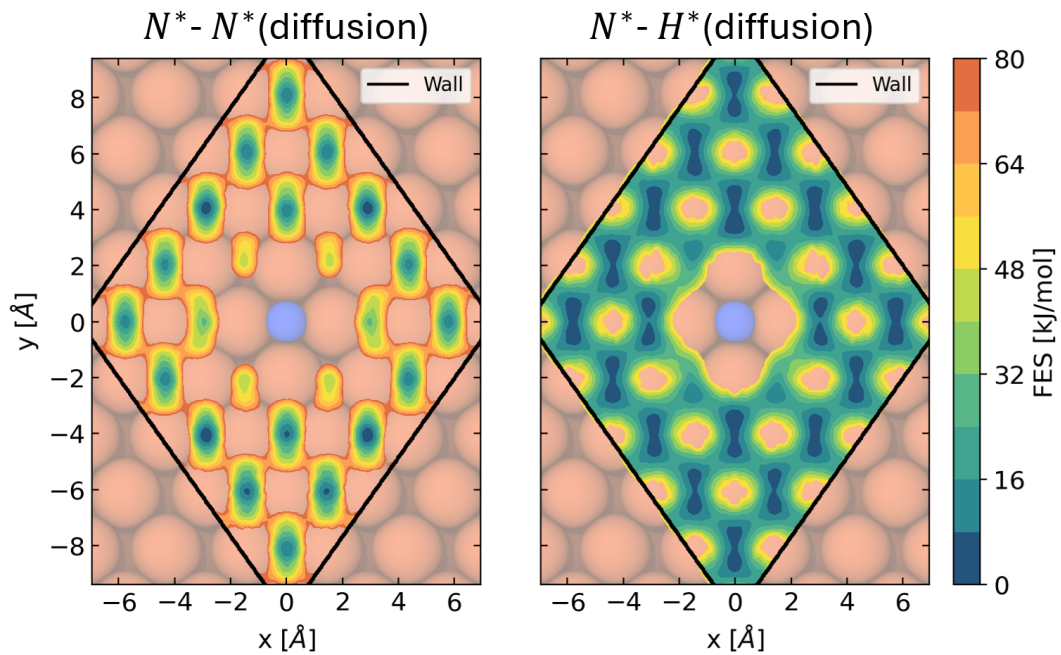


Figure 4.5: Free energy of N^* (on the left) and H^* (on the right) projected along the two crystallographic directions $[001]$ (x -axis) and $[1\bar{1}0]$ (y -axis), in the presence of one atomic nitrogen adsorbed in the center. For an effective sampling, a reflective wall was placed at $|v_{hkl}| < 2.55$, $hkl = 1\bar{1}1, 1\bar{1}\bar{1}$, displayed in the figure with a black line. Similarly, to avoid the diffusion of the central nitrogen, another wall at $|v_{hkl}| < 0.55$, $hkl = 1\bar{1}1, 1\bar{1}\bar{1}$.

only slightly modifies the free energy profile of the adjacent sites.

In light of these results, we focus on studying processes related to nitrogen, specifically its recombination and migration into the bulk. We leave other hydrogen-related processes for future studies as they have less impact on the catalyst. This choice is also motivated by experimental evidence. From transient NH_3 decomposition experiments, we observe that the release of H_2 into the gas phase occurs immediately following the dissociation of ammonia, unlike what happens for N_2 . This suggests that hydrogen released on the surface recombines and desorbs easily without dissolving into the bulk.

4.2 Nitrogen recombination

Recombination of nitrogen atoms on the Fe(110) surface leads to the formation of nitrogen molecules (N_2). This process is fundamental and necessary to allow the continuous decomposition of ammonia, as nitrogen adsorbed on the surface inhibits NH_3 adsorption.

To study the recombination process, it is necessary to clearly define the reactant and product states to identify the free energy barrier of the process correctly. In the case of $2N^*$ adsorption, as shown in Figure 4.5, the reactant state can be defined as one where the two nitrogen atoms occupy two adjacent adsorption sites (along the $[1\bar{1}1]$ direction and equivalents). This choice is generally the one adopted in static calculations [30, 38, 41, 203, 204]. Thus, we have simulated the recombination process, constraining two N^* to occupy two adjacent adsorption sites. However, to avoid forcing the system to explore only two sites, we constructed a reflective wall tailored to the crystal's symmetry and centered at the center of mass of the two nitrogen atoms. Specifically, we defined:

$$v_{hkl}^{abs} = v_{hkl}(\|d_x(N, N)\|, \|d_y(N, N)\|) , \quad hkl = 1\bar{1}1, 1\bar{1}\bar{1} \quad (4.1)$$

where v_{hkl} is defined by Equations 2.30, and $d_i(N, N)$ represents the x and y Cartesian components of the distance between the two N atoms. We placed a reflective wall when $v_{1\bar{1}1}^{abs} > 1.55$ and $\|v_{1\bar{1}\bar{1}}^{abs}\| > 0.55$. Referring to Figure 4.5, we can imagine sitting on a spatial reference system on one of the N atoms adsorbed in the hollow site to understand the wall's action. From this reference system, the second nitrogen can visit all four adjacent hollow sites. However, it cannot overcome the 12 barriers (three for each site) at the short-bridge sites to reach more distant adsorption sites.

In Figure 4.6, we show the free energy profile along the distance between the two nitrogen atoms $d(N, N)$, obtained from MD simulations at a temperature of 700 K.

Concurrent with the formation of the N-N bond, there is a release of charge to the

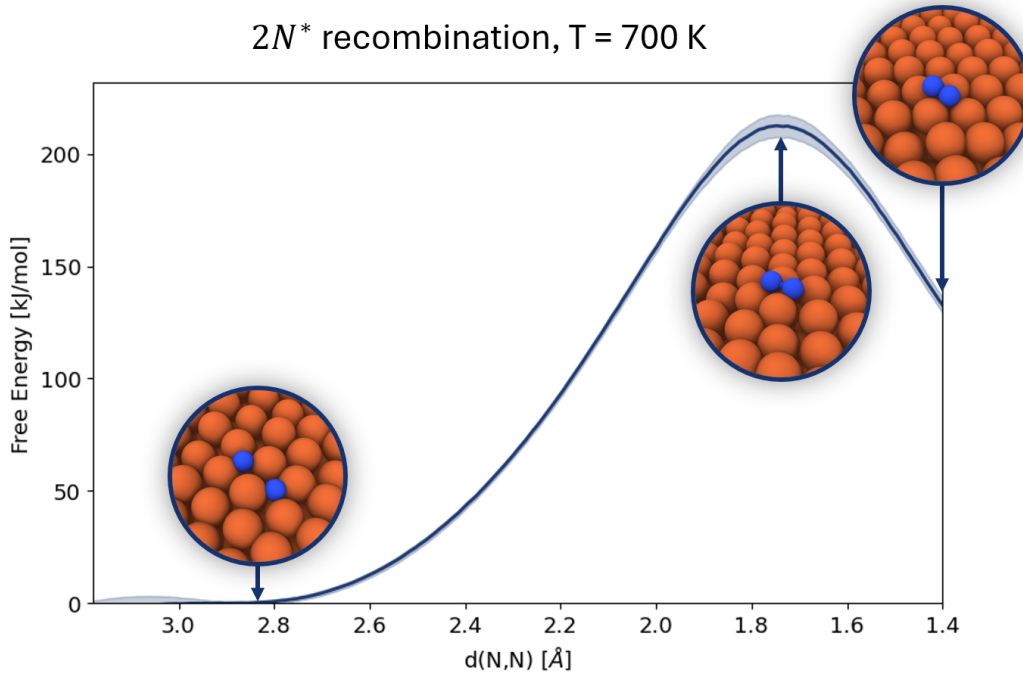


Figure 4.6: Free energy profiles of recombination steps on the Fe(110) surface at $T=700\text{ K}$. The free energy is projected along the distance between the two nitrogen atoms involved $d(N, N)$. To effectively sample the reaction from the precursor state (see text), a reflective wall was placed when $v_{111}^{abs} > 1.55$ and $\|v_{111}^{abs}\| > 0.55$. Similarly, to avoid molecular nitrogen from desorbing, we limit nitrogen to a highly coordinated adsorption state by placing another wall for $d(N, N) < 1.35$ (N_2 bonding length is directly correlated with its coordination with iron substrate [44, 202, 204]). Three snapshots captured from the simulation along a reactive event are also reported. The uncertainty along the free energy profile is computed as the standard deviation of four independent simulations performed with four different ML-potential trained. With this protocol, both the error originating from ML-potential and the one from sampling/reweighting are taken into account.

surface of about $1|e|$, which reduces the interaction strength between the adsorbate and the substrate, thus easily allowing for the desorption of N_2 . The obtained free energy barrier is approximately 216 kJ/mol . In the previous Chapter, we showed that the dehydrogenation steps preceding recombination have free energy barriers in the $75\text{--}120\text{ kJ/mol}$ range. Thus, recombination represents the rate-determining step in the decomposition of NH_3 on Fe, in agreement with recent literature [30, 43]. Moreover, considering nitrogen atoms must occupy adjacent hollow sites before recombination, the free energy barrier becomes even higher. We have seen that the free energy of adjacent hollow sites is about 40 kJ/mol higher than that of the more

stable adsorption sites. Experimental studies, such as TPD, confirm that nitrogen desorption occurs at elevated temperatures. However, the desorption peaks align with the decomposition of bulk nitrides, as we will discuss later.

4.3 Nitrogen migration inside the bulk

In addition to surface recombination, nitrogen can also migrate from the surface into the bulk of Fe(110), contributing to the formation of iron nitrides (Fe_4N , Fe_2N). The migration begins with nitrogen atoms penetrating surface layers and occupying octahedral interstitial sites within the iron lattice.

In Figure 4.7, we report the free energy of one nitrogen atom migrating inside the bulk at $T=700\text{ K}$, projected along the two crystallographic directions $[1\bar{1}0]$ (y coordinate on the x-axis) and $[110]$ (z coordinate on the y-axis). From this projection of the free energy, we can clearly identify the migration direction into the bulk, coinciding with the crystallographic direction $[0\bar{1}0]$, and effectively described by $v_{0\bar{1}0}$.

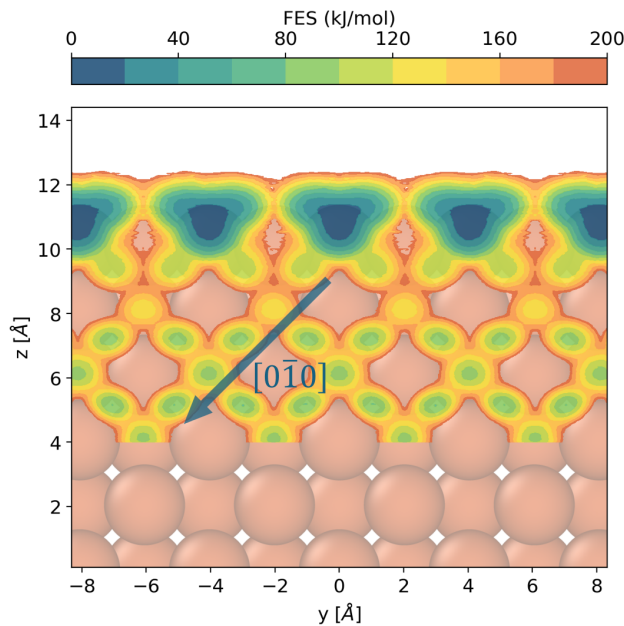


Figure 4.7: Free energy of N^* dissolution/segregation on Fe(110) slab at 700 K projected along the two crystallographic directions $[1\bar{1}0]$ (y) and $[110]$ (x). Local minima represent metastable states. The arrow denotes the $[0\bar{1}0]$ direction of migration inside.

To better analyze the migration process inside the bulk, we projected the free energy along one of these “diffusion channels”, as shown in Figure 4.8.

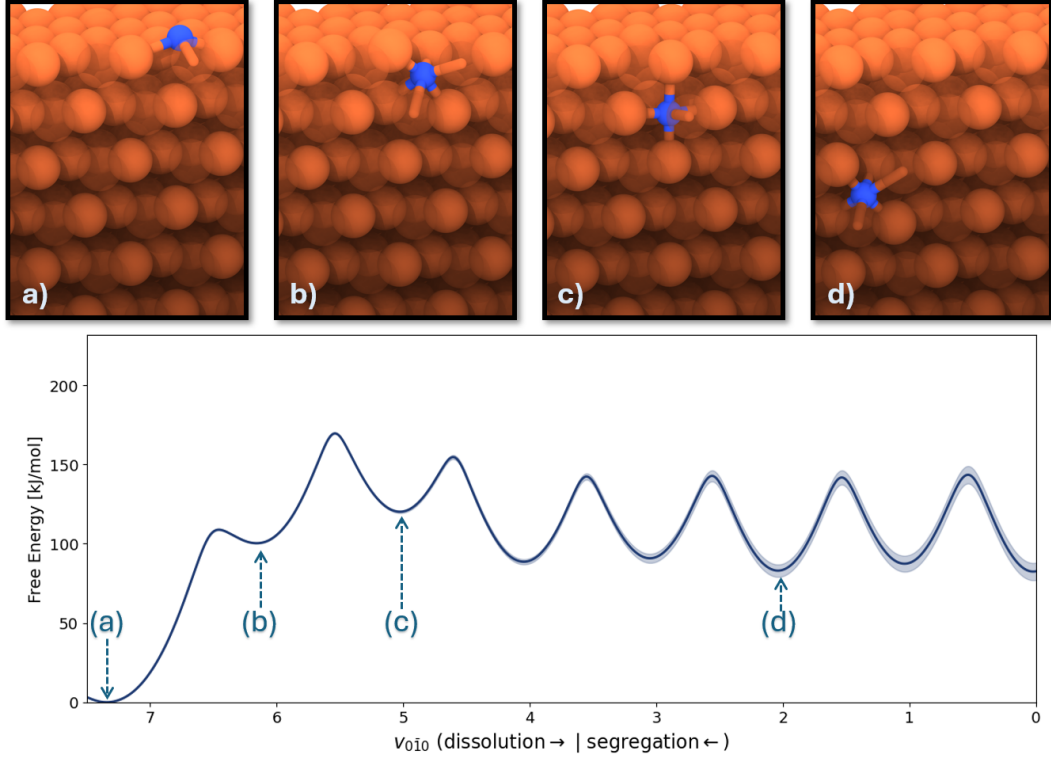


Figure 4.8: Free energy profiles of nitrogen dissolution/segregation on the Fe(110) slab at $T=700$ K. The free energy is projected along the $v_{0\bar{1}0}$ direction. To constrain the simulation to a single “diffusion channels”, a reflective wall is placed when $\|\tilde{v}_{1\bar{1}\bar{1}}\| > 0.55$ and $\|\tilde{v}_{\bar{1}\bar{1}\bar{1}}\| > 0.55$. Four snapshots captured from the metastable state are shown: a) hollow site, b-d) octahedral interstitial. The uncertainty along the free energy profile is computed as the standard deviation of four independent simulations performed with four different MLPs. This protocol considers both the errors originating from the MLP and the ones deriving from sampling/reweighting.

From the surface hollow adsorption site (Fig. 4.8a), the migration proceeds to the first metastable site (Fig. 4.8b). In this octahedral interstitial site, nitrogen is located in the middle of a BCC unit cell, forming two shorter bonds with the Fe atoms located at the vertices of that side (one from the first layer and one from the second layer) and four longer bonds with the Fe atoms in the center of the BCC (two from the first layer and two from the second layer). The migration then proceeds to a second octahedral site (Fig. 4.8c). Here, nitrogen is positioned at the center of the BCC face, forming two shorter bonds with the Fe atoms in the center of the BCC (both in the second layer) and four longer bonds with the Fe atoms at the face vertices (one from the first layer, two from the second layer, and one from the third layer). From here, it alternates to deeper octahedral interstitial sites (Fig. 4.8d). We emphasize that, geometrically, all these interstitial sites are equivalent. Additionally, the choice

to limit our study to the $[010]$ direction was made to reduce the complexity of the problem. Indeed, nitrogen can equivalently move from the octahedral interstitial site to any of the six adjacent octahedral sites along the $\langle 100 \rangle$ directions. We now focus on the energy barriers associated with the migration process without delving into the electronic details that were widely covered by static studies [205–207].

To migrate to the first interstitial site (Fig. 4.8b) from the surface hollow adsorption site (Fig. 4.8a), nitrogen must overcome a barrier of about 110 kJ/mol. Subsequently, to reach the first stable interstitial site (Fig. 4.8c), it must overcome an additional barrier of 70 kJ/mol (a total of 169 kJ/mol from the surface). From here, it can reach deeper interstitial sites by overcoming a barrier of 35 kJ/mol and diffusing into the bulk with barriers of approximately 60 kJ/mol. Once it has passed through the second Fe layer, nitrogen must overcome a barrier of about 85 kJ/mol to return to the surface. Considering that there are three net interstitial sites per bulk Fe atom (6 per unit cell) and the ratio between bulk and surface atoms, the probability that a nitrogen atom dissolved in the bulk will segregate back to the surface is very low.

MD simulations reveal that nitrogen migration into the bulk is clearly favored on pristine Fe(110). Before discussing the implications of these results, we take an additional step toward modeling a realistic system by introducing a finite coverage of nitrogen and hydrogen.

4.4 Lateral interaction effects

Lateral interactions between nitrogen atoms adsorbed on Fe(110) play a significant role in determining the adsorption, recombination, and migration processes. At low coverage, nitrogen atoms behave independently, but as coverage increases, repulsive interactions between adjacent nitrogen atoms become more prominent. These interactions influence the adsorption site preferences and can alter the free energy barriers for both recombination and migration. To confirm our previous findings in a model in a more realistic scenario that also considers lateral interactions, we performed simulations with finite coverage of N and H atoms (20% and 40% of a monolayer with a ratio N : H = 1 : 3). It should be noted that, unlike the approach commonly used to estimate the effects of coverage through static calculations on small systems, here we explicitly simulated the dynamics on large systems at the *operando* temperature of 700 K.

MD simulations at various coverage levels indicate that the barrier for recombination decreases with increasing coverage (Figure 4.9 left panel). Similarly, the migration barrier decreases as nitrogen coverage increases (Figure 4.9 right panel).

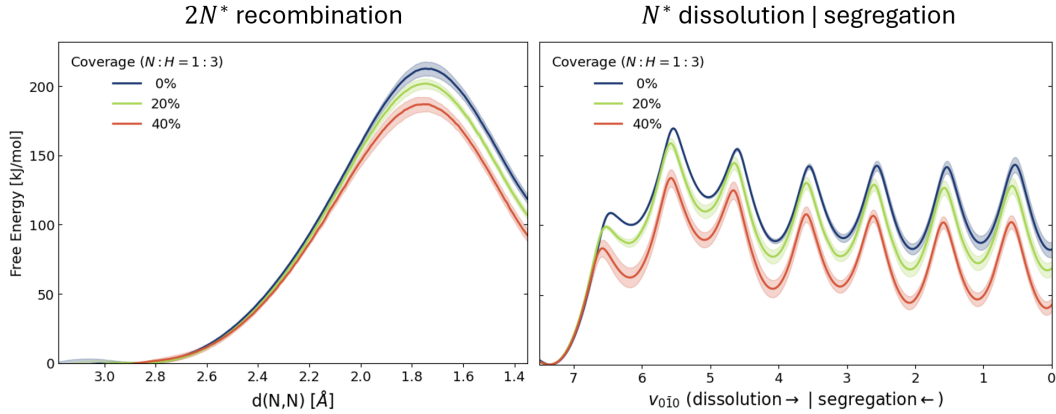


Figure 4.9: Free energy profiles of nitrogen recombination (left) and dissolution/segregation (right) on the Fe(110) slab at $T=700$ K at various coverages (0, 20, and 40% of a monolayer with an N : H = 1 : 3). The free energies are computed identically to the 0% coverage case, see Figures 4.6 and 4.8.

In particular, it is interesting to observe that there is a reduction in the free energy differences between reactants and products ($2N/N_2$ and N_{hollow}/N_{bulk} , respectively), particularly remarkable for the N in the deep octahedral interstitial which decreases from 85 kJ/mol to 45 kJ/mol. By focusing on the migration inside, we see that at low coverage (20%), there is a first reduction in the free energy barrier for entering inside the first interlayer interstitial octahedral site (Fig. 4.8a). However, the remaining free energy profile is mostly unchanged with respect to the pristine Fe case, suggesting that lateral interaction accounts for the destabilization of N in the hollow site (Fig. 4.8a). Instead, at higher coverage (40%), there is also the additional effect of stabilizing the first interlayer interstitial octahedral site.

In Figure 4.10, we report the overall free energy barrier for the two processes. The nitrogen recombination barrier decreases with increasing coverage, from 216 kJ/mol to around 196 kJ/mol. Similarly, the migration barrier decreases from 169 kJ/mol to 134 kJ/mol as nitrogen coverage increases. At this coverage ratio (N:H=1:3), the resulting slope of linear regression (dashed lines in Figure 4.10) are about -5 and -9 kJ mol⁻¹/10% coverage for the two processes, respectively, showing that lateral interaction has a more substantial impact on the latter process, with its barrier decreasing strongly. This makes the competition between the two processes even more unbalanced in favor of migration into the bulk rather than desorption through recombination.

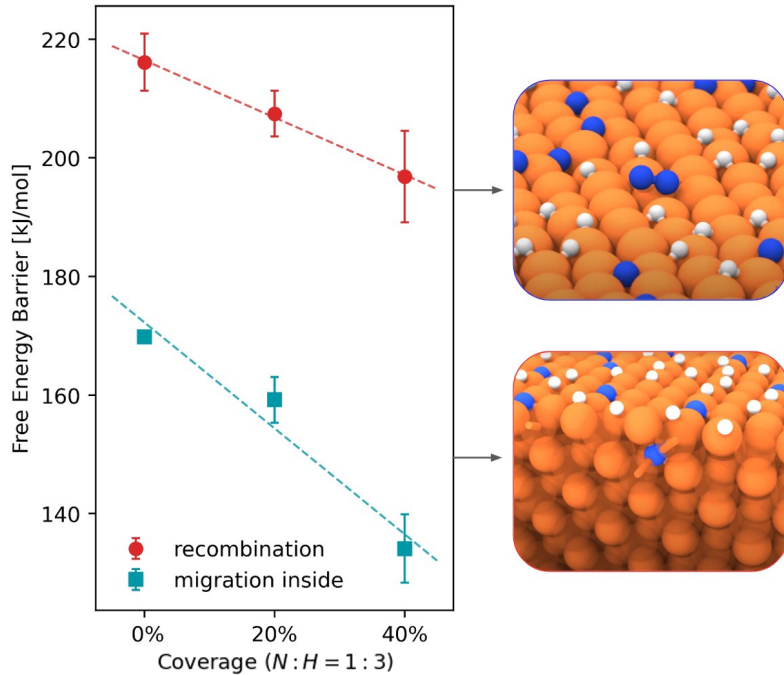


Figure 4.10: Free energy barrier of recombination (red circles) and migration inside (green squares) as functions of surface coverage, obtained from Figure 4.9. The surface coverage refers to N, H monolayer coverage with an N:H ratio equal to 1:3. The slope of linear regression (dashed lines) shows that lateral interaction has a more substantial impact on the latter process, with its barrier decreasing strongly. In the insets, two snapshots from simulations at 40% ML coverage are reported.

4.5 Discussion and experimental comparison

The full diagram of ammonia decomposition is shown in Figure 4.11, where we present the comprehensive behavior of nitrogen on Fe(110) at the *operando* temperature of 700 K. This provides the first complete microscopic map of ammonia decomposition obtained from MD simulation, highlighting the competition between nitrogen recombination and migration into the bulk.

Our MD simulations reveal that nitrogen prefers to migrate into the bulk of Fe(110) rather than recombining on the surface. This finding is consistent with experimental observations of nitride formation during ammonia decomposition, such as those presented in Ref. [208]. Of particular interest in comparing our results with the experimental data are the transient NH_3 decomposition measurements. These experiments show the immediate release of H_2 after ammonia decomposition, while nitrogen remains on the surface and migrates into the bulk, forming nitrides. The nitrogen is then released only after the formation and subsequent decomposition of

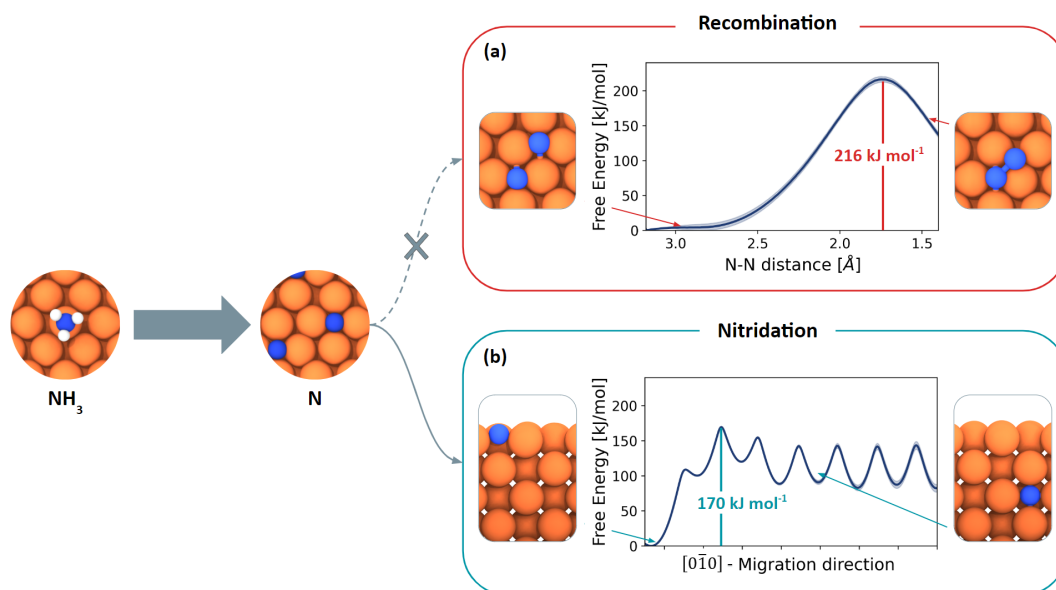


Figure 4.11: Diagram of NH_3 decomposition on $\text{Fe}(110)$, with a focus on the fate of nitrogen on the surface. Insets (a,b) show the free energy profiles for nitrogen recombination and migration into the bulk, respectively (see Figures 4.6 and 4.8 for details).

nitrides (Fe_4N and Fe_2N) at higher temperatures. Our computation results align well with these findings, both for H that lives on the surface and is easily released for the migration of nitrogen favored over recombination.

The combination of several experimental techniques, including transient kinetic measurements and *in situ* X-ray diffraction (XRD), allowed for the detailed observation of nitride formation and decomposition. The results of the transient NH_3 decomposition show an apparent activation energies of N_2 desorption due to nitride decomposition amounting to 157 kJ/mol (Fe_2N) and 172 kJ/mol (Fe_4N). In very good agreement the same apparent activation energy of 173 kJ/mol for NH_3 decomposition was derived, indicating that the decomposition of Fe_4N into Fe and N_2 is the rate-determining step of the reaction. Furthermore, temperature-programmed desorption (TPD) data on the spent catalyst shows distinct peaks for N_2 release at 737 K and 816 K, corresponding to the stepwise decomposition of Fe_2N and Fe_4N . *In situ* XRD measurements confirmed the stepwise iron nitride formation $\alpha\text{-Fe} \rightleftharpoons \gamma'\text{-Fe}_4\text{N} \rightleftharpoons \epsilon\text{-Fe}_3\text{N}_{1.5}$ and shed light on the different phase transformations that the catalyst underwent at different temperatures.

It is important to note that a direct comparison between the apparent activation energies for nitride decomposition and the computed free energy barriers for nitrogen migration and recombination is challenging due to the different nature of these

processes. The computed free energy barriers for nitrogen migration into the bulk (169 kJ/mol) and recombination (216 kJ/mol) on Fe(110) are not directly comparable to the experimental activation energies of nitride decomposition due to the complexity of the catalytic system. However, they provide a detailed microscopic picture of the underlying atomistic processes. Previous studies on the decomposition of Fe₄N to α -Fe and N₂ conducted by Ertl *et al.* [209] report activation energies in the range from 213 to 238 kJ/mol for unpromoted Fe wires and foils. They concluded that the decomposition of the surface nitride is the rate-determining step, which prevents the spontaneous decomposition of the metastable bulk iron nitrides under ambient conditions. Also, they derive the same conclusion that surface nitride decomposition is the rate-determining step of the recombinative desorption of N from differently oriented Fe single crystal surfaces because of the identical activation energies.

Additional flow-switching experiments were performed to obtain further insight into the active state of the Fe catalyst at different temperatures. At high temperature ($T = 873$ K), the active state of the catalyst is most likely α -Fe and not Fe₄N because hardly any release of N₂ was observed when switching the gas flow from diluted NH₃ to He. Other studies support this view. Arabczyk and co-workers [210, 211] found that nitride formation negatively affects the NH₃ decomposition rate. Additionally, Chen *et al.* [51] recently showed that alloying Fe with Co can improve catalytic efficiency and, more importantly, prevent nitride formation.

These findings have important implications for improving catalytic efficiency. Suppressing nitride formation while enhancing N₂ desorption rates can potentially lead to better-performing NH₃ decomposition catalysts, particularly under higher pressure. In Chapter 7, we will explore how alloying iron with cobalt could help in preventing the inward migration of nitrogen and reducing nitride formation, ultimately leading to enhanced catalytic activity.

Chapter 5

Fe(111) under high N^* coverage

As explored in the previous chapters, for a catalyst to be efficient and durable, it is crucial that the reaction products do not poison the catalyst. In this chapter, we study Fe(111) surface behaviors in the presence of atomic nitrogen (N^*) adsorbed. The presence of N^* coverage is expected both in the ammonia decomposition or synthesis context, respectively released on the surface after ammonia or molecular nitrogen cracking. Here, we focus on the Haber-Bosch synthesis process, where the rate-limiting step is believed to be the decomposition of nitrogen molecules. On Fe(111), N^* can occupy the active sites, and thus, it is important to ascertain how a high N^* coverage affects nitrogen dissociative chemisorption.

To answer this question, we study the properties of the Fe(111) surface at different N^* coverage both at room and *operando* temperature. In the latter regime, we already presented that Fe surface atoms exhibit high mobility, promoting the formation of adatoms and vacancies, and causing the catalytic centers to acquire a finite lifetime. We discover that the N^* coverage reduces but does not eliminate iron mobility. Remarkably, the N^* atoms stabilize triangular surface structures associated with the formation of vacancies, which are a sign of a frustrated drive toward a more stable Fe_4N phase, somehow expected in light of what just seen in Chapter 4. As a consequence, nitrogen atoms tend to cluster, reducing their poisoning effect. At the same time, the reduction in the number of catalytic centers is counteracted by an increase in their lifetime. The combined effect is that the dissociation barrier is not significantly altered in the range of coverages studied. Another time, these results bring to light the complex role that dynamics plays in catalytic reactivity under *operando* conditions.

5.1 Active site for N_2 dissociation

In the Haber-Bosch process, the (111) surface is indeed the most active one among the low miller index iron surfaces, and the cleavage of the di-nitrogen triple bond has been argued to be its rate-limiting step [24, 25, 31, 32, 202]. At low temperatures, the dissociative chemisorption process has been deeply experimentally and theoretically investigated

Nørskov *et al.* [202] conducted static Density Functional Theory (DFT) calculations, which confirmed Ertl's proposed mechanism while adding new atomistic insights. They identified two distinct vertical adsorption sites for N_2 (Fig. 5.1), termed γ and δ , based on whether the molecule was adsorbed on top of a first- or second-layer atom. Additionally, they predicted two horizontal adsorption sites (see Fig. 5.1). In one configuration, N_2 occupies a bridge position between first-layer atoms (α), while in the other, it resides in a hollow position atop a third-layer atom (α'). The α' site is particularly significant, as it is considered a precursor to dissociation. They suggest that this precursor state can be accessed either directly from the gas phase or via a transition through $\gamma \rightarrow \delta \rightarrow \alpha \rightarrow \alpha'$. Throughout this sequence, the nitrogen triple bond is gradually weakened. In the α site, a π bond is initially transferred to two surface atoms, followed by another in the α' site. At the α' position, the molecule interacts with seven-fold coordinated iron atoms (C_7), which are more effective in donating electrons to the N_2 molecule [42]. The increased catalytic activity of the Fe(111) surface has been attributed to its step structure, which stabilizes the α configuration and provides a high density of accessible C_7 atoms [31, 32]. The set of Fe atoms surrounding the N_2 molecule in the α' state is referred to as the χ_7 environment. This arrangement is considered a catalytic site because, as the N_2 molecule




Figure 5.1: Different adsorption sites for N_2 at 0 K. Adapted with permission from [44].

approaches the α' position, the charge transfer from the iron surface to the nitrogen increases. Once within this site, the N_2 molecule can rotate through different orientations until the electronic orbitals are optimally aligned for the reaction.

More recently, in Ref. [44], Bonati *et al.* have extended the theoretical investigations with MD simulations at the room ($T = 300$ K) and *operando* ($T = 700$ K) temperature regimes. Low-temperature simulations confirmed the static predictions, while high-temperature simulations showed that the C_7 sites, which assist the breaking of the N_2 triple bond, acquire a finite lifetime. This results in an increase in the free energy barrier for N_2 dissociation, which must necessarily pass through highly coordinated sites. Here, we go on to extend the work of Ref. [44], going on to study the influence of finite nitrogen coverage on surface dynamics and its effects on N_2 dissociation.

5.2 Effect of N^* coverage on surface morphology

Ertl and collaborators [24] observed that, at low temperatures, N^* orders into patterns commensurate with the underlying periodicity of the Fe(111) lattice. Experimental studies with multipromoted industrial catalysts indicate a steady-state coverage of nitrogen-containing intermediates around 50% at $T = 723$ K [212], and a coverage-dependent increase of the apparent activation energy to dissociation [213]. Therefore, it is important to ascertain whether the dynamic behavior of the (111) iron surface persists and whether and how the nitrogen poisons the catalyst, either by blocking active sites or inhibiting molecular N_2 adsorption [214]. To this effect, we simulate at room and at the *operando* temperature of 700 K the static, dynamic, and catalytic properties of Fe(111) for N^* coverages of up to 50%. As we learned in previous chapters, the catalyst behaves rather differently in these two regimes. The MD simulations are performed using the expanded version of the MLP, as described in Section 2.7.2.

At $T = 300$ K, we find that the ordered patterns of N^* [24] appear to be stable. In contrast, at $T = 700$ K, significant surface dynamics set in, leading to a structural heterogeneity in the nitrogen-covered surface. As for the clean one, this mobility results in the formation of defective structures. In particular, we observe triangular Fe motives described by the formula $Fe_3^*N_i^*$ $i = 1, \dots, 4$, where the Fe atoms form an equilateral triangle on top of a vacancy and the N^* sit either at the corners or at the center (see Figure 5.2). These Fe_3^* sites exhibited coordination with N^* ranging from 1 to 4. Triangular Fe structures can be fleetingly observed also on a clean surface, but they are stabilized by the presence of adjacent N^* atoms. These structures have already been described in Section 3.1 in the context of ammonia decomposition and

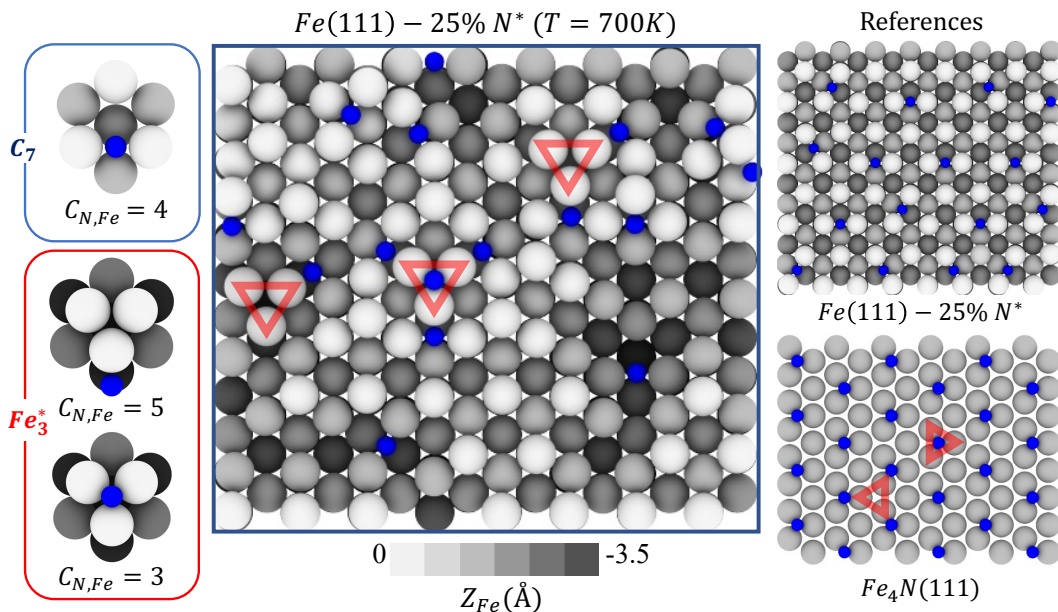


Figure 5.2: Snapshot of the $Fe(111)$ surface covered with 25% N^* during a molecular dynamics simulation at 700 K (center). Fe atoms are color-coded according to their height along the $[111]$ axis, the N^* atoms are colored in blue. The Fe_3^* sites, each surrounded by varying numbers of N^* atoms, are indicated by red triangles. Different coordination environments of N^* are illustrated on the left. In the top right panel is reported a reference $Fe(111)$ with 25% N^* coverage, obtained by optimizing the experimentally suggested structure at 300 K, in which all N^* occupy four-fold coordinated (C_7) sites. On the bottom right panel, a picture of the $Fe_4N(111)$ surface is presented, where the triangles analogous to those observed on the $Fe(111)$ surface at 700 K are highlighted.

are of particular relevance for the adsorption of NH (see Section 3.2.2).

For the triangular Fe_3^* moiety to be stable, at least one N^* needs to be coordinated to it. The coordinated N^* can sit either at the corner of the Fe_3^* triangle and have five-fold coordination or at its center and have three-fold coordination. The N^* atoms not coordinated to a Fe_3^* moiety are instead four-fold coordinated and occupy the C_7 catalytic centers; see the snapshots in Figure 5.2. Thus, we can distinguish the N^* adsorption sites according to their Fe coordination number. The distribution of the standard and triangular adsorption sites is reported in Figure 5.3, showing that about half of the N^* is associated with triangles, while the other half occupies the C_7 catalytic sites.




Figure 5.3: Fraction of N^* within different coordination environments: four-fold ($C_{N,Fe} = 4$), triangle ($C_{N,Fe} = 3$ or 5), and transitional (others) during molecular dynamics simulations at 700 K. The triangle group refers to the N^* coordination facilitated by Fe_3^* moieties, with light colors indicating $C_{N,Fe} = 3$ (center of Fe_3^*) and dark colors representing $C_{N,Fe} = 5$ (corners of Fe_3^*). Coordination numbers were computed using a cutoff of 2.5 Å. Error bars represent the fluctuation during the temporal evolution of the system at 700 K.

5.2.1 Triangular motives validation with DFT

To ensure that this phenomenon was not an artifact of our machine-learning potential, we validated it via DFT calculations performed on a small system (144 Fe atoms supercell) with the same orientation. We considered an initial structure obtained by introducing a single vacancy on a second layer beneath the surface, as depicted in Figure 5.4 (a). By performing a DFT relaxation, we observe the formation of a triangular Fe_3^* motif on top of the vacancy, see Figure 5.4 (b). This confirms that the triangular motives are metastable states induced by vacancies.

Subsequently, we systematically explored various N^* concentrations ranging from 1 to 4 N^* atoms and their all possible arrangements around the Fe_3^* motif. For each of them, we performed a DFT relaxation and we reported the results in Figure 5.5. The first column is the reference configuration, where N^* atoms reside on conventional four-fold coordinated sites, which are the preferred locations in the absence of Fe_3^* . The other columns describe the configurations in which N^* are around the triangular moiety. These calculations reveal that the arrangement of N^* atoms around Fe_3^* is energetically more favorable, and that the relative stability of this configuration increases as the number of N^* atoms increases.

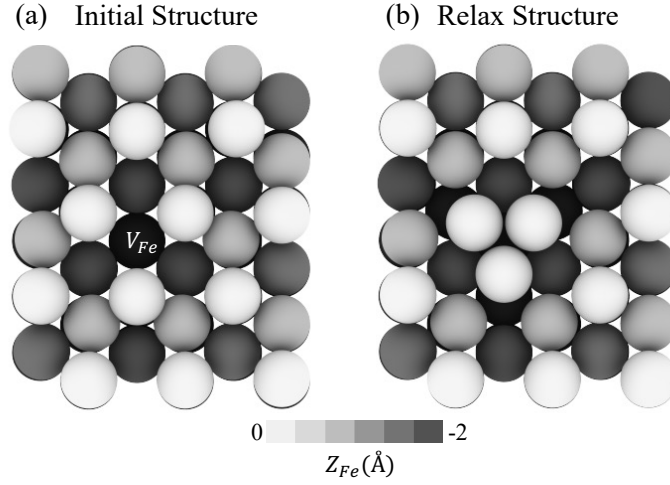


Figure 5.4: (a) The initial structure of BCC Fe(111) with the introduction of a single iron vacancy (marked as V_{Fe}) on the second layer. (b) The relaxed structure at 0 K obtained through density functional theory (DFT), illustrating the emergence of Fe_3^* on top of the vacant Fe site. Iron atoms are colored based on their z coordinate as shown in the scale bar.

5.2.2 $N^* - N^*$ pair interaction

We then further investigated the structural properties of the surface in the presence of adsorbed N^* . In particular, we measured the pairwise correlation of N^* and its relationship with the Fe substrate.

MD simulations reveal that triangular complexes induce clustering of adsorbed atoms, as evidenced by the presence of a peak in the distribution of N^* distances at values much lower than the average distance for a uniform distribution (see $T = 300$ K case in Figure 5.6). The strong short-range pair correlation between N^* atoms at *operando* temperature is consistently observed for all of the different coverages studied (12.5%, 25%, 50%), see Figure 5.7.

The strong short-range pair correlation between N^* atoms at *operando* temperature compared to room temperature, can be attributed to two possible origins:

- the formation of $Fe_3^*N_i^*$ motifs, leading to the clustering of N^* atoms;
- the diffusion of N^* atoms at *operando* temperature, allowing them to dynamically alter the $N^* - N^*$ pair interaction, which is absent at room temperature.

To better understand their respective contributions, we segmented the histogram into two parts: one representing the distribution from N^* atoms in standard four-fold coordination and the other from those belonging to triangular motifs, as shown in Figure 5.6 and Figure 5.7. It is evident that the formation of $Fe_3^*N_i^*$ plays a major

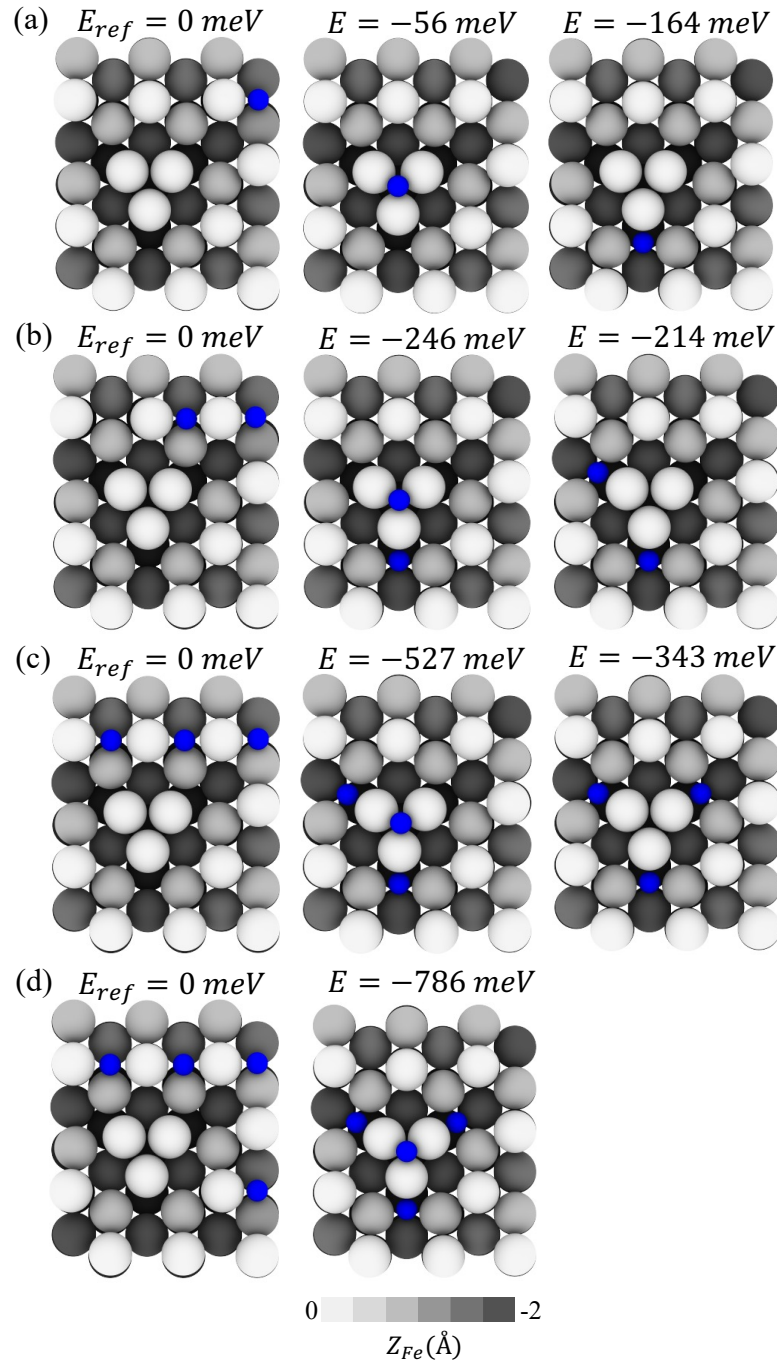


Figure 5.5: Relaxed structures and energetics of all possible configurations (a) 1 N^* , (b) 2 N^* , (c) 3 N^* , and (d) 4 N^* bound to Fe_3^* using density functional theory (DFT) calculations at 0 K. The reference structures are derived from arrangements in which all N^* are four-fold coordinated, representing their equilibrium configuration in the absence of Fe_3^* . The relaxed energies from these states are used as the energy reference for each respective N^* concentration.

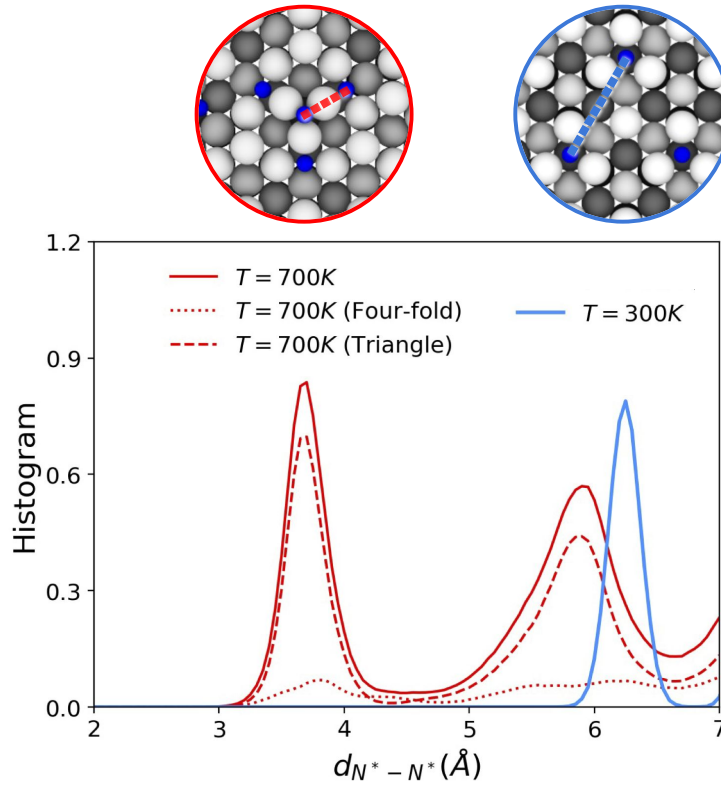


Figure 5.6: Average distribution of the $N^* - N^*$ distances on the 25% covered Fe(111) surface at 700 K (see Figure 5.7 for the other concentrations). In the figure are shown the histograms of distances; between all atoms (solid red line), between all pairs of atoms in which at least one is coordinated to a Fe_3^* triangle (dashed line), and between uncoordinated atoms ($C_{N,Fe} = 4$) using dotted line. We also contrast the $N^* - N^*$ histogram at 700 K (solid red line) with the corresponding histogram at 300 K (solid blue line). Above the picture, we report two snapshots highlighting the distances that contribute to the dominant peaks at low and high temperatures.




Figure 5.7: Average distribution of the $N^* - N^*$ distances on Fe(111) surface with various N^* coverages at 700 K. In the figure are shown the histograms of distances; between all atoms (solid red lines), between all pairs of atoms in which at least one is coordinated to a Fe_3^* triangle (dashed lines), and between uncoordinated atoms ($C_{N,Fe} = 4$) using dotted lines.

role in the clustering of N^* atoms. To further confirm this hypothesis, we simulated an artificial system for 25 % coverage, wherein the iron atoms remained at room temperature while the N^* atoms were allowed to diffuse by increasing their temperature to 1200 K, through the application two different thermostats. This strategy also solves a problem related to simulating the room temperature distribution. In fact, due to the time scale limitations, the distribution of pair distances at room temperature is solely determined by the initial arrangement of N^* , since diffusion can be considered a rare event at low temperatures. We find that the distribution for $N^* - N^*$ pair distances for this artificial system closely resembles the four-fold case, see Figure 5.8, and as such, it cannot explain the intensity of the first peak.

We can thus conclude that the $N^* - N^*$ effective interaction is modulated by the substrate. It also warns against the danger of treating the surface as static, even if the reagents are allowed to move.




Figure 5.8: Average distribution of the $N^* - N^*$ distances on the 25% covered Fe(111) surface at 700 K. In the figure are shown the histograms of distances; between all atoms (solid red lines), between all pairs of atoms in which at least one is coordinated to a Fe_3^* triangle (dashed lines), and between uncoordinated atoms ($C_{N,Fe} = 4$) using dotted lines. We also contrast the histogram of N^* pair distances (dashed blue line) for an artificial system where iron atoms are maintained at room temperature while N^* adsorbates are heated at 1200 K to allow them to diffuse. The distribution for this artificial system is similar to the four-fold case.

5.2.3 Effective N^* coverage

This clustering leads to a decrease in the effective N^* surface coverage, as evidenced in Figure 5.9, where the effective coverage at 700 K is shown to be lower than the actual N^* coverage. This is determined by accounting for the surface area that remains free from Fe_3^* clusters and by excluding N^* atoms that are part of Fe_3^* motifs, which have coordination numbers $C_{N,Fe} \neq 3$ or 5. The solid circles and error bars in the figure represent the average effective coverage and fluctuations observed during molecular dynamics simulations. Thus, at high temperatures, the clustering effect reduces the




Figure 5.9: Effective N^* coverage as a function of the actual N^* coverage at 700 K, determined by considering the surface area free from Fe_3^* and excluding N^* atoms that are part of Fe_3^* ($C_{N,Fe} \neq 3$ or 5). Solid circles and error bars represent the average value and the fluctuations during molecular dynamics simulations.

poisoning of the surface, leaving a greater portion of the surface available for nitrogen dissociation.

Moreover, the probability of finding an N^* atom near a C_7 site is significantly lower at 700 K compared to 300 K, as depicted in Figure 5.10. In the figure, we report the histogram of the distances between C_7 sites and N^* atoms for various N^* concentrations at both temperatures. The histogram is normalized based on the total number of C_7 at each timeframe t evaluated, $\sum_t N_{C_7}(t)$. The $C_7 - N^*$ distance is much greater at 700 K (represented by the red line) compared to 300 K (blue line). This indicates that at higher temperatures, the surface dynamics contribute to a reduced likelihood of N^* atoms occupying the vicinity of C_7 sites, further alleviating the poisoning effect observed at lower temperatures.




Figure 5.10: Histogram illustrating the distribution of $C_7 - N^*$ distances on the Fe(111) surface covered with various N^* concentrations during molecular dynamics (MD) simulations at 300 K and 700 K. The histogram values are normalized based on the total number of C_7 sites ($\sum_t N_{C_7}(t)$). The blue and red lines are used for 300 K and 700 K, respectively.

It is important to notice that these triangular motifs are characteristic of the Fe_4N (111) surface, which corresponds to the stable nitride phase under *operando* conditions [209]. The observation of these triangular motifs on the Fe(111) surface under high nitrogen coverage can be interpreted as early-stage precursors of a transition towards the more stable nitride phase, though this transition remains frustrated under the conditions studied. This behavior is reminiscent of the hypothesis proposed by Ertl and coworkers [24], who suggested that “surface nitrides” structures related to but distinct from bulk Fe_4N , are formed during nitrogen dissociation on Fe surfaces.

This behavior turns out to be, although different, intrinsically related to what was shown in the previous Chapter in the context of ammonia decomposition on Fe(110) surface. In that case, the “hardness” of the surface does not allow the formation of the surface nitride phase, and the transition occurs in the bulk, as observed experimentally. Here, the “softness” of the surface allows a nitride phase to emerge directly at the surface.

5.2.4 Dynamical properties

Having described the structural changes induced by N^* , we focus on the dynamical properties of the surface. From a visual inspection of the simulations, the atoms involved in the $Fe_3^*N_i$ motifs have slower dynamics compared to the other atoms. Thus, we calculate the diffusion coefficient D of Fe surface atoms and N^* according to Equation 2.38 (see section 2.6.2.2). In Figure 5.11 we report Fe surface atoms diffusion coefficient $D_{Fe_{surface}}$ as a function of N^* surface coverage at 700 K. The diffusion coefficient $D_{Fe_{surface}}$ clearly decreases monotonically as the value of coverage increases.

However, the average iron diffusion remains sizable, allowing surface dynamics to affect the system behavior also at finite coverages.




Figure 5.11: Diffusion coefficient of Fe surface atoms as a function of N^* surface coverage at 700 K.

Given the limited number of N^* , we were not able to get enough statistics to pin down the value of the N^* diffusion coefficient with acceptable statistical confidence, but for the case of the 50% coverage where we could estimate a diffusion coefficient value of $D_{N^*} = (1.7 \pm 0.1) 10^{-8} cm^2/s$ which is one order of magnitude higher than surface Fe atoms at the same N^* coverage.

To study the behavior of local structures during simulations, we employ the Environment Similarity metric as described in Section 2.6.2.1. As the concentration of N^* increases, the density of the Fe_3^* triangles increases while that of the C_7 -like environments decreases, as can be seen in Figure 5.12. Both surface features have a finite lifetime, as reported in Figure 5.12 (bottom panel). The lifetime distributions are non-Gaussian, with long tails towards long lifetimes. It is important to note for the catalytic activity that the average lifetime of the C_7 catalytic environments increases as a function of coverage.

At the end of this section, we want to make a remark on the complex dynamical properties of the surface, especially under varying N^* coverage. First, the dynamics of the iron atoms enable the formation of triangular $Fe_3^*N_i$ motifs, where nitrogen atoms tend to cluster, while still leaving some C_7 catalytic sites free, thus effectively reducing the surface coverage and mitigating the poisoning effect. Second, as the N^* coverage increases, the mobility of surface atoms decreases, leading to a general reduction in surface dynamism. However, this reduction in mobility extends the lifetime of the remaining free C_7 catalytic sites, allowing them to remain active

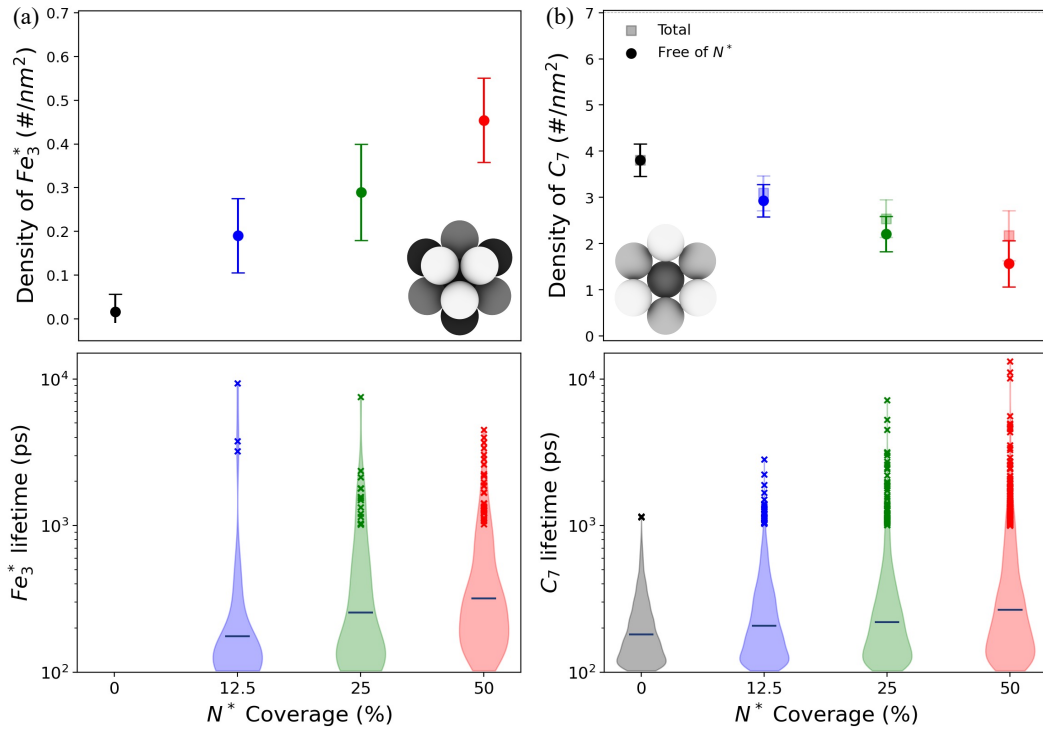


Figure 5.12: Surface density and violin plot of lifetime distribution for a) Fe_3^* and b) C_7 sites. These are calculated using the Environment Similarity metric as described in Section 2.6.2.1. For the surface density of C_7 , faded squares indicate the overall density of C_7 , while solid circles denote the density of C_7 sites unaffected by N^* poisoning. In the lifetime distribution plots, only Fe_3^* and C_7 sites with lifetimes of at least 100 ps are considered. Sites with lifetimes surpassing 1 ns are indicated by x markers.

for longer periods despite the high nitrogen coverage. This dynamic balance between clustering, mobility reduction, and catalytic site preservation is key to sustaining catalytic activity at elevated coverages, as we will address in the following section.

5.3 Effect of N^* coverage on N_2 adsorption and decomposition

In this section, we study how this complex dynamics affects the ability of the N^* -covered surface to adsorb and dissociate an incoming N_2 molecule. In Ref. [44], we found that two collective variables could well describe the behavior of this process: the N_2 bonding length d and the charge transferred from the metallic surface to the nitrogen molecule q . As we extensively presented in the previous Chapters, the changes in q are a rather reliable indicator of the underlying chemical process and

allow classifying the adsorption states without having to specify explicitly the geometrical arrangement. This quantity is particularly effective also for describing the N_2 dissociative chemisorption at high temperatures [44]. The q values were computed by training a graph neural network on the data collected in Ref. [44] supplemented by the new ones obtained here for the covered surface, see Section 2.5.




Figure 5.13: The free energy is plotted as a function of N - N distance ($d(N, N)$) and the partial charge of N_2 ($q(N_2)$) at 300 K for various N^* covered surfaces. The local minima on the plot correspond to metastable states, and the white dashed lines indicate the minimum free energy pathway (s_m) in this plane. The reference is set to 0 for the lowest energy state (α).

In Figures 5.13 and 5.14, we report the free energy profiles along $d(N, N)$ and $q(N_2)$ variables ($F(d, q)$) obtained at the various N^* surface coverages (0%, 12.5%, 25%, 50%) at 300 K (Figure 5.13) and 700 K (5.13). These free energy surfaces provide a comprehensive description of how the adsorption and dissociation of N_2 evolve as a function of both the N - N bond distance $d(N, N)$ and the charge transferred to the nitrogen molecule, $q(N_2)$. Before going into a detailed description of the process, we want to note how remarkable it is that the same variables can also be used profitably in the more complex situation described here. The minima observed along the $d - q$ plane correspond to different adsorption configurations, each with its specific stability. The metastable states reflect various stages of molecular adsorption before full dissociation. The free energy landscape suggests that at 300 K temperature (Figure 5.13), dissociation barriers increase with N^* coverage, indicating that nitrogen




Figure 5.14: Analogous of Figure 5.13, at 700 K.

poisoning effectively hinders the adsorption process. In this scenario, the surface is less dynamic, and N^* atoms occupying active sites impede further N_2 dissociation. At 700 K, as depicted in Figure 5.14, the scenario changes significantly. The increased temperature leads to enhanced surface mobility, which, as discussed in earlier sections, mitigates the poisoning effect caused by N^* atoms. Despite the presence of N^* , the dissociation barrier remains relatively unchanged, as shown by the similar free energy profiles across different coverages. This can be attributed to the dynamic surface rearrangements that create transient adsorption sites for N_2 .

The white dashed lines along the free energy surfaces represent the minimum energy pathway s_m on the space (d, q) , which indicates the most probable trajectory that N_2 follows during adsorption and dissociation and results very little depending on the system studied. For clearer readability, we have reported the free energy profiles along minimum energy pathways $F(s_m)$ in a single panel (Figure 5.15b) in which we also report again the 2D profile at 300 and 700 K for the coverage of 25% N^* (Figure 5.15a). At $T = 300$ K, one can recognize the classical adsorption sites of the low-temperature studies of nitrogen chemisorption on iron [32, 202, 215]. As the charge transfer increases, the adsorption sites go from a vertical γ/δ to a horizontal α , and to the precursor horizontal α' state in which the molecule sits in the second layer just above the C_7 iron atoms. At $T = 700$ K, the overall qualitative behavior is similar with the notable exception of the instability of the α' state due to the surface

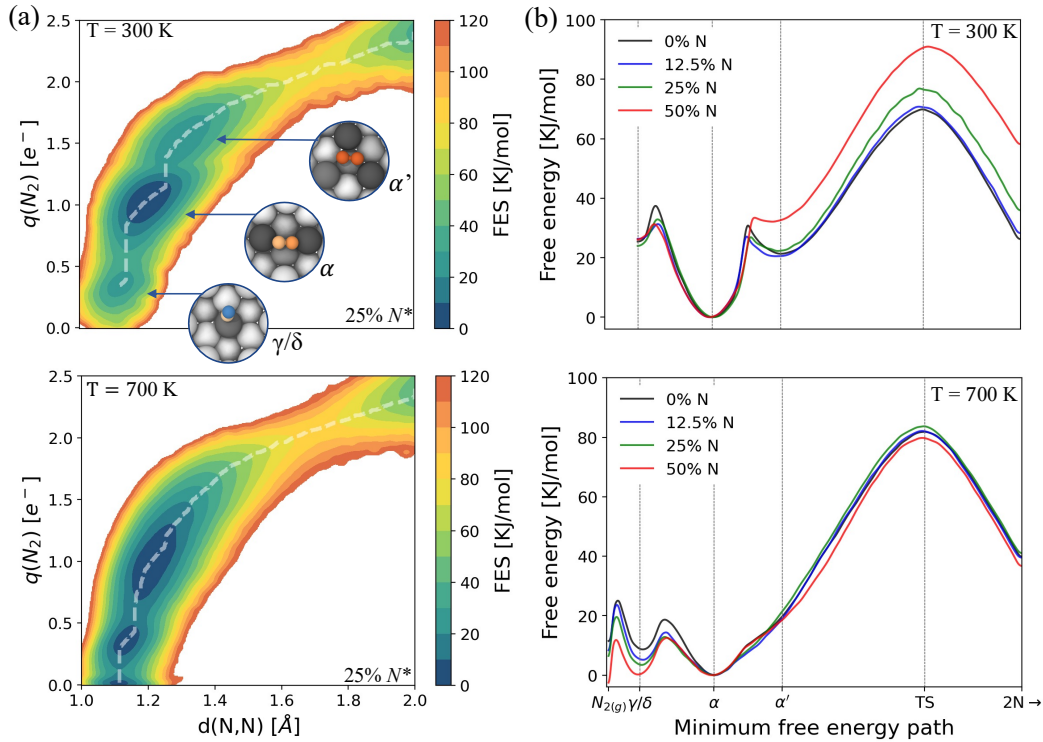


Figure 5.15: (a) The free energy, $F(d, q)$ is plotted as a function of N-N distance and the nitrogen charge $q(N_2)$ (defined as the sum of the partial charges of N atoms) at 300 K and 700 K for a surface covered with 25% N^* . Results for the other concentrations are reported in Figures 5.13 and 5.14. The local minima on the plots correspond to metastable states, and the white dashed lines indicate the minimum free energy pathway (s_m) in this plane. The snapshots of various adsorption sites (γ/δ , α , and α') are shown in the top panel of (a). (b) The minimum free energy profile ($F(s_m)$) computed along the s_m in $d - q$ space for the Fe(111) surface at 300 K and 700 K. In plotting the $T = 700$ K, we have not considered those transitions that take place on the Fe_3^* triangles whose energy barrier is much higher (120 kJ/mol) than the one associated with the C_7 sites. See Figure 5.17 for the free energy profiles computed using all pathways. The reference is set to 0 for the lowest energy state (α). Note that, for 10 bar N_2 partial pressure, we do not find $N_2(g)$ to be a metastable state at 300 K.

dynamics, already noted in Ref. [44] for the clean surface.

A more punctual comparison between the systems at different temperatures and coverages can be made via the minimum free energy paths, see Figure 5.15b. In this representation, the difference between the low and the high-temperature behavior is more striking. In the $T = 300$ K case, where the atom mobility is close to zero and the surface structure is ordered, the free energy barrier between the α state and dissociation increases with coverage reflecting the poisoning effect of the N^* atoms,

as observed by Ertl experiments at low temperature on Fe(111) [24]. In noticeable contrast, at $T = 700$ K, the free energy barrier relative to the α state hardly changes. Two main physical effects conspire to yield this result:

- the surface atoms mobility, which is reduced but not quenched by the N^* . This pushes the lifetime of the C_7 sites to higher values.
- the looming transition to a nitride phase. This leads to the formation of $Fe_3^*N_i^*$ complexes and to a clustering of N^* atoms, which reduces the number of N^* blocking the active sites, see Figure 5.12.

These two effects compensate for the decrease in available C_7 environments, leading only to very small changes in the free energy barrier. On the other hand, the free energy difference between the gas phase and the α state is affected by the coverage, leading for the 50% case to an increase of the apparent activation energy of ~ 15 kJ/mol, in agreement with the estimate based on the analysis of an industrial catalyst behavior [213]. Once again, as in the case of the clean surface [44] and the ammonia decomposition (see Chapter 3), we find that there are two different behaviors depending on temperature, and in particular that the effect of lateral interactions with N^* at high temperature cannot be derived by extrapolating the results obtained at low ones.

5.3.1 Reaction pathway analysis

From a detailed observation of the reactive trajectories at *operando* temperature, we observe that the adsorption and dissociation of N_2 can occur through two distinct and independent pathways. One pathway closely resembles the low-temperature route (referred to as the C_7 pathway), while the other takes place atop the Fe_3^* moieties (designated as the Fe_3^* pathway). To discern their respective $F(d, q)$ profiles, we divide the configurations based on the minimum distance between the N_2 molecule and all three Fe atoms of Fe_3^* . If this minimum distance falls within 3.5 \AA of all three atoms of Fe_3^* , the configuration is categorized as part of the Fe_3^* pathway; otherwise, it pertains to the C_7 pathway. In Figure 5.16, we depict the $F(d, q)$ curves along these distinctive pathways for 12.5% N^* coverage case. Similar results have been obtained for the other coverages for which we report only the one dimensional $F(s_m)$ (see Figure 5.17). For the Fe_3^* pathway, only one adsorption state emerges as a metastable equilibrium state, with a $q(N_2)$ value resembling the α state in the C_7 pathway. On the other hand, the C_7 pathway resembles the general (all pathways) free energy landscape, previously presented in Figure 5.14.

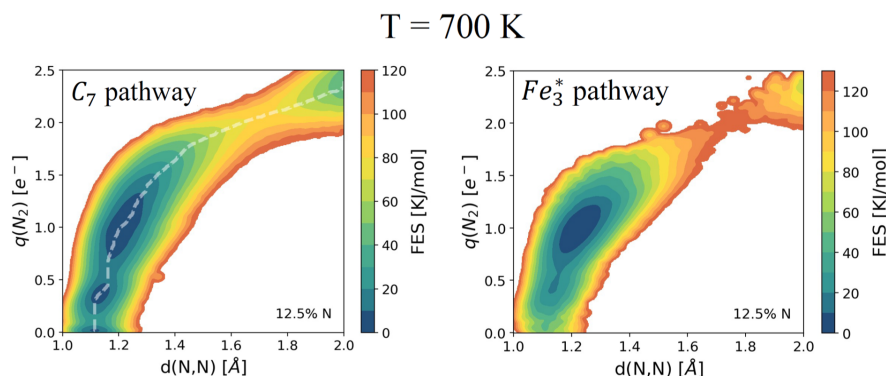


Figure 5.16: The free energy is plotted as a function of N - N distance ($d(N, N)$) and the partial charge of N_2 ($q(N_2)$) for the reaction pathway on the C_7 (left) and Fe_3^* (right) at 700 K for various 12.5% N^* coverage. The configurations in which the minimum distance between N_2 and any atom of Fe_3^* is more than 3.5 \AA are used for the free energy calculation of C_7 pathway, the remaining ($\min(d(N_2, Fe_3^*)) < 3.5 \text{ \AA}$) are used for the Fe_3^* pathway. The local minima on the plot correspond to metastable states, and the white dashed lines indicate the minimum free energy pathway (s_m) in this plane. The reference is set to 0 for the lowest energy state.

The global $F(s_m)$ profile, along with the individual $F(s_m)$ profiles for the two pathways, are shown in Figure 5.17. Notably, the dissociation barrier for the Fe_3^* pathway is significantly higher ($\approx 120 \text{ kJ/mol}$) compared to the C_7 pathway ($\approx 80 \text{ kJ/mol}$). Consequently, as it is anticipated, the reaction predominantly proceeds through the C_7 pathway.




Figure 5.17: The free energy profile computed along the minimum free energy pathway in the 2-dimensional $d(N, N)$ $q(N_2)$ space for the Fe(111) surface at 700 K (left panel) using all reactive pathways and (right panels) resolved for C_7 and Fe_3^* pathways.

Discussion

In conclusion, in this Chapter we showed that the overall catalytic activity of the Fe(111) surface is due not only to the existence of catalytic sites but also to the properties of the surface as a whole, its dynamics, and its fluctuations. The closeness to a phase transition toward the nitride phase helps to induce the kind of heterogeneous static and dynamical behavior observed here and to sustain the fluctuations needed for the catalytic process to be stabilized. A beneficial effect on the catalyst resistance to poisoning comes from a combination of surface dynamics and chemical processes-induced changes. Similar observations of the dramatic effects of surface dynamics and reaction-induced changes can be found in the literature, with notable examples being the activation of the copper surface by carbon monoxide [56] and the destabilization of lithium imide [66] and barium hydride [65] surfaces under ammonia decomposition. Once again, our study demonstrates that a static description of one such catalyst is inadequate. For the Haber-Bosch reaction, this was suggested some time ago by Spencer [54], when he argued that for an industrial catalyst to be as stable as it is, a steady state dynamical process needs to be set up. He indicated two possible ways in which this could happen: either the surface atoms becoming mobile or the reaction itself dynamically changing the surface. One can recognize that in this process, the two scenarios combine into one.

It's also interesting to compare these results with what we reported in Chapter 4. For the case of Fe(110), we have observed how nitrogen dissolves in bulk, explaining the formation of bulk nitrides. Instead, for the Fe(111), the surface species and iron mobility favor the formation of a nitride pattern directly on the surface. In the case of Haber-Bosch ammonia synthesis, this implies not only the aforementioned beneficial effect on the active sites for N_2 dissociation but also allows for exposing N^* on the surface without dissolving in bulk. The latter is crucial for the N^* to be hydrogenated. In contrast, the N dissolution and consequent bulk nitride formation observed for the Fe(110), will not lead to hydrogenation steps and, thus, the reaction to proceed. This might offer an additional explanation for the dissimilar activity observed across different surfaces[31]. Moreover, it further demonstrates how, in heterogeneous catalysis, the role of the catalyst is much more complex than can be modeled with simplified models.

Finally, we note that the picture depicted here could be experimentally checked with modern *operando* surface spectroscopy techniques and these experiments could also be used for studying the nucleation of the nitride phase [56, 216]. Notwithstanding the economic and time burden of obtaining and studying a Fe(111) catalyst obtained from a single crystal, we hope that this can be done in the near future.

Chapter 6

A Data-Efficient protocol for modeling catalytic reactivity

As we have shown in previous Chapters, studying catalytic reactivity under operative conditions poses a significant challenge due to the dynamic nature of the catalysts. These studies were possible thanks to the combination of machine learning and enhanced sampling, which allows for the simulation of dynamics at a fraction of the cost. Still, to build accurate and robust MLPs, which are needed for modeling catalytic reactivity *in operando* conditions, all relevant configurations, particularly reactive ones, have to be included in the dataset. As seen in Section 2.3, active learning techniques can effectively accomplish this task, particularly when combined with enhanced sampling. We have demonstrated this in the construction of potential for ammonia synthesis/decomposition on Fe surfaces, where we applied active learning through the *query-by-committee*, nowadays considered as a standard approach. However, this procedure is effective but not efficient, and it could prevent studying systems whose electronic structure description is more expensive. In fact, the total amount of DFT calculations performed for Fe studies is of the order of hundreds of thousands (see Section 2.3.1), which could be prohibitively large.

This is, for instance, the case of Fe-Co alloys, which we were interested in studying as they showed an improved catalytic performance for NH_3 decomposition with respect to Fe. We were thus prompted to search for new strategies to efficiently construct reactive potentials with a minimum number of QM calculations. More generally, this is also critical to enable their widespread use. In this regard, recent developments in data-efficient architectures can help in reducing the number of points needed to train a robust model. These advances include equivariant GNNs (Section 2.2.3.2) or transfer learning approaches [217] leveraging foundational models [218, 219]. However, these techniques still do not solve the crucial issue of identifying the few relevant

configurations to be included in the reference dataset set.

In the following, we present a scheme we developed to construct a reactive potential for FeCo in a data-efficient way. This is achieved by combining enhanced sampling methods first with GP to discover transition paths and then with GNN to obtain a uniformly accurate description. The necessary configurations are extracted via a novel active learning procedure we named Data-Efficient Active Learning (DEAL) based on local environment uncertainty. The protocol has been developed with the aim of studying the NH_3 decomposition on FeCo, but it can be applied to general reactions. At the end of the chapter, we will provide an extended validation of the protocol and its general features, leaving the presentation of simulation results of the new catalyst to the next chapter.

6.1 A data-efficient protocol

Despite the success of enhanced sampling-based active learning, these studies often required many iterations and numerous DFT calculations, leaving room for improvement. Two aspects should be considered to make the construction of MLPs more efficient. At first, when there is little data available, incremental learning [220] approaches should be used in which the potential is frequently updated. Otherwise, by pushing the system out of equilibrium with enhanced sampling, we risk extrapolating poorly and sampling the wrong configurational regions. On the other hand, we also need an efficient criterion for selecting structures for DFT calculations from those obtained in the active learning process.

Here, we introduce a new scheme to efficiently construct reactive potentials, leveraging both enhanced sampling methods and on-the-fly selection of relevant structures. This is achieved following a two-stage protocol: an exploratory phase to harvest an initial pool of reactive configurations and a second stage in which we obtain a uniformly accurate description of the transition pathways (see Fig. 6.1 for a diagram of the protocol). For the second step, we developed a new active learning scheme, named Data-Efficient Active Learning (DEAL), which allows us to identify a non-redundant set of structures to be added to the training set. As an enhanced sampling method, we used OPES [157] (see Section 2.4.2) with its different variants that can be used to explore or converge the free energy landscape. Furthermore, another important ingredient is Gaussian processes (GPs) [221], which we used first to learn the potential energy surface on the fly and then to identify novel local environments to build a minimal data set within our active learning scheme.

We illustrate these methodological advances by studying several reactions related to ammonia decomposition over FeCo alloy catalysts. The ferromagnetic and

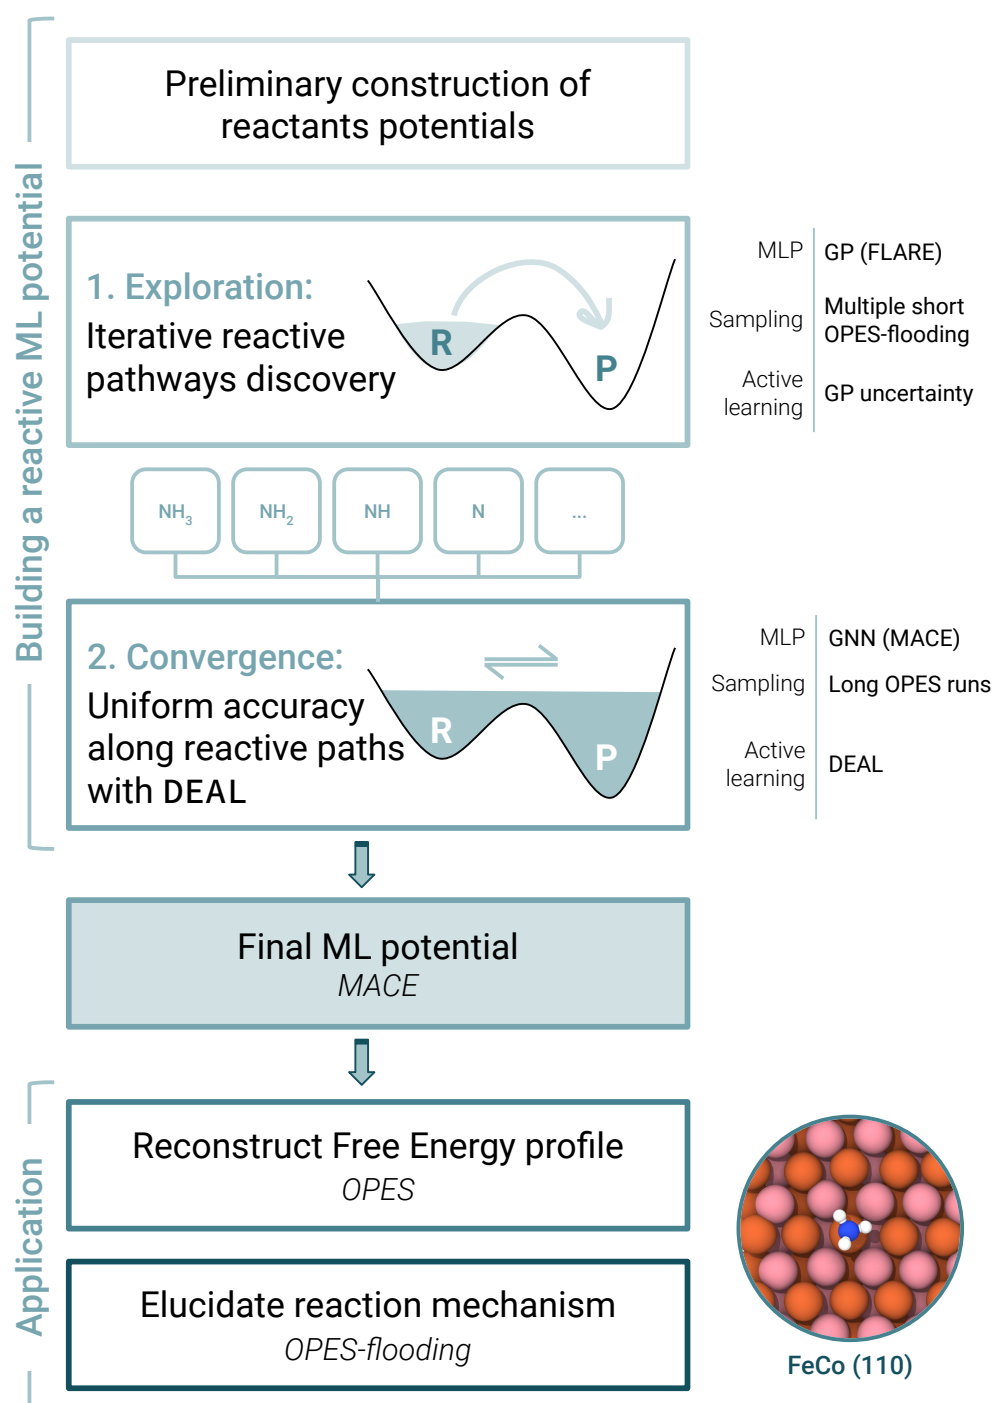


Figure 6.1: **Diagram of the protocol** used to obtain uniformly accurate and data-efficient reactive potentials. On the right, we have reported some details regarding the ML potential, enhanced sampling scheme, and selection criterion used in the different stages.

metallic properties of the surface make the description of the electronic structure particularly computationally expensive, and this is why we needed to develop this computationally efficient methodology. Although it is a new system, it shares many structural properties with Fe, so many of the variables and expertise developed in the iron studies will be exploited here.

We focus here on the description of the method and the construction of the MLP for FeCo alloys, while the results on the catalytic behavior and activity of FeCo will be presented in Chapter 7.

6.1.1 Preliminary construction of reactant potentials

In our pursuit to model the catalytic process of ammonia decomposition FeCo surfaces, we initiated our investigation by constructing ML potentials for the reactants. This involved gathering configurations to characterize the pristine 110 surface (see Fig. 6.2) and the different intermediate species adsorbed on it (NH_3 , NH_2 , NH , N , N_2 , H ,...). To efficiently accomplish this task, we employed GPs to learn on-the-fly




Figure 6.2: NH_x adsorbates on the FeCo (110) surface, top and left view.

the potential energy surface, with the sparse implementation of FLARE [115] using the Atomic Cluster Expansion (ACE) descriptors [116], see Section 2.7.4.2 for details. Recognizing the limitations of GPs with large training datasets [221], we trained separate models for each intermediate species.

Initially, we generated a dataset by conducting a set of uncertainty-aware molecular dynamics (MD) simulations based on GPs at the *operando* temperature of $T=700$ K. Subsequently, we performed simulations at higher temperatures (up to 900 K) to diversify our configurations and capture surface dynamics. In addition,

to obtain an exhaustive coverage of the reactant space, we carried out enhanced sampling simulations to explore the various adsorption sites and the diffusion of the molecules on the surface. This preparatory stage produced about 2500 configurations for all the different intermediates.

Given the complexity of the reaction pathways and the potential existence of multiple channels, we approached the collection of reactive configurations through a two-step process. In the first, the reactive pathways are discovered, while in the second, the description of these pathways is improved until high accuracy is achieved.

6.1.2 Reactive pathways discovery via uncertainty-aware flooding simulations

The initial step in harvesting reactive configurations is the discovery of reaction pathways and transition state structures. While for gas phase reactions, a simple interpolation between reactants and products can provide a reasonable guess, for catalyzed reactions, this is not possible, especially at high temperatures. Indeed, the active site is not known beforehand or may change due to dynamics [44, 222] or, again, there may be multiple reaction pathways that need to be sampled [222]. To address this challenge, we used OPES to perform a set of “flooding-like” [223] simulations together with uncertainty-aware MD (see Section 2.7.4.2). As we have seen, OPES-flooding introduces an external potential to fill the reactant basin and then let the reactive event occur as spontaneously as possible along the low free-energy pathways. This allowed us to sample reactive processes with minimal knowledge of the reaction mechanism. The only requirement for these methods is the definition of collective variables (CVs) that can distinguish between reactant and product states. We note that, in the case where the products of the reaction to be studied are not known, generic CVs could be used to discover the possible products [224]. In addition, the integration with uncertainty-aware MD brings two significant advantages: it allows for an efficient selection of reactive configurations based on the uncertainty of the local environments, and it updates the potential energy surface model in an incremental manner, correcting wrong extrapolations to nonphysical regions of phase space.

To illustrate our workflow and demonstrate its effectiveness in sampling reactive pathways, we start with the case of N_2 dissociation, a process that we have extensively examined on Fe surfaces (see Ref. [44] and Chapter 4). Nitrogen molecules can adsorb onto the (110) surface in two main configurations, either parallel (α) or perpendicular (γ) to the surface (see Fig. 6.3a). *A priori*, we do not know from which adsorption site the molecule may break. To find out, we performed a set of iterative OPES-flooding

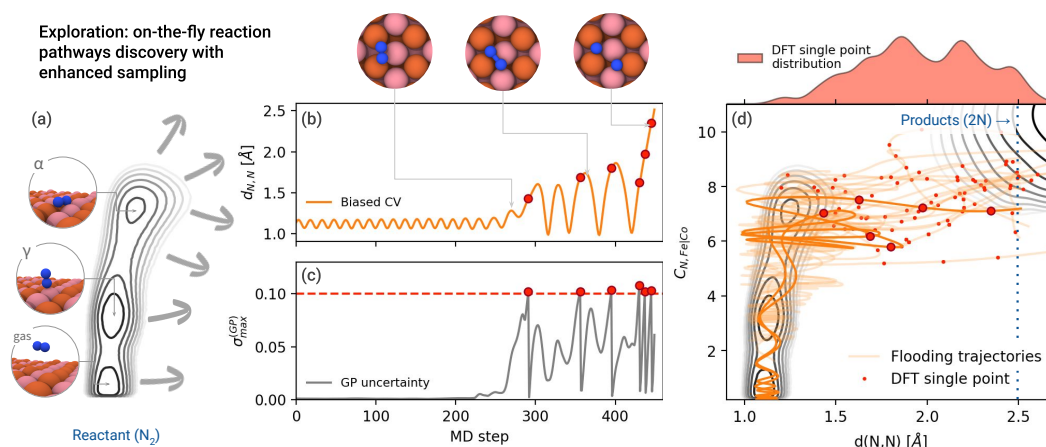


Figure 6.3: **Exploration stage: iterative discovery of reaction paths through flooding simulations and GP-based MD.** We report a schematic presentation for the case of N_2 dissociation. a) The reactant free energy landscape constructed at the end of the preliminary phase, highlighting the metastable minima (gas phase and two adsorption states). Arrows indicate that the reaction paths toward the product state are unknown *a priori*. Several iterations are then initiated from the reactant basin. The time evolution of the distance between N atoms $d_{N,N}$ (b) and the GP maximum uncertainty on the local environments (c) are shown for a specific iteration. Every time the uncertainty of the GP exceeds the threshold of 0.1 (red dotted line), energy and forces are recomputed at the DFT level (red dots), and the GP is updated. Panel (d) shows all iterations in the 2D space defined by $d_{N,N}$ and the coordination number between N and surface atoms $C_{N,Fe|Co}$, with the trajectory shown in (b,c) highlighted. Above the top axis, we report the distribution of points where DFT calculations were performed, projected along $d_{N,N}$.

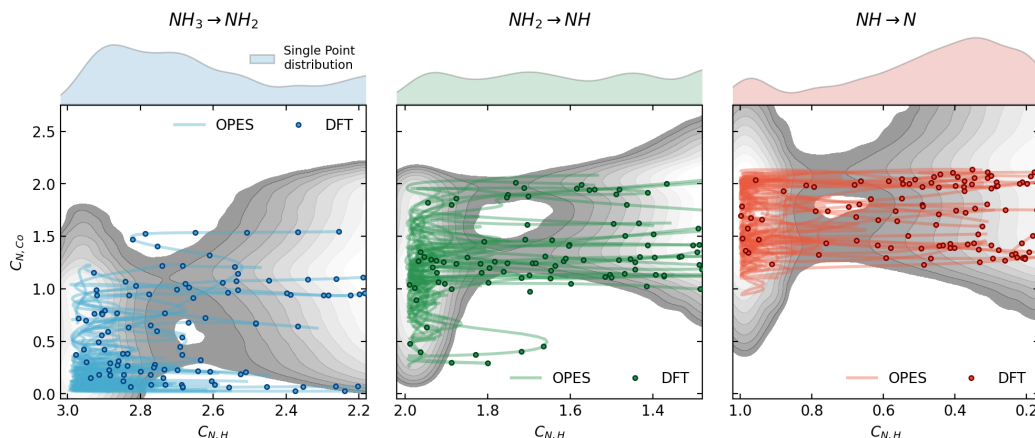


Figure 6.4: **Reaction path discovery for the hydrogenation steps.** The trajectories are visualized in the space defined by $C_{N,H}$ and $C_{N,Co}$. To facilitate the understanding of the sampled paths, we have shown in the background the free energy surface obtained from the final simulations, which are shown up to $\sim 2 k_B T$ above the barriers, highlighting the minimum free energy paths for each reaction (see Fig. 6 for the full colored version). The dots represent the structures recalculated with DFT, whose distributions along $C_{N,H}$ are projected above the upper axis.

simulations, using the same CVs of Ref. [44], namely the N-N distance $d_{N,N}$ and the coordination of N with surface atoms $C_{N,Fe|Co}$. As it can be seen in Fig. 6.3b, the bias potential promoted the exploration of new regions of the potential energy surface. Whenever the uncertainty of the local environments (Fig. 6.3c) exceeded a predefined threshold (*i.e.*, it was structurally different from the environments in the training set), a new DFT calculation was performed, and the GP model was updated. The reactive pathway of N_2 dissociation can be effectively visualized in the plane defined by the two CVs (thick orange curve in Fig. 6.3d), illustrating the progression from the gas phase through the adsorption sites γ and α and finally to the transition state region. Whenever the system reached the product state (vertical dashed line), the simulation was stopped, and a new one started. In this way, an ensemble of reactive pathways was iteratively collected (Fig. 6.3d), ensuring continuous refinement of the potential energy surface and inclusion of a diverse set of reactive environments. Indeed, each pathway exhibited geometric diversity covering a wide range of coordination with surface atoms while originating from the same horizontal adsorption site α .

The need for a flexible and adaptive approach is even more important for modeling the three dehydrogenation steps from ammonia to atomic nitrogen. To harvest such reactive structures, we followed the same protocol, using uncertainty-aware flooding simulations starting from the GP potentials optimized for each reactant. In previous work on the Fe surfaces (see Chapter 3), despite being three steps of the

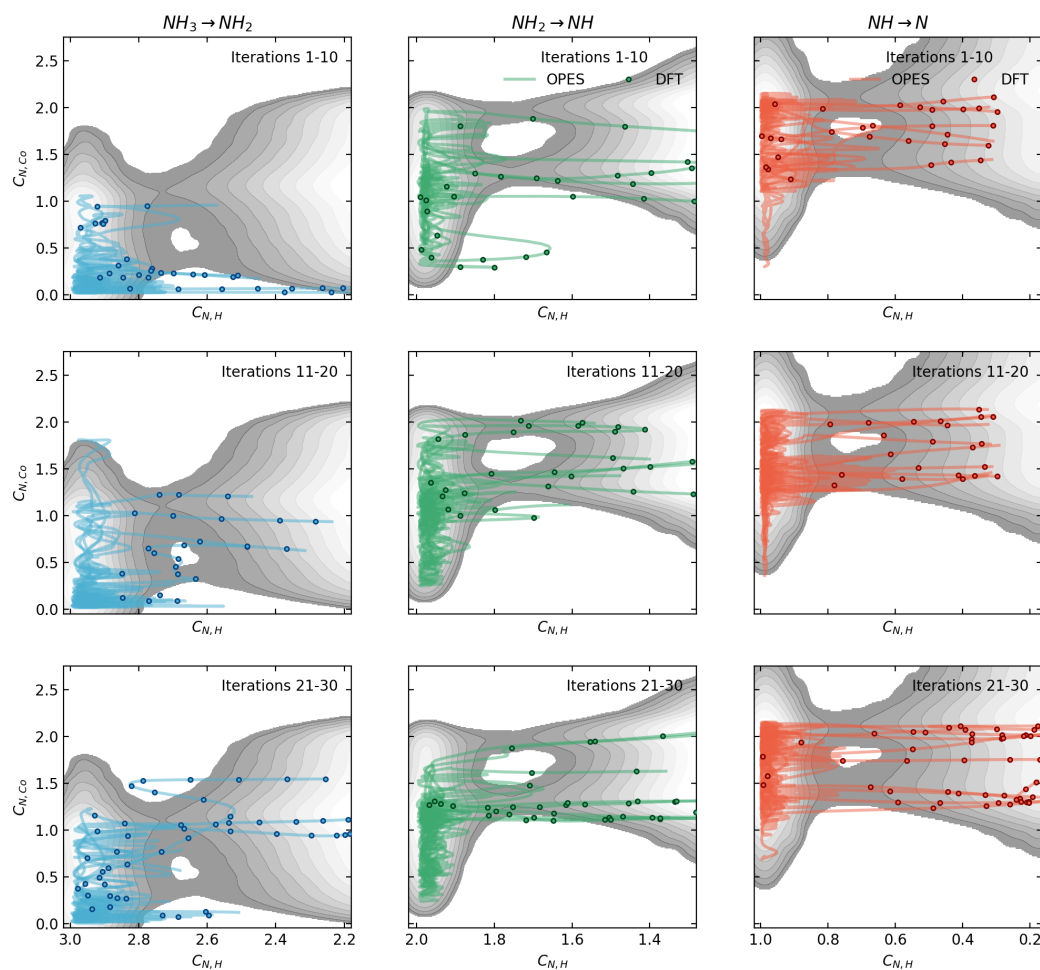


Figure 6.5: **Evolution of the OPES-flooding reactive discovery iterations** for the dehydrogenation reactions. This is the same plot as in Fig. 6.4, divided here into three rows according to the number of iterations. The sampled paths in the first few iterations mostly involve regions of higher free energy. As the flooding iterations proceed and the potential is updated with these new data points, the sampled paths converge to lower free energy regions, emphasizing the importance of an incremental approach. This is particularly evident in the case of the first step of dehydrogenation of NH_3 , which initially occurs at very low coordination numbers between N and Co and then converges to the minimum free energy pathway (that is around 1).

same chemical process (dehydrogenation), the different interactions of the adsorbates with surface atoms had necessitated the use of different CVs to converge the free energy calculations. Since we are in an exploratory phase here, we wanted to use instead a single generic CV (*i.e.*, the planar distance between N and H) for all reaction steps to demonstrate our approach's ability to find the accessible pathways with minimal knowledge. The resulting trajectories are visualized in the plane defined by nitrogen-

hydrogen and nitrogen-cobalt coordination (Fig. 6.4), whose distribution revealed a broad spectrum indicative of the existence of at least two distinct reaction channels. To put our results into context, we superimposed the reactive pathways on the free energy surface derived later from our study (analogous to Fig. 7.1). The remarkable alignment between the sampled and minimum free energy pathways confirmed the effectiveness of our methodology in accurately sampling the crucial reactive configurations without prior assumptions on the adsorption sites or the transition states. Although each dehydrogenation step occurred from different adsorption sites characterized by increasing coordination with the surface, thanks to the flooding scheme, we were able to find reactive pathways using a generic CV that only distinguished reactants from products. Furthermore, an analysis of the evolution of the sampled paths as a function of iterations demonstrated the importance of an incremental approach: the transition pathways initially pass through high free-energy regions and then converge to low free-energy regions as new data are added (Fig. 6.5).

In total, about 400 reactive configurations were collected in this exploratory phase, demonstrating that our approach enables systematic and efficient sampling of reaction pathways under *operando* conditions, even when multiple ones are available.

6.1.3 Uniform accuracy along reactive paths through GNNs and Data-Efficient Active Learning (DEAL)

Having collected the first dataset of reactive structures, we moved on to the second phase of our approach, aiming to converge the accuracy along the reactive paths. For this purpose, we used equivariant graph neural networks (GNNs) [225] to represent the potential energy surface. As we discussed in Section 2.2.3.2, these architectures are more flexible than GPs [226] and have demonstrated remarkable data efficiency, enabling robust simulations as early as hundreds/thousands of training points [227]. In particular, we used MACE [111], which integrates the descriptors used in our previous approach (*i.e.*, ACE) in a message-passing scheme. Resorting to NN allowed us to consolidate all previously collected data into a single potential, overcoming the limitations imposed by having separate GPs for each intermediate. In addition, this allowed us to conduct long (\sim ns) enhanced sampling simulations instead of the many short (\sim ps) simulations of the exploratory phase. This way, we sampled hundreds of different reactive events, providing a more thorough sampling of reaction pathways under dynamic conditions.

To enrich the training dataset further and improve the accuracy of our model, we turned our attention to selecting additional structures from the new MD trajectories. As before, we illustrate our method while focusing on the nitrogen recombination

Data-Efficient Active Learning (DEAL)

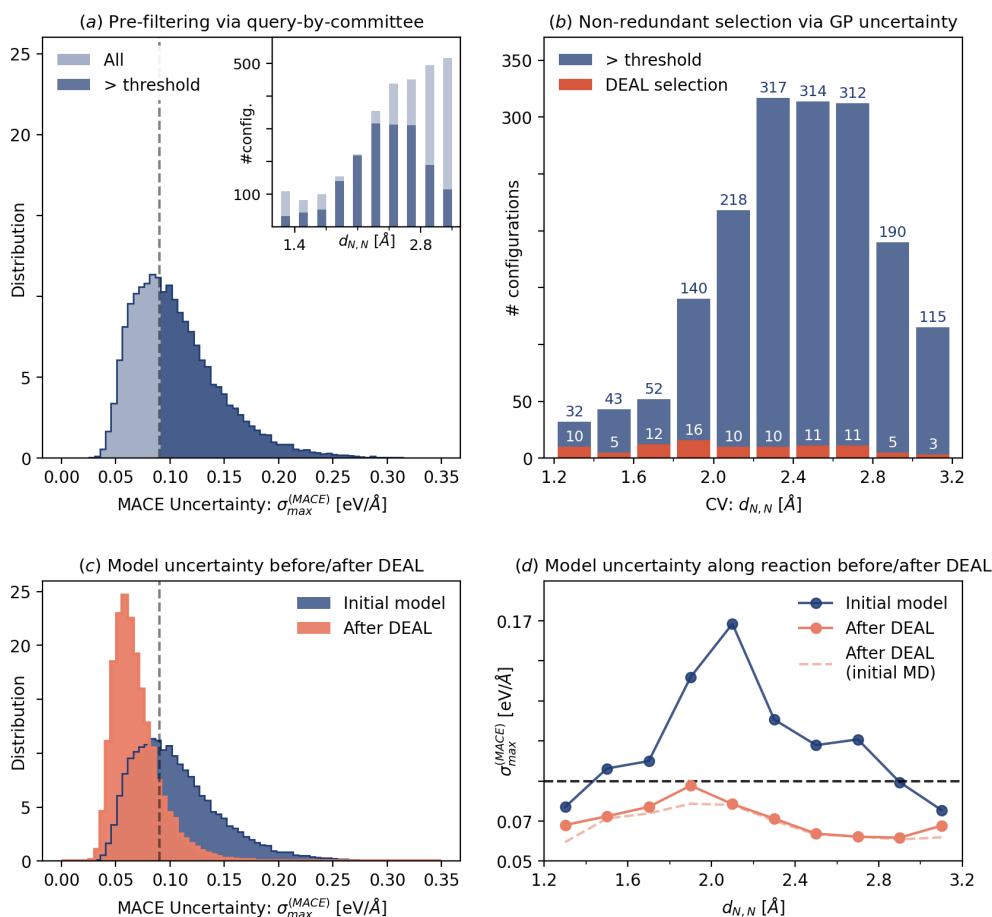


Figure 6.6: **Convergence stage: data-efficient active learning of the GNN.** We illustrate the GP-based active learning of the MACE structures for the N recombination and dissociation process. a) Distribution of the maximum uncertainty on the forces $\sigma_{\max}^{(\text{MACE})}$ (calculated from a committee model) on the simulation performed with MACE before the active learning. The grey dotted line at $90 \text{ meV } \text{\AA}^{-1}$, represents the chosen threshold for the query-by-committee selection. In the inset, we report the distribution of the same configurations along the collective variable $d_{N,N}$. b) Distribution of the DFT single point performed with DEAL (red bars), among the ones pre-selected via query-by-committee (blue bars). The number of configurations in each bin is reported on top of the bars. c) Distribution of the maximum uncertainty on the forces $\sigma_{\max}^{(\text{MACE})}$ for initial model (blue histogram) and after DEAL (red histogram). The second distribution is obtained from a new MD with identical parameters performed with the MACE model after DEAL. d), Average uncertainty along the reaction coordinate $d_{N,N}$ for the two simulations generated respectively with the initial MACE model (blue line) and after DEAL (red line).

and dissociation process. Our first step is to examine the uncertainty of the MACE model, evaluated using the maximum standard deviation of a committee of models (see Equation 2.14). In Figure 6.6a, we report the distribution of the maximum uncertainty per configuration peaks around $80 \text{ meV } \text{\AA}^{-1}$. We note that this value is high due to the large magnitudes of the N forces involved in the reaction: 90 percent of these forces have a relative uncertainty of less than 20 percent, see Fig. 6.7.




Figure 6.7: **MACE relative uncertainty before/after DEAL.** We plotted the distribution of the relative maximum force uncertainty $\sigma_{max}^{(MACE)}/f$ for the initial MACE model (top) and after (bottom) the DEAL the active learning scheme described in Main Text Fig. 4. On the right axis, the corresponding cumulative fraction is reported.

To assess the quality of the potential in describing the reactive event, it's revealing to analyze the average uncertainty as a function of the collective variable, which describes the progress of the reaction (Fig. 6.6d). Indeed, this quantity showed a strongly nonuniform behavior, with significantly higher uncertainty in the region between the reactant and product states (*i.e.*, $1.5 < d_{N,N} < 2.2$). Furthermore, by comparing the mean uncertainty with the distribution of the selected configurations (Fig. 6.6c) we realized that the query-by-committee criterion would lead to an imbalance between the region of the space which is most sampled ($2.2 < d_{N,N} < 2.7$), but whose configurations are only slightly above the threshold, and the one that really needs to be improved. In previous works, it was necessary to manually select multiple

configurations with high uncertainty [63] or from a specific region in CV space [129] to accurately model reactive events.

Our objective is to avoid these issues and systematically achieve uniform accuracy along the entire reaction path with a minimal number of DFT calculations. To this end, we propose an active learning scheme that uses the uncertainty of local environments to select a minimal set of representative configurations. In the following, we refer to this scheme as Data-Efficient Active Learning (DEAL). DEAL consists of two parts: first, a pre-selection of structures with high uncertainty is made using the query-by-committee criterion (see Fig. 6.6a). Next, we extract a representative subset of non-redundant configurations using Gaussian processes. Specifically, we use the GP models trained in the previous stage to decide which configurations to recalculate at the DFT level, by measuring the uncertainty of the local environments of the pre-selected structures. In this way, the uncertainty estimate is updated on the fly after each single-point calculation, allowing us to filter out all the redundant configurations. In the case of the $N_2 \rightleftharpoons 2N$ reaction, this approach allowed us to perform only ~ 100 DFT calculations, which corresponds to about 5% of those selected using the query-by-committee criterion (Fig. 6.6b). Moreover, this percentage varies significantly as a function of CV, ranging from 2-3% in the most sampled region to 12-25% in the TS region, demonstrating how our algorithm can automatically target the relevant configurations without requiring any input on the reaction coordinate.

The addition of the configurations selected by DEAL to the MACE model effectively lowered the uncertainty of the configurations generated in the MD simulations (Fig. 6.6c). In fact, already after only one cycle of active learning, the uncertainty along the reactive coordinate drops below the chosen threshold (Fig. 6.6d). Otherwise, we would have performed another cycle of active learning with the same protocol. We have applied the same scheme to the three dehydrogenation steps, obtaining results similar to those we reported in Figure 6.8.

Before turning to the validation of the potential, we note that this active learning scheme has two limitations: the need for a pre-trained GP model and the seriality of DFT calculations. Therefore, we searched for alternatives that could address these two aspects. To this end, recall that our goal is to extract a representative set of configurations from those reported with high uncertainty. We have found that this objective can also be achieved with a variant of DEAL in which the GP model is trained from scratch using the energy and forces predicted by the potential as labels (see Section 6.3). This is possible because the uncertainty provided by the GP is based on the similarity of the local environments and not on the labels used. In this variant, once the selection of structures is completed, DFT calculations can be performed in an embarrassingly parallel way, as in the standard query-by-committee

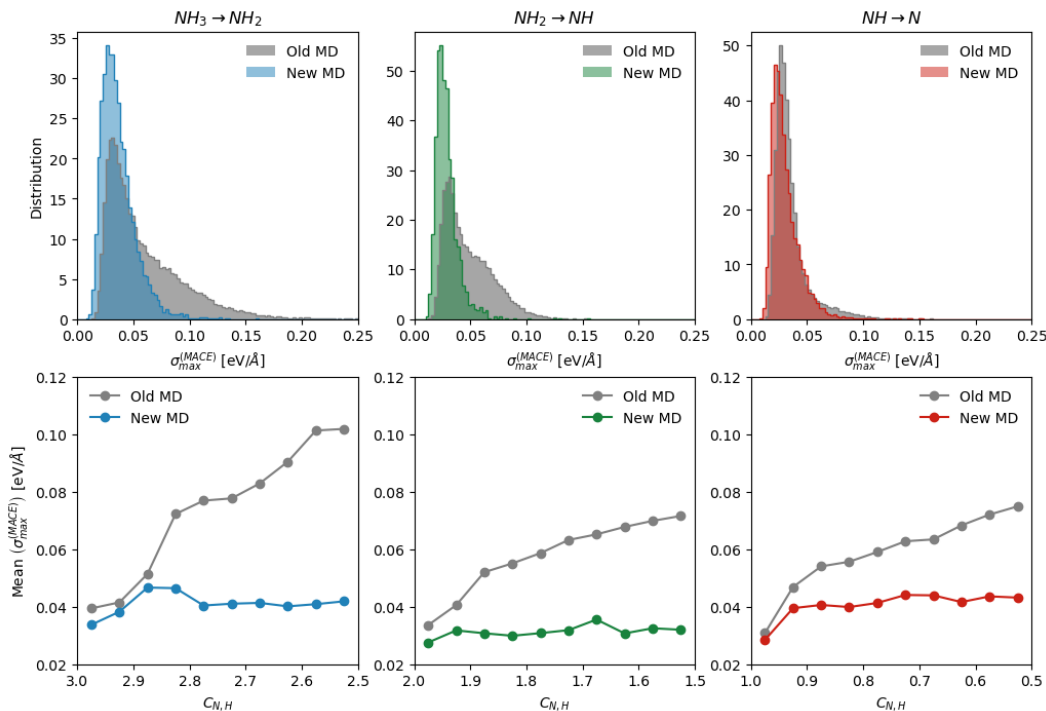


Figure 6.8: **Uncertainty before/after DEAL (dehydrogenation reactions).** We present here the results of the dehydrogenation steps. MD simulations are performed with MACE and the OPES bias to reversibly sample the transition ($NH_x \rightleftharpoons NH_{x-1} + H$, $x=1,2,3$). The first row reports the distribution of the maximum uncertainty on the forces $\sigma_{\max}^{(\text{MACE})}$ on the MD simulations performed with the MACE models before (grey histograms) and after (colored histograms) the active learning procedure. The second row shows the distribution of the average (maximum) uncertainty along the reaction coordinate $C_{N,H}$ for the configurations extracted from the MD simulations before (grey lines) and after (colored lines) the active learning. After the active learning, the uncertainty was uniformly below the threshold ($50 \text{ meV}/\text{\AA}$). Unlike the case of N_2 recombination, the uncertainty in the product regions (low coordination) was even higher than the TS. This was due to the absence of configurations with short-range interaction between the products in the preliminary dataset.

approach. We validated this approach by comparing the distribution of the selected structures with that obtained previously with the pre-trained GP models (Fig. 6.9). In addition, we trained a MACE potential using the newly collected structures and verified that the description along the entire reaction pathway was equally accurate (Fig. 6.10).




Figure 6.9: **DEAL selected configurations by the different variants (pre-train GP vs from scratch)**. Here we compare the distribution of configurations selected by DEAL variants along the CV. (top) Distributions of high uncertainty configurations pre-selected with the query-by-committee criterion. (bottom) Configurations selected by the GP pre-trained on the data collected in steps 0 and 1 and those trained from scratch only on the configurations generated by MD, using MACE energy/forces as labels (see Section 6.3). In particular, since training a model from scratch would result in selecting more configurations from the first part of the dataset, we tested two different schemes: shuffling the trajectory or sorting the configurations according to uncertainty in order to select more configurations with high uncertainty. In both cases, we found remarkable agreement with the selection made with the pre-trained model.




Figure 6.10: **DEAL selection: force accuracy along the reactive pathway for the two variants** (pretrain GP vs from scratch). We compared the accuracy of the force predictions of the N atoms (as in Fig. 6.11) along the reactive pathway using DEAL-selected configurations with a pre-trained or trained ex novo GP model. In both cases, a uniformly good description is obtained along the entire reactive path.

6.2 Validation of the workflow

In this section, we assess the quality of the final potential, as well as investigate its evolution through the different stages. The root mean square error (RMSE) of the final MACE model on the validation dataset was 0.4 meV/atom for energies and 17 meV \AA^{-1} for forces (corresponding to a relative error of less than 3 percent). Figure 6.13 reports the accuracy of force predictions divided per chemical species, from which we see that the adsorbate atoms (N,H) have the largest error, also partly due to the higher magnitude of the forces.

To confirm the accuracy and reliability of our potential in describing reactive processes and to better understand the importance of each stage, we conducted an extensive *post-hoc* analysis. To this end, we trained three different MACE models using the configurations collected at the end of each of the three stages (stage 2 thus represents the final potential). We constructed an additional test set composed of configurations extracted from the production simulations and recalculated at the DFT level. Specifically, we selected 200 structures for the $N_2 \rightleftharpoons 2N$ reaction and 150 for each dehydrogenation step, uniformly distributed along the CVs. We then evaluated how the accuracy of the predictions of the energies and forces (of the adsorbates) changes when expanding the training set as the workflow progresses. This analysis shows that our procedure is able to achieve uniform accuracy along the reaction paths for both energies (Fig. 6.14) and forces (Fig. 6.11 and Fig. 6.12). It also clarifies the

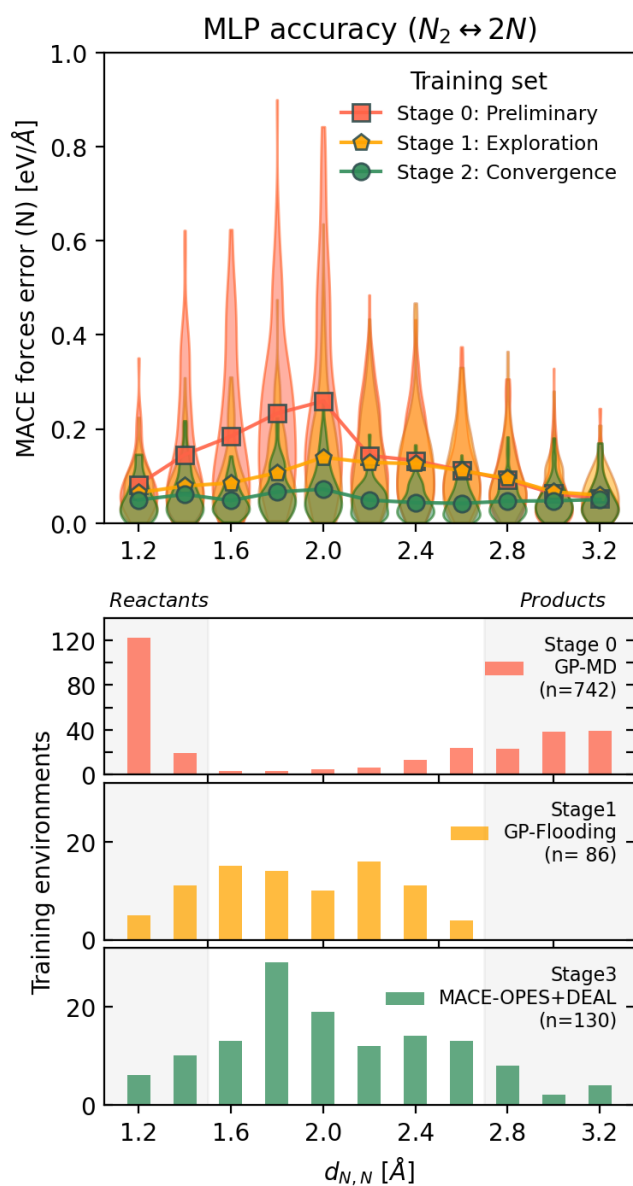


Figure 6.11: **Evolution of force accuracy along the reactive pathways** in the different stages of the workflow, for the case of $N_2 \rightleftharpoons 2N$. (top) Force accuracy for the N forces along the CV (N - N distance), with the violin plot denoting the distribution within each bin and the solid lines representing the average along the CV. Each line/violin plot corresponds to the accuracy of a MACE model optimized using as a training set the configurations collected until a given stage (bottom). Number of configurations in the reactive range collected at each stage. For each stage, it is specified the type of simulation, while the number and length of the simulations are reported in Section 2.7.4.5. The results for the other dehydrogenation steps are reported in Fig. 6.12

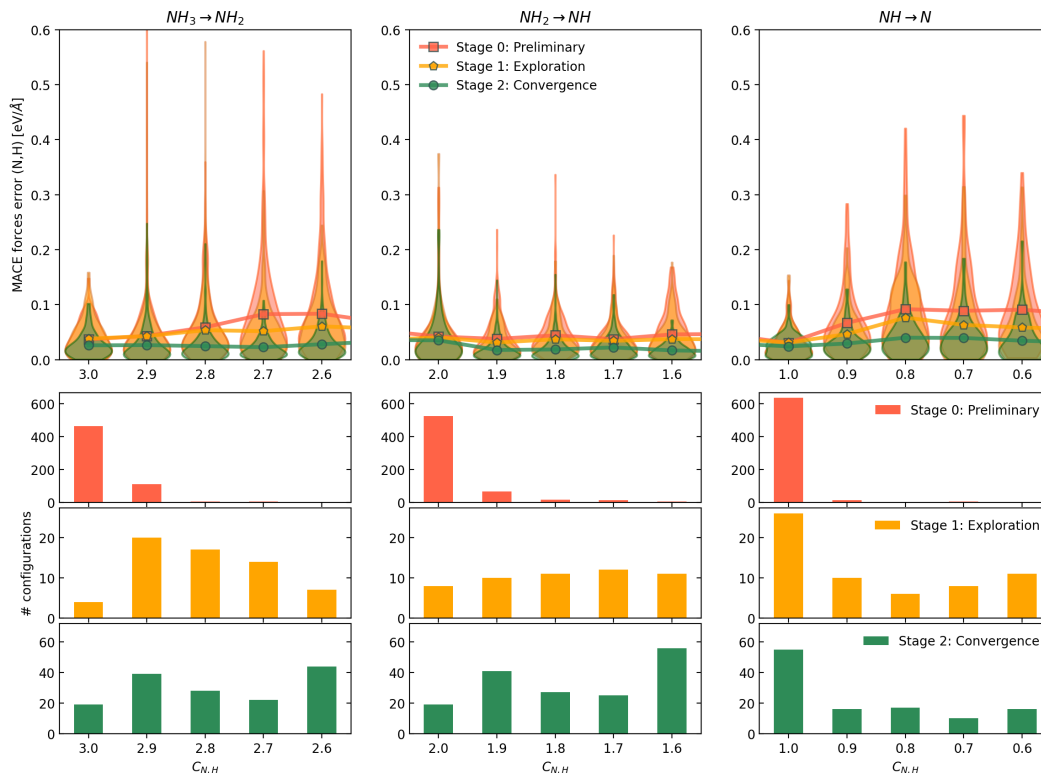


Figure 6.12: *MACE* force accuracy along the reactive pathways of $\text{NH}_x \rightleftharpoons \text{NH}_{x-1} + \text{H}$, $x=1,2,3$. Analogous of Fig. 6.11 for the dehydrogenation steps.

relevance of each phase. In fact, focusing on the force predictions (Fig. 6.11, top panel), we can observe that, at the end of the preliminary step, a significant fraction of the configurations have errors greater than $0.5 \text{ eV } \text{\AA}^{-1}$. Therefore, using directly the MACE potential to perform reactions by skipping the exploratory phase could be dangerous, as we risk extrapolating to unphysical configurations.

We also observed how the relationship between the uncertainty $\sigma_{\text{max}}^{(\text{MACE})}$ and the error on the forces changes in the different stages (Fig. 6.15). Interestingly, in the initial stage, the committee-based uncertainty strongly underestimates the error in the transition region, while this gap decreases as more points are added. Additionally, we also investigated the possibility of using a standard feed-forward NN potential such as DeepPMD [107, 109] instead of MACE. However, Fig. 6.16 shows that, after the exploratory phase, the DeepPMD errors are still large up to more than $1 \text{ eV } \text{\AA}^{-1}$, which leads the simulations to be unstable. This implies that DeepPMD requires larger training datasets to obtain stable simulations, in agreement with what was observed in Ref. [227], and confirms the importance of using data-efficient architectures in our scheme. At the same time, this analysis shows that using the dataset of configura-

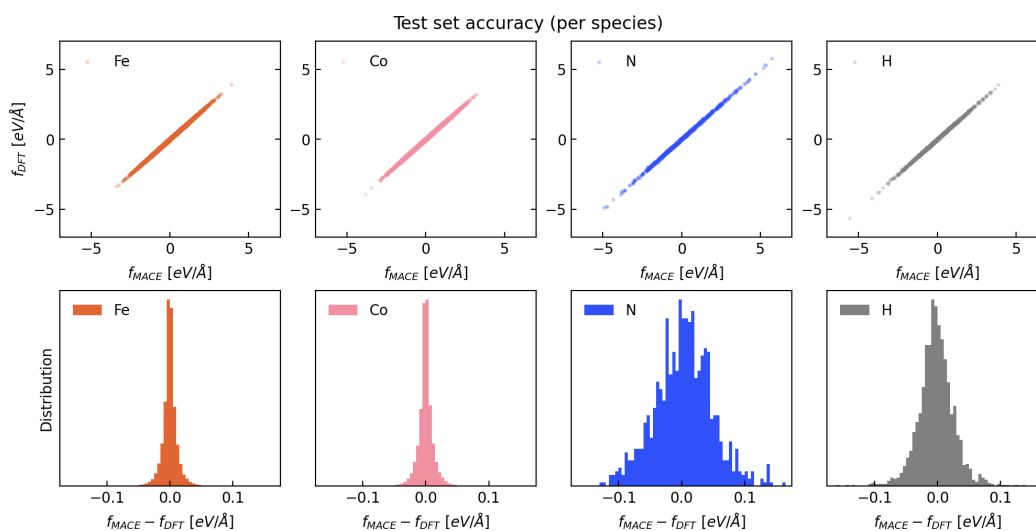


Figure 6.13: **Forces accuracy of the final MACE model divided per chemical species** (top row) MACE forces prediction versus DFT reference, resolved per atomic species. (bottom row) Distribution of the error per each species. The test dataset is composed of a subset of configurations extracted from the final dehydrogenation simulations and recalculated with DFT single-point calculations.




Figure 6.14: **MACE energy accuracy along the reactive pathway** for $N_2 \leftrightarrow 2N$, as a function of the N-N distance. The violin plots represent the distribution of the errors, and the lines denote the average over the points belonging to each bin.

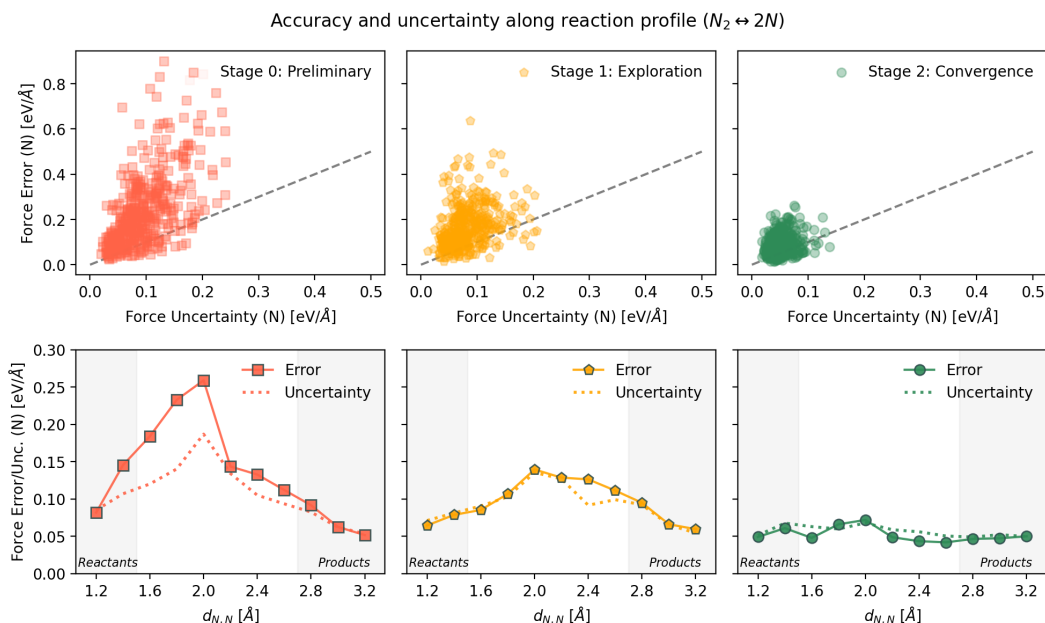


Figure 6.15: **Accuracy vs uncertainty for the MACE GNN models.** Top row: scatter plot between the MACE uncertainty and the MACE error for the predictions of the N forces, calculated on the reactive pathway of the $N_2 \rightleftharpoons 2N$ reaction. Bottom row: Average MACE error/uncertainty along the CV (N - N distance), highlighting the reactants and product regions. Each column corresponds to the MACE models trained with the training set collected at the different stages.

tions collected with our workflow with other architectures still yields uniform (though higher) accuracy along the entire reactive path, further confirming the quality of the data selection scheme. Moreover, this suggests that the use of architecture which are based on the same descriptors is not a strict requirement, leading to a potential substitution of the GNN with an equivalent or even better highly data-efficient architecture.

6.2.1 Data-efficiency

Finally, before concluding, it's worthwhile to comment on the number and distribution of the reference training structure. The number of configurations collected in each stage is reported in Table 6.1, while the distribution along the CVs is reported in the bottom panels of Figs. 6.11 and 6.12. It is worth noting that these distributions are complementary and collectively result in uniform coverage of the reactive path. Remarkably, this uniform distribution was achieved without using the CVs in the selection process, but simply by providing our active learning schemes with reactive trajectories. If we focus on nitrogen dissociation/recombination, our approach

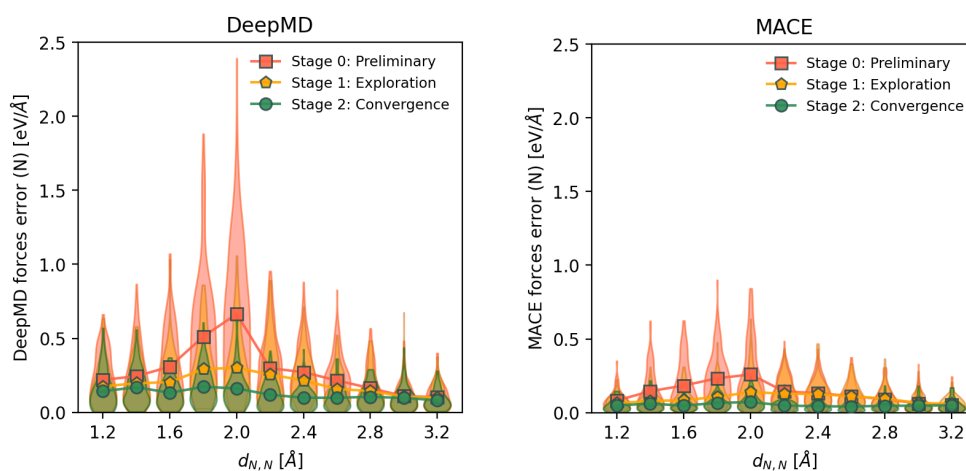


Figure 6.16: **Comparison between DeepMD and MACE for the forces accuracy at the different stages** for the $N_2 \rightleftharpoons 2N$ reaction. To evaluate how important it is to use a data-efficient GNN architecture in our scheme, we optimized another set of ML potentials using the DeepMD scheme on the same data collected at different stages. We notice that the errors are systematically larger: after the initial exploratory phase (Stage 1) DeepMD errors still reach more than 1 eV/Å, which leads the simulations to be highly unstable. As a result, it's not possible to use a potential with DeepMD with so few reactive configurations. At the same time, this analysis shows us that, at the end of the sampling process (Stage 2), the accuracy of DeepMD is also uniform (although much higher) throughout the entire reactive path, confirming the goodness of the data selection scheme.

produces a robust potential with less than 1k configurations (958). Considering the entire catalytic cycle, which consists of 6 different intermediates, a total of 5k DFT calculations were used. Compared to the number of configurations used in previous studies on Fe surfaces (30k for N_2 dissociation [44] and 110k for ammonia decomposition, see Table 2.1), we obtained more than 20-fold improvement in data efficiency. This was possible not only because of data-efficient architectures but, more importantly, due to an efficient protocol for sampling reactive configurations and selecting the most relevant ones.

Dataset composition

System	Stage	MD	System	# config.
FeCo	Preliminary	AIMD ¹	T=700 K	124
		FLARE	T=800 K	6
				130
N_2 / FeCo	Preliminary	FLARE	MD	651
		FLARE	OPES (adsorption/diffusion)	91
	Explore	FLARE	OPES-flooding	86
	Convergence	MACE	OPES	130
			958	
NH_x / FeCo	Preliminary	FLARE	MD	1184
		FLARE	OPES (adsorption/diffusion)	286
	Explore	FLARE	OPES-flooding	323
	Convergence	MACE	OPES (reactive)	450
		MACE	OPES (products) ²	1681
			3924	
TOTAL				5012

Table 6.1: **Dataset composition.** Summary of the configurations collected during the different stages, divided based on the system, the MD driver, and the type of sampling.

¹: Short (400 fs) AIMD simulation, with configurations extracted every 3 steps. These configurations were used to initialize the GP model.

²: Configurations from the product states extracted from the MACE simulations using the query-by-committee criterion.

6.3 DEAL: general features and applicability

In this final section, we want to summarize in a schematic and general way the elements of DEAL, which make it a scheme applicable far beyond the specific case

presented earlier. DEAL is an active learning scheme that aims to improve the accuracy of an ML potential in a data-efficient way. Given a set of configurations (typically obtained from MD simulations using the potential itself), we want to identify a non-redundant subset to be recomputed at the DFT level and added to the training set of the potential. DEAL is composed of two steps: I) pre-filtering via query-by-committee and II) non-redundant selection of the local environments using the uncertainty provided by GPs.

Pre-filtering via query-by-committee. In the first step, we select a set of configurations that are not properly described by the ML potential, as measured by the uncertainty computed by a committee of NNs. That is, we select all the configurations whose uncertainty is above a given threshold. We suggest choosing the threshold not based on the absolute value of uncertainty but in such a way as to select the subset of configurations that are described worse, as routinely done for the standard query-by-committee approach, see Section 2.3.1. In fact, especially in the early stages, the uncertainty could significantly underestimate the error (see Fig. 6.15), just as the absolute value depends greatly on the architecture of the NN (see Fig. 6.16). Also, it is appropriate to set a limit on the maximum uncertainty (σ_u , see Section 2.3.1) to avoid including non-physical configurations, especially in the first stages.

Non-redundant selection via GP uncertainty. The second step is the identification of a non-redundant subset from those with high uncertainty. To achieve this goal, we train a GP model on the filtered configurations, using the maximum uncertainty on the local environments $\sigma_{\max}^{(\text{GP})}$, see 2.2.4. This quantity correlates quite well with $\sigma_{\max}^{(\text{MACE})}$ (Fig. 6.17). This allows the uncertainty estimate to be updated on the fly and, in turn, selects a non-redundant set of configurations to be calculated at the DFT level. In practice, this can be accomplished in two different ways: (a) using the pre-trained GP models on the previously collected configurations and performing a DFT calculation before updating the GP or (b) training a GP model from scratch only on the filtered configurations and using the MD energy/forces as labels (selection only). This is possible because the uncertainty provided by the GP is based on the similarity of the local environments and does not depend on the actual labels used for training. We note that the second strategy has several practical advantages. First, the training time for the GP is greatly reduced because the size of the training set remains small. More importantly, DFT calculations can be performed in parallel after selection. The only downside is the fact that it might select structures that are already in the ML potential’s training set, but the query-by-committee pre-selection ensures that all these structures are not described well by the ML potential. Another




Figure 6.17: **Correlation between GP and GNN uncertainties.** We compared the GP maximum local atomic uncertainty $\sigma_{\max}^{(GP)}$ versus the MACE maximum force standard deviation $\sigma_{\max}^{(MACE)}$, for the $N_2 \rightleftharpoons 2N$ reaction step, for both the adsorbate molecule/atoms (N) and the surface (Fe|Co).

element to consider is that, when building a sparse GP model from scratch, more of the first parsed structures will be selected. For this reason, we tested two different schemes: shuffling the trajectory or sorting the configurations by their uncertainty (Fig. 6.9).

Finally, we note that while this scheme was designed to achieve uniform accuracy along reactive paths (by applying it to enhanced sampling simulations), it does not require any knowledge of the reaction coordinate and is thus applicable as a general active learning scheme also to non-reactive simulations.

We have successfully applied DEAL using the GP from scratch approach to study the impact of finite surface coverage on the reactions. As we will demonstrate in Chapter 7, in 5 iterations and the addition of only ~ 3000 single point calculations, we manage to accurately describe the two fundamental processes of ammonia decomposition, namely nitrogen recombination and nitrogen dissolution/segregation.

Discussion

The training set construction typically involves active learning procedures [117, 119, 120]. In particular, it is corroborated that, by integrating these sampling techniques with active learning strategies, it becomes possible to construct MLPs reliable for simulating a wide variety of rare events [44, 63, 64, 125–137], in particular catalytic processes [44, 63, 64] as we did in the case of Fe.

However, this strategy is effective but not efficient and cannot be applied to more computationally expensive systems. Thus, we focus on the development of a

method that requires as few QM calculations as possible, which are essential to enable the widespread use of ML-based potentials. In this Chapter, we introduced a data-efficient scheme for constructing reactive potentials. We demonstrate that this method can effectively capture complex reaction mechanisms composed of multiple channels without prior knowledge of the adsorption sites or the transition states. Remarkably, the number of needed training points is at least one order of magnitude less than the ones required for the iron catalyst. These advances can facilitate the simulation of more realistic systems, both in complexity and underlying level of theory. Also, it is important to note that the proposed approach is highly modular. For example, if one has already collected reactive configurations, it is possible to move directly to the second stage of active learning with DEAL. Or, one could leverage foundational models instead of GPs to collect the initial set of reactive configurations. In addition, it will be interesting to test the parallel variant of DEAL as a tool to reduce the size of existing databases, as we are going to do in the case of Fe ones. These features make our approach a general and versatile framework for constructing reactive potentials, applicable not only to catalytic processes but also to the broader field of chemical reactions and materials. Indeed, we believe that having a data-efficient way to construct ML potentials is critical to enable the study of increasingly complex systems, both in terms of the environments (such as reactions in the presence of solvent or surface dynamics) and the level of theory of the reference calculations. Finally, this workflow could be further integrated with ML techniques to design collective variables[228] in a data-driven way, thereby reducing the need to identify process-specific CVs. Similarly, path sampling methods could be integrated into the active learning scheme to investigate the mechanism in a more unbiased way [130]. For all these reasons, we believe that our scheme can be effectively used to construct reactive potentials for a wide variety of processes occurring under dynamic conditions, gaining insights into their workings at the atomistic level. We are currently applying this novel strategy to study many complex systems at an affordable computational and time cost, from the nitride iron phases to pre-melted catalysts. These studies are still ongoing and not included in this Thesis, but we hope to include them in future publications.

Chapter 7

Ammonia cracking on iron-cobalt

In this last chapter, we present the study of ammonia decomposition on iron-cobalt (FeCo) alloy catalysts, which has been suggested as a promising alternative to pure iron due to its superior catalytic performance (see Section 1.2.2). From the modeling perspective, even if this is a new system that has never been studied, it has the same crystal structure of α -Fe (BCC). Thus, we can leverage the knowledge obtained from the previous chapters. As we have described in Chapter 6, the study of ammonia decomposition on FeCo was made possible by the development of the protocol for the construction of an MLP that accurately describes chemical reactions in a data-efficient manner. In the following, all the results will be obtained with this MLP; we refer the reader to Section 2.7.4 for the computational details.

The first sections will be on unraveling the free energy landscape and the elementary steps of the reaction, as well as on investigating the effects of nitrogen migration inside the bulk of the alloy; comparing the results with the one obtained on pure iron. Then, we will show how we effectively expanded the MLP to study lateral interaction and the role of finite surface coverage on the two more relevant processes, which we have shown to explain the nitride formation in Fe catalysts (Chapter 4), providing a microscopic interpretation on the high catalytic activity of FeCo.

7.1 Unraveling the free energy landscape and reaction mechanism

Following the approach adopted for the Fe catalyst (Chapters 3 and 4), we study here the elementary steps of ammonia decomposition on the (1 1 0) surface of FeCo. XRD measurements indicate that the active FeCo samples are those presenting a body-centered cubic (BCC) alloy arrangement, with the (1 1 0) surface predominantly exposed [30, 51]. On this surface, iron and cobalt atoms are exposed in a 1:1 ratio,

as shown in Fig. 6.2.

Before discussing the reactions, all simulated at the operating temperature of 700 K as for Fe, it is important to recall the behavior of the catalyst at this temperature. As with the case of Fe, this close-packed surface is particularly stable. In fact, even though we are above the Hüttig temperature (~ 600 K), the energy required for defect formation is greater than the entropy gain obtained from it.

7.1.1 Dehydrogenation steps

We start focusing on ammonia dehydrogenation steps, which eventually lead to atomic nitrogen and hydrogen release.

Similarly to Fe, we first performed a set of long OPES simulations, which allowed us to calculate the free energy profiles of the three dehydrogenation steps (as depicted in Fig. 7.1). The resulting free energy barriers exhibited varying heights, with $\text{NH}_2 \rightarrow \text{NH} + \text{H}$ being the lowest, followed by $\text{NH}_3 \rightarrow \text{NH}_2 + \text{H}$, and finally $\text{NH} \rightarrow \text{N} + \text{H}$, similarly to what was observed for Fe(110). We have reported the tabulated results for each elementary step in Table 7.1. It's noteworthy to notice that the free energies

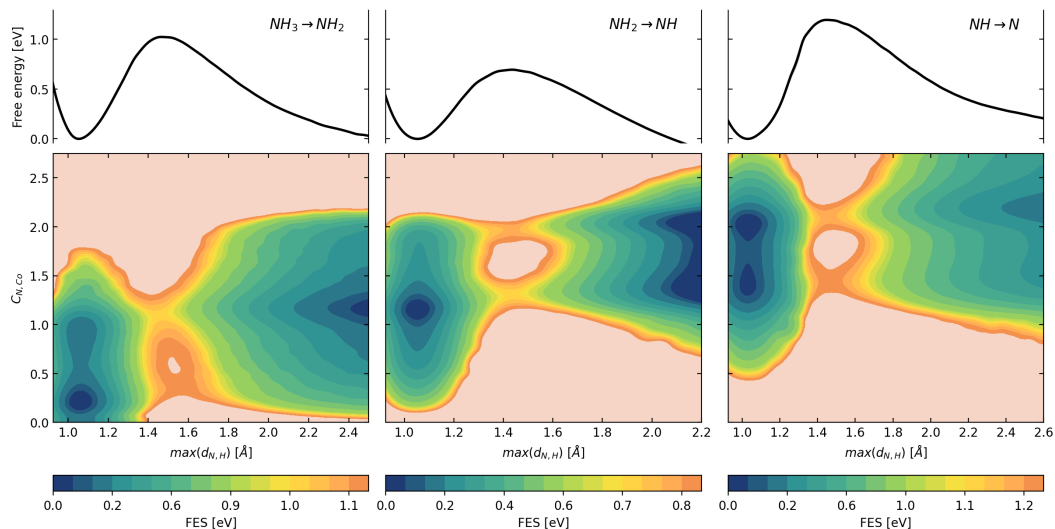


Figure 7.1: **Free energy profile of dehydrogenation steps** projected along the maximum distance between N and H atoms $\max(d_{N,H})$ (top), and on the 2D plane defined by $\max(d_{N,H})$ and $C_{N,Co}$ (bottom). The uncertainties on the 1D free energy profiles (shadow region, ~ 0.01 eV) are calculated with a weighted block average [157].

are accurately reconstructed, with a sampling error smaller than 20 meV, demonstrating the goodness of the developed protocol and the resulting MLP. Given the presence of two different metal species at the surface, it is interesting to perform a deeper analysis to ascribe the potential role played by Fe and Co.




Figure 7.2: **Free energy surfaces of dehydrogenation steps** $NH_x \rightleftharpoons NH_{x-1} + H$ ($x=1,2,3$) projected along the maximum distance between N and H atoms $\max(d_{N,H})$ and the coordination with surface atoms ($C_{H,Fe|Co}$, top panels), cobalt ($C_{H,Co}$, central panels), and iron ($C_{H,Fe}$, bottom panels). The free energy surfaces are shown up to $3k_B T$ above the barriers for each reaction. In terms of coordination with the surface, the minimum free energy pathways pass through a value of $C_{H,Fe|Co} \approx 1$ (top panels). However, if we resolve the coordination with respect to Fe and Co, we see a marked tendency for the reactive hydrogen to form bonds with Co rather than Fe, indicated by the minimum free energy path passing through $C_{H,Co} \approx 1$ and $C_{H,Fe} \approx 0$.

Analyzing the 2D free energy surfaces projected onto the maximum distance between N and H atoms, $\max(d_{N,H})$, and the coordination number between N and Co atoms, $C_{N,Co}$, reveals further insights on the reaction mechanism. Multiple pathways are clearly present in all three reaction steps, as already observed in the exploration phase. Notably, these paths exhibit very similar barrier heights but distinct geometric configurations, underscoring the necessity of considering all possible routes. In the case of NH_3 dehydrogenation, we observed two predominant paths characterized by the nitrogen-cobalt coordinations being around 0.5 or 1, with a preference for the latter. Similarly, for NH_2 , the reaction occurred at either coordination 1 or 2, albeit with a slight favor towards the former. Finally, for NH , there was substantial parity between the two paths at coordination values 1 and 2. If we look instead from the perspective of hydrogen, its total coordination with the surface in the TS region is always around 1. Still, the reaction occurs preferentially when the reactive hydrogen is in contact with cobalt atoms (Figure 7.2).

To gain a deeper understanding of the reaction paths, we conducted a new set of flooding simulations using the same scheme that we have adopted for iron, namely enhancing only the reactant region, guaranteeing that no bias is added in the TS and product regions (see Section 2.6.1). We thus obtain an ensemble of 100 reactive trajectories per step, which are reported in Figure 7.3. From the distributions reported on the right of each panel, we can learn a qualitatively different behavior between NH_3 and NH_2/NH . Notably, NH_3 trajectories exhibited a broadened distribution in the coordination between H and Co, $C_{H,Co}$, reacting uniformly in contact with either iron ($C_{H,Co} = 0$) and cobalt ($C_{H,Co} = 1$). Conversely, NH_2/NH trajectories displayed a greater flux over higher coordination values, indicating a propensity for dehydrogenation when hydrogen is cobalt-coordinated.

To focus more on the transition state (TS) region, we report in Figure 7.4 a detailed analysis of TS configurations in terms of nitrogen and hydrogen coordinations. As we observed for iron (see Figure 3.9), the TS configurations form an ensemble of structures rather than a well-defined TS. Projecting the results in terms of these variables, it is evident that nitrogen has a greater affinity for forming bonds with iron. In fact, $C_{N,Fe} > C_{N,Co}$ for the majority of the configurations collected (76, 82, and 90% for the three dehydrogenation steps, respectively). When looking at the total coordination of NH_x at the TS (considering both the coordination with metal atoms and with hydrogen), it is always greater than 4, indicating that the reaction occurs in an over-coordination regime with the metal, similarly to iron (Figure 3.10). On the other hand, this analysis also highlights what was mentioned earlier, namely the higher propensity of hydrogen to bond with cobalt rather than with iron ($C_{H,Co} > C_{H,Fe}$ in 55, 69, and 70% of the reactive events for the three

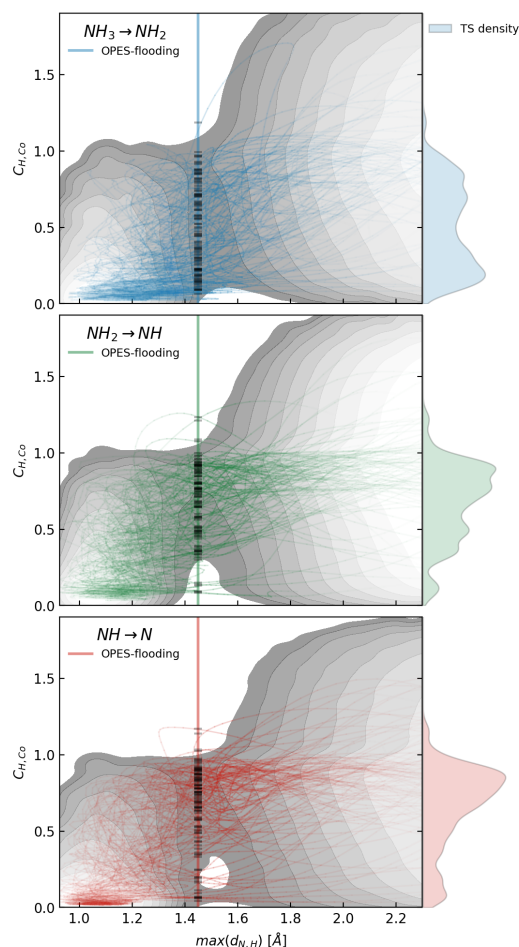


Figure 7.3: **Reactive path analysis for the dehydrogenation steps.** An ensemble of 100 trajectories is calculated with the OPES-flooding scheme and visualized in the 2D plane defined by $\max(d_{N,H})$ and $C_{H,Co}$. Only the 40 fs around the reactive event (determined by $\max(d_{N,H})$ crossing 1.45 \AA , denoted by a vertical line) are reported for each trajectory. For convenience, we report in the background the free energy surfaces projected on the same coordinates obtained from the OPES simulations (Fig. 7.2), which are shown up to $\sim 2k_B T$ above the barriers for each reaction. On the right axis, we projected the distribution of the coordination number $C_{H,Co}$ when crossing the line at $\max(d_{N,H}) = 1.45 \text{ \AA}$ (black dashes).

dehydrogenation steps, respectively). The broader distribution exhibited by the first dehydrogenation step can be explained by the fact that, as discussed in Section 3.2.1, the nitrogen molecule can be considered physisorbed on the surface and, therefore, quite mobile. Additionally, a large portion of reactive events occurs due to the high kinetic energy the molecule gains during adsorption, as also observed by Gerrits *et al.* on Ru(0001) [59].

This comprehensive analysis underscores the multiplicity of reaction pathways

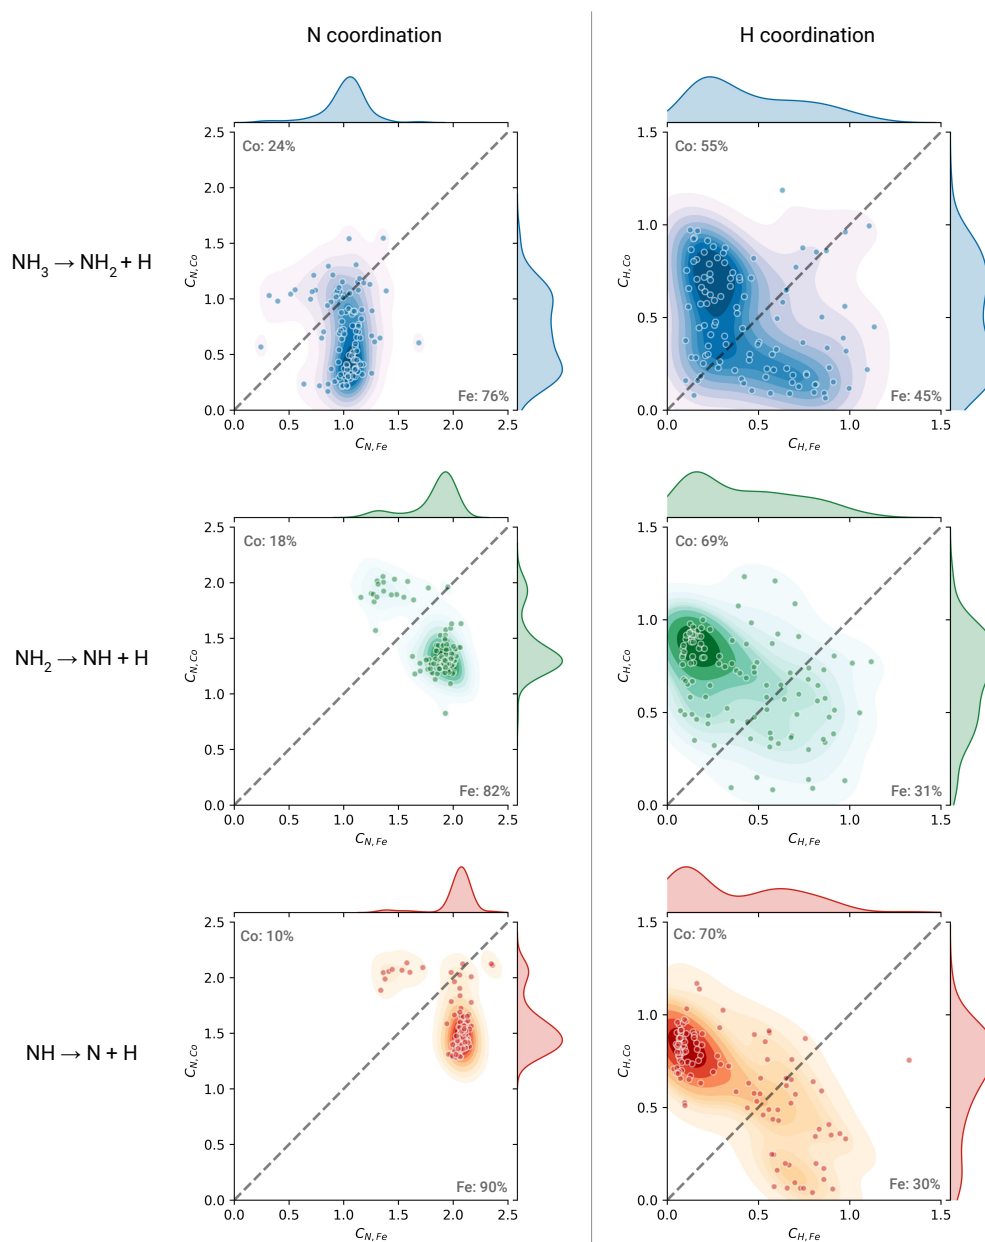


Figure 7.4: **Transition state analysis.** Analysis of geometries related to the transition state configurations extracted from the OPES-flooding trajectories (see Fig. 7.3) for the three dehydrogenation steps. The left column shows the coordination number between iron (x) and cobalt (y) with N, while the right column shows the one with H. The distribution shows that, at the transition state, N is more coordinated with Fe in most of the reactive events, while H is more closely coordinated with Co. This can be quantified by the percentage of reactive events for which the coordination with N/H is higher with Co (Fe), which is reported in the top left (bottom right) corner of each plot.

facilitated by high-temperature dynamic conditions. It also demonstrates the effectiveness of the data-efficient approach presented in Chapter 6, not only in parsimoniously constructing the potential but also in investigating the process under operative conditions with the same simulations, thus elucidating the reaction mechanism and reconstructing the free energy profile of the different catalytic steps.

7.1.2 Nitrogen adsorption and recombination

We now focus on studying the behavior of atomic nitrogen released on the surface (N^*). Here, we repeat the studies conducted for Fe (see Section 4.1) to understand its adsorption and to study its surface diffusion.

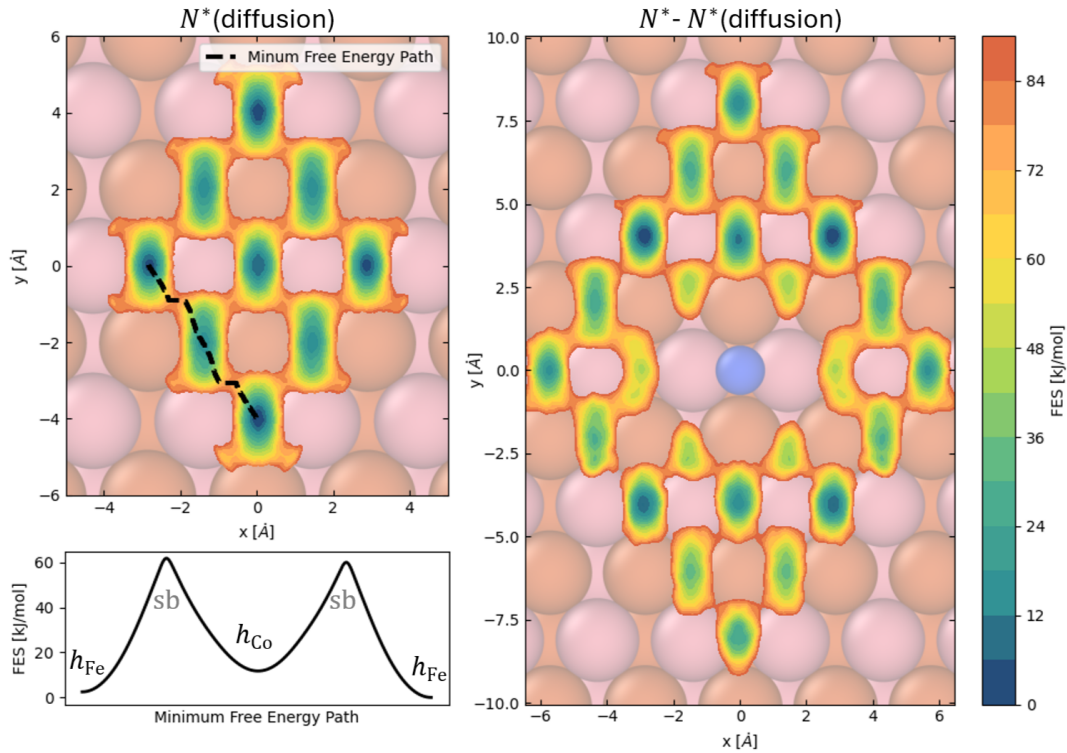



Figure 7.5: On the left panel, free energy of N^* projected along the two crystallographic directions $[001]$ (x -axis) and $[1\bar{1}0]$ (y -axis). Local minima represent metastable states, and black dotted lines denote the minimum free energy diffusion pathways. On the bottom panel, free energy is projected along the minimum free energy pathways. On the right, the same free energy surface, in the presence of another atomic nitrogen adsorbed in the center. For an effective sampling, a reflective wall was placed at $|v_{hkl}| < 2.55$, $hkl = 1\bar{1}1, 1\bar{1}\bar{1}$, at 1.55 in the absence of the second N^* atom. Similarly, to avoid the diffusion of the central nitrogen, another wall at $|v_{hkl}| < 0.55$, $hkl = 1\bar{1}1, 1\bar{1}\bar{1}$.

In the top left panel of Figure 7.5, we present the 2D free energy surface (FES)

projected along the two in-plane crystallographic directions $[001]$ (x) and $[1\bar{1}0]$ (y). We can immediately observe that nitrogen is preferentially adsorbed at the hollow sites (h). However, the FeCo(110) surface features two distinct hollow sites, h_{Fe} and h_{Co} , due to the asymmetry caused by the two different chemical species. In the h_{Fe} site, nitrogen is coordinated with two cobalt atoms (in the first layer along the x direction) and three iron atoms (two in the first layer along the y direction and one in the layer below). Similarly, in the h_{Co} site, the coordination is the same but with iron and cobalt swapped. From the FES, we can see that h_{Fe} is slightly favored over h_{Co} , which has a free energy that is approximately 15 kJ/mol higher.

From the FES, we can also analyze the diffusion mechanism and the associated free energy barrier. The minimum free energy path (black dotted line) for diffusion between adjacent sites passes through a short bridge (sb); in this case, all sb sites are equivalent. Specifically, to move from one h_{Fe} site to the next, nitrogen passes through an h_{Co} site, crossing two sb sites:



with free energy barriers (ΔG^\ddagger) of approximately 65 and 50 kJ/mol, respectively (Figure 7.5, bottom left panel).

Furthermore, we repeat the instructive study of nitrogen diffusion in the presence of another adsorbed N^* at the center (Figure 7.5, right panel). This study is again made possible by employing the variables $v_{1\bar{1}1}$ and $v_{1\bar{1}\bar{1}}$ (see Equations 2.30 and Section 2.7.4.6).

As can be observed, the resulting FES is practically superimposable with that obtained for Fe, with the four hollow sites adjacent to the occupied central site being approximately 40 kJ/mol higher in energy than the most stable sites. Once again, these sites are of considerable interest because they represent the precursor states of recombination. It should be noted that in the specific case depicted in the figure, the central nitrogen is placed in h_{Fe} , and the four adjacent sites are h_{Co} . However, the result remains the same if the central atom is placed in h_{Co} , as the presence of two nitrogen atoms in two adjacent h sites resolves the asymmetry of the problem.

Having defined and characterized the precursor site, we can now focus on the recombination of two N^* starting from this site. In Figure 7.6, we report the FES as a function of the distance between two nitrogen atoms $d(N, N)$ and the coordination between nitrogen and the metal atoms $C_{N, \text{Fe|Co}}$. It is important to note that, unlike what was shown in Section 6.1.2, here we are focusing on the recombination $2N^* \rightarrow N_2$. The reactant state consists of two N^* atoms adsorbed at two adjacent h sites,

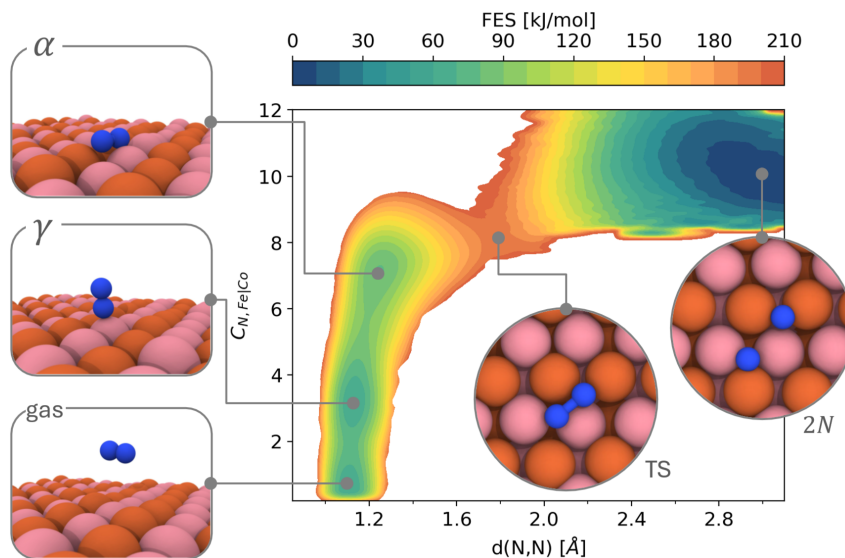


Figure 7.6: Free energy surface of nitrogen recombination/dissociation ($2N^* \rightleftharpoons N_2$) as a function of N-N distance $d(N,N)$ and the coordination between nitrogen and the metal atoms $C_{N,Fe|Co}$ at 700 K. The local minima on the plots correspond to metastable states. Some snapshots along the reaction are reported. From left to right, the reactants ($2N$), transition state (TS), the various N_2 adsorption sites (α , and γ) and products (N_2 gas phase) are shown.

one at h_{Fe} and one at h_{Co} , while the transition state (TS) is located near the sb site (see snapshots on the right). For this reason, we do not report results resolved for the two coordinations $C_{N,Fe}$ and $C_{N,Co}$, as they are identical by symmetry. Instead, we emphasize that, in the product states, the adsorption states for N_2 , horizontal (α) and vertical (γ), are degenerate and can be resolved based on the two different metals. However, similarly to what was observed for nitrogen adsorption (Figure 7.5), these states are close in energy and are of little interest to the reaction, since once N_2^* is recombined, it can easily desorb from the surface. Regarding the free energy required for this process, the free energy barrier at the TS is approximately 190 kJ/mol. This represents the rate-determining step of the reaction, even in the case of FeCo.

7.1.3 Comparison of elementary step with Fe

Having studied all the elementary steps of the reaction, we can now compare our results with those previously obtained for Fe. In Table 7.1, we report the free energy barriers (ΔG^\ddagger) for all the elementary steps of the reaction on the two different catalysts, obtained from MD simulations at the operating temperature of 700 K.

We can immediately observe that the great agreement we found in the reaction mechanism is also reflected in the free energy barriers. Specifically, for the three

Free energy barriers ΔG^\ddagger [kJ/mol] at T=700 K		
	Fe(110)	FeCo(110)
$\text{NH}_3^* \longrightarrow \text{NH}_2^* + \text{H}^*$	97	106
$\text{NH}_2^* \longrightarrow \text{NH}^* + \text{H}^*$	76	77
$\text{NH}^* \longrightarrow \text{N}^* + \text{H}^*$	121	127
$2\text{N}^* \longrightarrow \text{N}_2$	216	192

Table 7.1: Estimated free energy barrier for each elementary step of ammonia decomposition on Fe(110) and FeCo(110) surfaces at T=700 K. The values are derived from the associated free energy profiles. Considering the accuracy of the ML potentials and the sampling uncertainty of the free energy calculations, we estimated an uncertainty within the range [5-10] kJ mol⁻¹.

dehydrogenation steps, ΔG^\ddagger is always slightly higher for FeCo, but with a difference of at most 9 kJ/mol. This result is difficult to interpret given the higher affinity of hydrogen for Co. Further investigations will be required to understand the cause, which could stem from a slight difference in the surface dynamics of the intermediates or be related to the pseudopotentials used in the DFT calculations.

On the other hand, for nitrogen recombination, which is clearly the rate-limiting step for both catalysts, a significantly lower barrier was obtained for FeCo compared to Fe, with $\Delta(\Delta G^\ddagger) = -24$ kJ/mol. Although not dramatic, this reduction in the rate-limiting step has a notable effect on the reaction rates, which are exponentially dependent on ΔG^\ddagger . Under the reasonable assumption that the pre-exponential factors are the same for both reactions, from Equation 2.21 we have:

$$\frac{t_{1/2}^{\text{Fe}}}{t_{1/2}^{\text{FeCo}}} = \frac{k^{\text{FeCo}}}{k^{\text{Fe}}} = \exp\left(-\frac{\Delta G_{\text{FeCo}}^\ddagger - \Delta G_{\text{Fe}}^\ddagger}{RT}\right) = \exp\left(-\frac{\Delta(\Delta G^\ddagger)}{RT}\right) \approx 60$$

That is, at this temperature, the reaction proceeds approximately 60 times faster on FeCo than on Fe.

Our results indicate that the reaction mechanism and free energy barriers on FeCo are similar to those on Fe, challenging the hypothesis that Co induces a radically different mechanism [30]. Moreover, our findings are in line with those reported in Ref. [51], where the authors measured similar apparent activation energies between Fe and FeCo.

However, we have not yet considered one of the reactions that proved crucial in explaining the Fe catalyst: nitrogen migration into the bulk, which will be the focus of the next section.

7.2 Nitrogen migration inside the bulk

We now proceed to study the migration process of a nitrogen atom into the bulk material. As we have seen, in the case of the iron catalyst, this process was able to explain the nitridation of the material during ammonia decomposition. This represents the first step toward the formation of bulk nitrides. Therefore, we repeat these studies on FeCo to understand its behavior as the reaction progresses. Again, we must emphasize that Fe and FeCo share many structural properties under operating conditions. Thus, we can utilize the same tools that we implemented to study the dissolution/segregation of nitrogen on the $(1\ 1\ 0)$ surface of iron.

In fact, on FeCo, nitrogen migration inside the bulk proceeds along the same “diffusion channels” of Fe. We have thus employed the collective variables (CVs) described in Equations 2.31, using $v_{0\bar{1}0}$ as the bias variable and applying reflective walls when $\|\tilde{v}_{1\bar{1}\bar{1}}\| > 0.55$ and $\|\tilde{v}_{\bar{1}\bar{1}\bar{1}}\| > 0.55$ to constrain the simulation within the channel.

In Figure 7.7, we show the free energy profile along $v_{0\bar{1}0}$ for both Fe (blue line, same as in Fig. 4.8) and FeCo.

Focusing on FeCo, starting from the surface hollow adsorption site (specifically the most stable h_{Fe} in this case), nitrogen migration proceeds to the first octahedral interstitial site, which corresponds to the small shoulder at $v_{0\bar{1}0} \approx 6$, and then to a second octahedral site, which is the first metastable one. From there, nitrogen alternates between deeper octahedral interstitial sites.

It is important to note that in the case of FeCo, not all octahedral sites are equivalent. Until now, we have intentionally described the FeCo crystal lattice imprecisely as BCC to highlight its similarity to Fe. However, because Fe and Co are two different chemical species, the correct description is a primitive cubic lattice with a base, with a Fe atom at $(0,0,0)$ and a Co atom at $(a/2, a/2, a/2)$, or equivalently, two interpenetrating primitive cubic lattices. This structure, known as “cesium chloride” or simply B2 structure, breaks the symmetry along the $\langle 1\ 1\ 1 \rangle$ direction, resulting in a change in chemical species. As a result, there are two types of octahedral interstitial sites: the first type has nitrogen at the center of the face of the Fe simple cubic lattice and, consequently, at the midpoint of the side of the Co simple cubic lattice (Figure 7.7, inset a); the second type is identical but with Fe and Co swapped (Figure 7.7, inset b). From a chemical perspective, the two sites are different, as in the first case, nitrogen forms two shorter bonds with Co atoms and four longer bonds with Fe, while in the second case, the opposite occurs. As can be seen from the free energy profile, the first site is more stable than the second, similar to what was observed for the two hollow sites on the surface, confirming the lower affinity of nitrogen for Co.




Figure 7.7: Free energy profiles of nitrogen dissolution/segregation on Fe(110) and FeCo(110) slabs at $T=700$ K. The free energy is projected along the $v_{0\bar{1}0}$ direction, and for Fe is reprinted from Figure 4.8. The same constraints of Fe (Figure 4.8) were adopted for FeCo. In the insets, we report two snapshots of the different FeCo octahedral interstitials where N is placed at a) the center of the face of the Fe simple cubic lattice and b) at the center of the face of the Co simple cubic lattice. The uncertainty along the free energy profiles is computed as the standard deviation of four independent simulations performed with four different MLPs. Additionally, to indicate and highlight the value of the free energy inside the bulk, the linear interpolation of the metastable free energy minima ($v_{0\bar{1}0} < 4.5$) is reported (dashed lines).

As for iron, the choice to limit our study to the $[010]$ direction was made to reduce the complexity of the problem. Indeed, nitrogen can equivalently move from one octahedral interstitial site to any of the six adjacent octahedral sites along the $\langle 100 \rangle$ directions, while only two are considered here.

If we now focus on the free energy profile, for nitrogen to migrate to the first metastable interstitial site at $v_{0\bar{1}0} \approx 5$, it faces a barrier of approximately 210 kJ/mol. Subsequently, it must overcome a first barrier of about 45 kJ/mol and a second of 30 kJ/mol to reach the “stable” interstitial site at $v_{0\bar{1}0} \approx 3$. From this point, migration into the bulk proceeds with barriers of about 60 and 30 kJ/mol.

It is interesting to note that the barriers for dissolution are approximately 5 kJ/mol higher than those for segregation. This is also reflected in the average values of the metastable free energy minima inside the bulk, whose value increases as migration progresses into the bulk (represented by the orange dashed line). This means that if nitrogen successfully dissolves below the first layers of FeCo, the probability that the process continues is lower than that of the opposite process of segregation. This is in stark contrast to what was observed for Fe (blue line).

Additionally, when observing the free energy values in the interstitial sites, these are significantly higher in FeCo. The minimum difference is as much as 80 kJ/mol when considering nitrogen in an interstitial site of the third layer ($v_{0\bar{1}0} \approx 3$).

If we compare the barriers on the two surfaces, for nitrogen to migrate into the bulk ($v_{0\bar{1}0} \geq 3.5$), it must overcome a free energy barrier of 223 kJ/mol in the case of FeCo, while for Fe, this barrier is 169 kJ/mol, with $\Delta(\Delta G^\ddagger) = 54$ kJ/mol. Similarly to what was observed for nitrogen recombination, if we compare the $t_{1/2}$ of the same process in the two catalysts, we find that the probability of observing dissolution in FeCo with respect to Fe is approximately 1:10000.

Before comparing this process with other dehydrogenation steps, we want to introduce lateral interactions, which, as we observed in the case of Fe, have a considerable impact on both processes.

7.3 Lateral interactions

In this section, we introduce the effect of finite coverage on the surface and its impact on the two most important reactions: recombination and bulk migration. Before studying these reactions, we had to refine the MLP by expanding the training set to simulate these phenomena accurately. We report this methodological step here because we consider it of great value in demonstrating the developed method's effectiveness and potential applications.

7.3.1 Efficiently refining the ML potential with DEAL

To expand the training set and accurately capture lateral interactions, we performed five active learning cycles with DEAL, using the version where the Gaussian Process (GP) is trained from scratch each time, and the DFT calculations are executed in an embarrassingly parallel manner, as described in Section 6.3.

Although lateral interactions represent a new and highly complex phenomenon to model, it was possible to start directly with DEAL since the potential already contained configurations of adsorbates on the surface, albeit isolated. Therefore, it

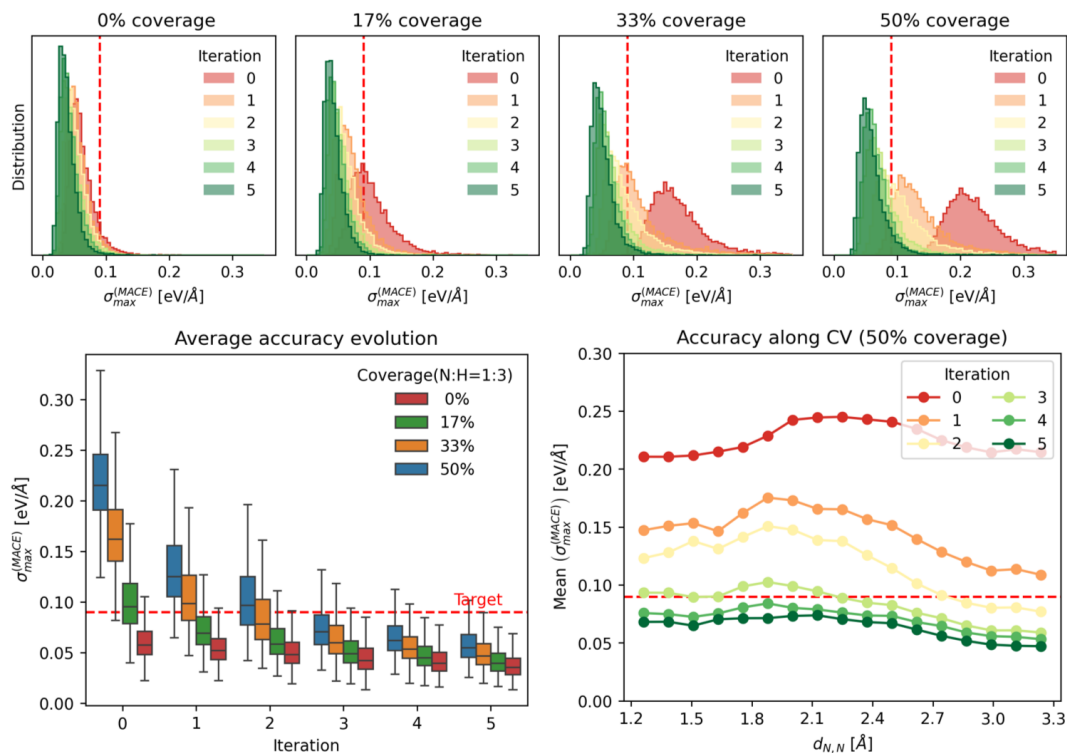


Figure 7.8: Evolution of MLP accuracy during active learning iterations for N-N recombination reaction at various coverage (17%, 33%, and 50%, monolayer coverage with ratio N : H = 1 : 3) at $T=700$ K. The evaluation is performed along the same trajectory for each coverage, sampled along reactive CV with well-tempered probability using the final model (MLP5). On top panels, the distributions of the maximum force uncertainty $\sigma_{\max}^{\text{MACE}}$. In the bottom panels (left) the boxplot of the same distributions and (right) the average value of $\sigma_{\max}^{\text{MACE}}$ along the reaction coordinate ($d_{N,N}$) for the highest coverage studied.

was able to produce reasonable predictions from the beginning.

In the first cycle, we used the initial potential (MLP0, the production model of previous studies but unrefined for lateral interactions) to perform a first set of short unbiased simulations (100 ps) where we varied the concentration of adsorbates on the surface to 17%, 33%, and 50%, considering only H, only N, and N and H in a 1:1 and 1:3 ratio. In this initial iteration, all the generated configurations were new to the potential and could be included in the training set without redundancy. Therefore, we preselected the configurations simply by extracting one frame every 2.5 ps for each trajectory, rather than using the *query-by-committee* criterion, which is less reliable in this regime. We then proceeded according to the DEAL framework, extracting a non-redundant subset of configurations through the GP, recalculating these configurations at the DFT level, including them in the training set, and training

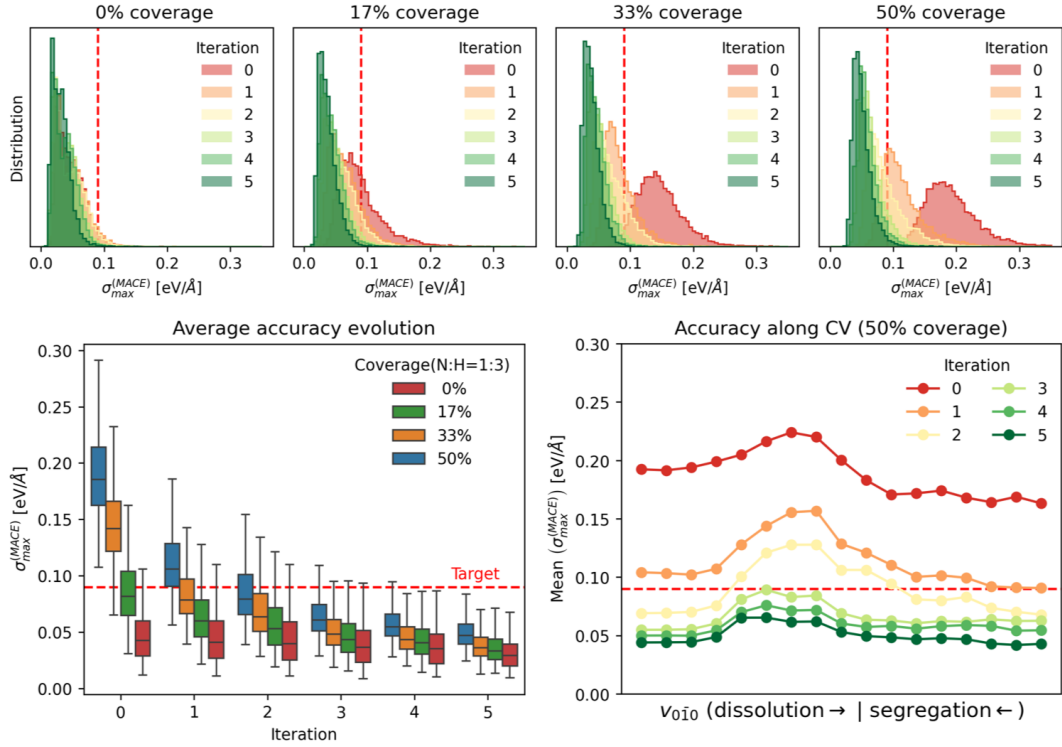


Figure 7.9: Analogue of Figure 7.8 for the N dissolution/segregation reaction.

a new committee of potentials (MLP1).

We ran the same simulations with MLP1 and performed a DEAL cycle, including the preselection via *query-by-committee*, obtaining MLP2, which was already capable of accurately describing the surface in the absence of reactions.

We then focused on the two reactions of interest, recombination and bulk migration, and performed three additional DEAL cycles, starting from longer reactive simulations (at least 1 ns) and focusing only on surface coverages with an $N:H$ ratio of 1:1 and 1:3.

To validate the approach, we measured the accuracy of the different MLPs during converged reactive simulations obtained with the final potential (MLP5) at different coverages (Figures 7.8 and 7.9). From this study, it is interesting to observe that, in the case of high coverage, many active learning cycles are required to bring the potential to the target accuracy, demonstrating the complexity of the phenomenon being studied.

The evolution of the potential's accuracy along the reaction coordinate for the simulations at the highest coverage (50%), shown in the bottom right panel of Figure 7.8 for recombination and Figure 7.9 for bulk migration, is also noteworthy. This graph clearly shows that during the first two iterations (performed on non-reactive

simulations), the accuracy of the potential in the region of the reactants and products reaches the target level. However, the region around the transition state (TS) is still poorly described. The subsequent cycles, which are carried out on reactive simulations, are necessary to achieve uniform accuracy along the entire reaction pathway.

This demonstrates the effectiveness of DEAL in identifying the configurations that should be included in the training set to achieve the required accuracy.

Selected configuration with Data-Efficient Active Learning (DEAL)

Iteration	1	2	3	4	5	Total
Pre-selection	3863	2606	8089	24297	11086	49941
GP selection	194	252	911	755	917	3029

Table 7.2: **DEAL selection.** Selected configuration with Data-Efficient Active Learning (DEAL) during the 5 iterations executed during the study of finite coverage effects on reactivity. The GP threshold parameter was set to 0.1 for all the iterations except for iteration 4, where it was set to 0.12 since the large number of configurations preselected (longer simulations). See text for details.

In Table 7.2, we report the number of preselected and selected configurations during the five iterations. The most remarkable number is the total number of configurations collected, calculated at the DFT level, and included in the potential: 3029.

Considering also the configurations previously collected, the total number of DFT calculations used to study the entire catalytic process of ammonia decomposition on FeCo, including finite coverage, amounts to 8200. This number is to be contrasted with the approximately 400k calculations done for Fe. We leave it to the reader to draw the obvious conclusions about the effectiveness and savings achieved.

Before moving on to the presentation of the results, we would like to make one last consideration regarding the practicality of the method applied here. The active learning cycles performed with DEAL in this manner are very fast. The initial configurations can be collected from many different simulations performed in parallel. In this specific case, the 1 ns simulations on the reduced-size system (necessary for recalculating it with DFT) require about 10 hours on a single modern GPU. Similarly, the GP trained from scratch can process around 1000 configurations in less than 24 hours on a few dozen CPUs. Finally, the DFT calculations are embarrassingly parallelizable, so, with sufficient computational resources, the time required for the relabeling step is equivalent to the time needed for a single-point DFT calculation. Retraining the GNN potential also requires little time, as the number of configurations in the training set remains small. In practice, each cycle presented took only 1

to 2 days to be completed.

7.3.2 Nitrogen recombination and dissolution

Having built a reliable potential, we can now study the effect of finite coverage on the recombination and bulk migration processes.

In fact, we know that lateral interactions between adsorbed atoms on Fe(110) play a significant role, influencing the adsorption site preferences and altering the free energy barriers for both recombination and migration (Section 4.4). To directly compare with Fe, we simulated the exact same setup, performing simulations with finite coverage of N and H atoms (20% and 40% of a monolayer with an N:H ratio of 1:3). We want to emphasize that, unlike the commonly used approach of estimating the effects of coverage through static calculations on small systems, here we explicitly simulated the dynamics on large systems at the *operando* temperature of 700 K.




Figure 7.10: Free energy profiles of nitrogen recombination (left) and dissolution/segregation (right) on the FeCo(110) slab at $T=700$ K at various coverages (0, 20, and 40% of a monolayer with a ratio N : H = 1 : 3). The free energies are computed identically to the 0% coverage case. The uncertainty is computed from a set of different configurations, see text for details.

MD simulations at various coverage levels indicate that the barrier for recombination decreases with increasing coverage (Figure 7.10, left panel). Similarly, the migration barrier decreases as nitrogen coverage increases (Figure 7.10, right panel). This is the same behavior we observed for Fe, although it is less drastic in this case.

Focusing on migration into the bulk, we observed that, with increasing coverage (20%), there is a reduction in the free energy barrier for entering the first interlayer interstitial octahedral site at $v_{0\bar{1}0} \approx 6$, which becomes slightly metastable at the maximum coverage investigated.

A final note on the uncertainty of the simulations is necessary (represented by the shaded area in Figure 7.10). For the simulations at 0% coverage, we estimated the uncertainty from the standard deviation of four independent simulations, each run with a different MLP model (trained on the same training set) for 6 ns. For the simulations at finite coverage, we performed 32 independent 6 ns simulations, initializing 8 different random adsorbate configurations and running each of them 4 times independently, using a different model for each. We used this strategy because MACE, like all GNN models in general, is about 10 times slower than DeePMD, making it impractical to run simulations of hundreds of nanoseconds. However, these time scales would be necessary to sample the process ergodically, as nitrogen atom diffusion on the surface occurs on the nanosecond scale at 700 K ($\Delta G^\ddagger \approx 60$ kJ/mol). To address this, we circumvented the problem by running many replicas, performing a total of approximately 200 ns of simulation for each coverage level. Of course, since the sampling is not ergodic (not with respect to the reaction coordinate but the surface configurations), each simulation converges to a slightly different free energy value, resulting in increasing uncertainty with higher coverage. Nonetheless, we note that these variations are small, and the uncertainty estimate provided, which accounts for both MLP uncertainty (by using different models) and sampling uncertainty (by performing independent simulations), is conservative, as we could have used more strict estimators like the standard deviation of the mean.

7.4 Experimental observations of no nitride formation

Having reconstructed all the free energy profiles of the entire ammonia decomposition process on the (1 1 0) surfaces of Fe (see Chapters 3 and 4) and FeCo through MD simulations, we are now able to rationalize the results obtained in light of experimental evidence.

We observed that, for both catalysts, the dehydrogenation steps are not rate-limiting, while the nitrogen recombination step is the limiting one. This last reaction must be compared with another competitive process for removing nitrogen from the surface: its dissolution into the bulk. In Figure 4.10, we report the overall free energy barrier for these two processes for Fe (left panel) and FeCo (right panel), considering the more realistic scenario of both finite temperature and coverage.

As already discussed in Section 4.5, our MD simulations reveal that nitrogen prefers to migrate into the bulk of Fe(110) rather than recombining on the surface. This finding is consistent with experimental observations of nitride formation during ammonia decomposition [208]. In particular, the experiments show the immediate release of H₂ after ammonia decomposition, while nitrogen remains on the surface

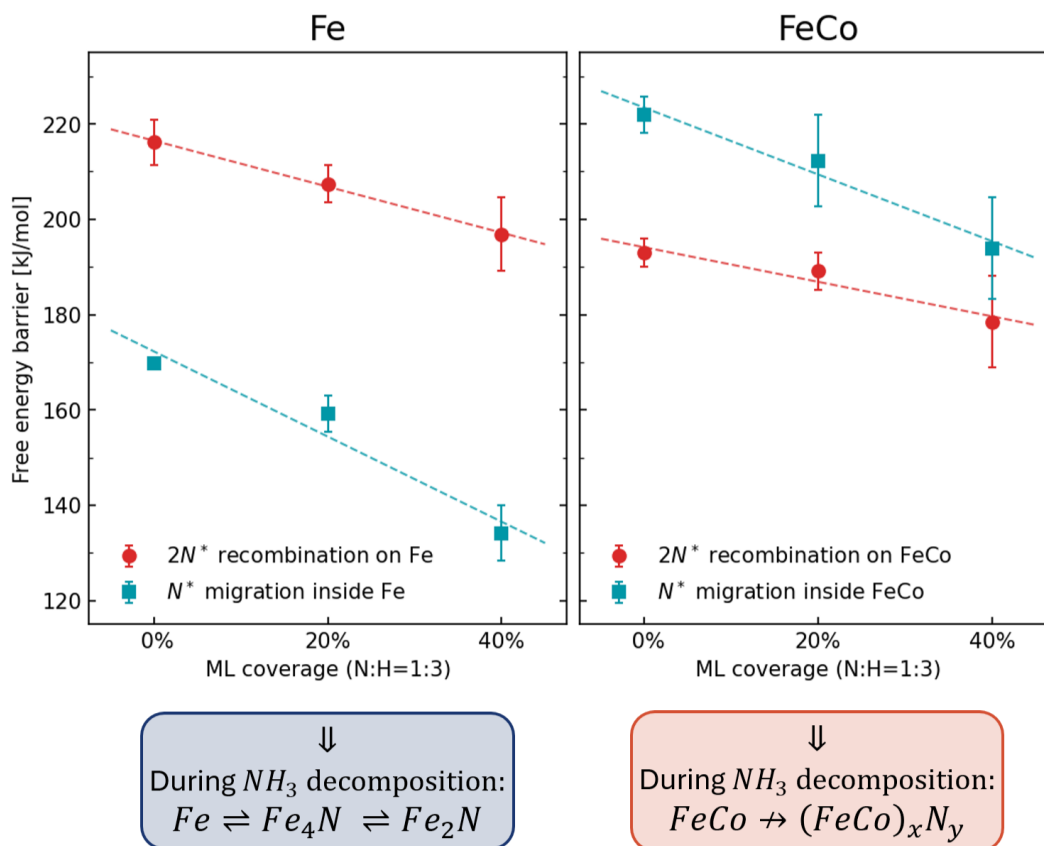


Figure 7.11: Free energy barrier of recombination (red circles) and migration inside (green squares) as functions of surface coverage for Fe (left) and FeCo (right), obtained from Figure 4.9 and Figure 7.10, respectively. The surface coverage refers to N, H monolayer coverage with an N:H ratio equal to 1:3. The dashed lines are obtained from linear regression. Given the relationships between the free energy barriers on the two different catalysts, different processes occur during ammonia decomposition, as illustrated in the boxes.

and migrates into the bulk, forming nitrides. After the formation and breakdown of nitrides (Fe_4N and Fe_2N) at elevated temperatures, nitrogen is finally released.

In contrast, our results show that this tendency is exactly the opposite in FeCo(1 1 0), with recombination being favored over migration into bulk. This remains true even at high coverage. Furthermore, as we discussed in Section 7.2, nitrogen dissolution is always favored over segregation even after the first migration barrier inside the bulk is overcome (Figure 7.7). From these results, we can infer that in FeCo, the formation of bulk nitrides is not expected, as it is blocked in the early stages.

However, this is not the only beneficial effect of adding cobalt to form an alloy. In fact, the free energy barrier of the rate-limiting step of the reaction, *i.e.*, N_2 release via N^* recombination, is consistently reduced on FeCo. The process occurs

approximately 50 times faster on the alloy.

Together with our experimental partners, we are gathering evidence of this dual promotional effect of the cobalt-iron alloy, although important information can already be derived from the present literature [30, 51]. From the study of the elementary steps of the reaction, our results indicate that the reaction mechanism and free energy barriers on FeCo are similar to those on Fe. This challenges the hypothesis that cobalt induces a radically different mechanism [30]. Instead, our findings are in line with those reported in Ref. [51], where the authors measured similar apparent activation energies between Fe and FeCo, as well as they agree with the long-time recognized structural sensitivity [31, 32]. More importantly, *in operando* XAS and XPS measurements suggest that FeCo alloys remain stable during ammonia decomposition without transitioning to a nitride phase, contrary to what was reported for the Fe catalyst. In conclusion, our results perfectly explain the experimental evidence of the improved catalytic activity of the FeCo alloy for ammonia decomposition [30, 51], providing a mechanistic explanation of all the reaction steps and the lack of nitride formation.

Furthermore, this demonstrates that simple explanations using descriptors [19, 21, 51, 52] are not always sufficient, and the processes involved in heterogeneous catalysis are far more complex, requiring a realistic modeling that includes dynamics and lateral interactions.

Conclusions and Perspectives

Catalytic processes are fundamental to both biological systems and industrial operations, driving everything from essential life-sustaining reactions to large-scale chemical production. Atomistic simulations have long been regarded as an important tool for gaining insight into these complex processes at the atomic and molecular levels. With recent advancements based on machine learning techniques and the increasing availability of computational resources, the field of atomistic simulations is expanding its scope toward *in silico* experiments of catalytic processes under realistic, *operando* conditions.

In this thesis, we have taken significant steps in this direction by presenting a comprehensive computational study of ammonia decomposition on iron and iron-cobalt alloy catalysts. Our work addressed critical challenges in understanding the atomic-level mechanisms driving catalysis, shifting from a static representation of catalytic processes to a dynamic and statistical approach using state-of-the-art molecular dynamics simulations.

Our results demonstrate that dynamic effects critically influence all stages of decomposition of ammonia, particularly on the Fe(111) surface, where adsorption, active site formation, and reaction mechanisms evolve with time and temperature. We showed that treating reactions dynamically is essential, as these reactions exist as ensembles of pathways rather than a single, well-defined path. This highlights the importance of integrating dynamical equilibrium concepts into the study of catalytic processes and materials, aligning with a perspective of the catalyst as a functional material [1, 2, 54].

Additionally, our investigation into nitridation underscored the complexity of catalytic behavior, challenging traditional models that rely on simple descriptors, such as the binding energy. Our findings indicate that it is the delicate balance between different surface processes that explains the actual formation surface or bulk iron nitrides and why monometallic iron becomes nitrated under certain conditions while iron-cobalt alloys do not. The results shed light on the microscopic origin of cobalt's role in preventing bulk nitridation, which is driven by the synergistic

effects of structural and electronic promotion. Collectively, these findings contribute to a deeper understanding of catalytic processes critical to ammonia decomposition, offering a rational basis for designing future catalysts for industrial applications, such as leverage alloying to prevent deactivation due to nitridation.

Furthermore, the novel machine learning potential protocol developed here offers a promising approach for studying complex catalytic systems with high accuracy and reduced computational cost. By focusing on data-efficient techniques, we demonstrated the ability to simulate key reaction mechanisms under realistic conditions, paving the way for routine use of MD simulations to study complex reaction networks.

Despite the progress made in this study, several avenues remain open for future research. While our work focused on iron and iron-cobalt alloys, exploring alternative catalysts, such as other bimetallic systems, and investigating the role of promoters could reveal further improvements in catalytic efficiency and durability. These advancements are well within reach thanks to our efficient and practical protocol, which could be further improved by integrating new enhanced sampling methods and more efficient architectures or even used to fine-tune foundational models.

Another natural extension of this work toward simulations directly comparable with experiments would be the inclusion of pressure effects through grand canonical sampling as well as the integration with microkinetic models. The latter could be exploited to obtain the (dynamical) equilibrium conditions at which to run MD simulations, thus reconstructing the free energy profiles under these conditions. These free energies could then be used to refine the microkinetic model in an iterative way.

Machine learning continues to hold great potential for advancing the study of catalytic processes, and future efforts should explore more sophisticated algorithms to generate and optimize catalysts or to identify reaction mechanisms. For instance, our atomic charge model successfully captured the driving force of the reaction, suggesting that ML techniques can be expanded not only as analysis tools but also to drive reactive simulations, further extending the scope of enhanced sampling techniques from a geometry-based description to a chemical one.

In conclusion, this work contributes to the ongoing efforts to develop green hydrogen technologies by demonstrating how machine learning and molecular dynamics can accelerate the search for solutions for sustainable futures. This goal has never been more urgent.

List of Publications

Articles published in peer-reviewed journals

1. Perego, S., Bonati, L., (2024), “Data efficient machine learning potentials for modeling catalytic reactivity via active learning and enhanced sampling.” *npj Comput Mater*, 10, 291. DOI: 10.1038/s41524-024-01481-6
2. Perego, S., Bonati, L., Tripathi, S., Parrinello, M., (2024) “How Dynamics Changes Ammonia Cracking on Iron Surfaces.” *ACS Catalysis*, 14 (19), 14652-14664. DOI: 10.1021/acscatal.4c01920
3. Purcel, M., Berendts, S., Perego, S., Bonati, L., Müller, A., Lerch, M., Parrinello, M., Muhler, M., (2024), “Iron Nitride Formation and Decomposition During Ammonia Decomposition over a Wustite-based Bulk Iron Catalyst.” *ACS Catalysis*, 14(18), 13947-13957. DOI: 10.1021/acscatal.4c04415
4. Tripathi, S., Bonati, L., Perego, S., Parrinello, M. (2024). “How Poisoning Is Avoided in a Step of Relevance to the Haber–Bosch Catalysis.” *ACS Catalysis*, 14(7), 4944-4950. DOI: 10.1021/acscatal.3c06201

Preprints and articles submitted or in preparation

5. Perego, S., Purcel, M., Chen, S., Muhler, M., Behrens, M., Parrinello, M., Bonati, L., (2024), “Unraveling the promotional effect of cobalt-alloying on Fe catalysts for ammonia decomposition via machine learning-driven molecular dynamics and experiments” In preparation

Bibliography

- [1] R. Schlögl. “Heterogeneous Catalysis”. In: *Angewandte Chemie International Edition* 54.11 (2015), pp. 3465–3520.
- [2] R. Schlögl. “The functional interface in catalysis”. In: *Arkivoc* 2024.3 (2024).
- [3] J. Humphreys, R. Lan, and S. Tao. “Development and Recent Progress on Ammonia Synthesis Catalysts for Haber–Bosch Process”. In: *Advanced Energy and Sustainability Research* 2.1 (2021), p. 2000043.
- [4] S. Giddey, S. P. Badwal, and A. Kulkarni. “Review of electrochemical ammonia production technologies and materials”. In: *International Journal of Hydrogen Energy* 38.34 (2013), pp. 14576–14594.
- [5] M. Neurock. “The microkinetics of heterogeneous catalysis. By J. A. Dumesic, D. F. Rudd, L. M. Aparicio, J. E. Rekoske, and A. A. Treviño, ACS Professional Reference Book, American Chemical Society, Washington, DC, 1993, 315 pp.” In: *AIChE Journal* 40.6 (1994), pp. 1085–1087.
- [6] M. Bowker. “Modelling of ammonia synthesis kinetics”. In: *Catalysis Today* 12.2-3 (1992), pp. 153–163.
- [7] P. Stoltze and J. K. Nørskov. “The surface science based ammonia kinetics revisited”. In: *Topics in Catalysis* 1.3-4 (1994), pp. 253–263.
- [8] International Energy Agency. *Ammonia Technology Roadmap*. Tech. rep. IEA, 2021.
- [9] S. Ornes. “Green ammonia could produce climate-friendly ways to store energy and fertilize farms”. In: *Proceedings of the National Academy of Sciences of the United States of America* 118.49 (2021), e2119584118.
- [10] H. Ishaq and C. Crawford. “Review of ammonia production and utilization: Enabling clean energy transition and net-zero climate targets”. In: *Energy Conversion and Management* 300 (2024), p. 117869.
- [11] B. S. Zainal, P. J. Ker, H. Mohamed, H. C. Ong, I. M. Fattah, S. M. Rahman, L. D. Nghiem, and T. M. Mahlia. “Recent advancement and assessment of green hydrogen production technologies”. In: *Renewable and Sustainable Energy Reviews* 189 (2024), p. 113941.
- [12] D. R. MacFarlane, P. V. Cherepanov, J. Choi, B. H. Suryanto, R. Y. Hodgetts, J. M. Bakker, F. M. Ferrero Vallana, and A. N. Simonov. “A Roadmap to the Ammonia Economy”. In: *Joule* 4.6 (2020), pp. 1186–1205.

- [13] N. Morlanés, S. P. Katikaneni, S. N. Paglieri, A. Harale, B. Solami, S. M. Sarathy, and J. Gascon. “A technological roadmap to the ammonia energy economy: Current state and missing technologies”. In: *Chemical Engineering Journal* 408 (2021), p. 127310.
- [14] A. Valera-Medina, H. Xiao, M. Owen-Jones, W. I. David, and P. J. Bowen. “Ammonia for power”. In: *Progress in Energy and Combustion Science* 69 (2018), pp. 63–102.
- [15] I. Lucentini, X. Garcia, X. Vendrell, and J. Llorca. “Review of the Decomposition of Ammonia to Generate Hydrogen”. In: *Industrial and Engineering Chemistry Research* 60.51 (2021), pp. 18560–18611.
- [16] C. Zamfirescu and I. Dincer. “Using ammonia as a sustainable fuel”. In: *Journal of Power Sources* 185.1 (2008), pp. 459–465.
- [17] A. Klerke, C. H. Christensen, J. K. Nørskov, and T. Vegge. “Ammonia for hydrogen storage: challenges and opportunities”. In: *Journal of Materials Chemistry* 18.20 (2008), pp. 2304–2310.
- [18] T. E. Bell and L. Torrente-Murciano. “H₂ Production via Ammonia Decomposition Using Non-Noble Metal Catalysts: A Review”. In: *Topics in Catalysis* 2016 59:15 59.15 (2016), pp. 1438–1457.
- [19] D. A. Hansgen, D. G. Vlachos, and J. G. Chen. “Using first principles to predict bimetallic catalysts for the ammonia decomposition reaction”. In: *Nature Chemistry* 2010 2:6 2.6 (2010), pp. 484–489.
- [20] J. C. Ganley, F. S. Thomas, E. G. Seebauer, and R. I. Masel. “A priori catalytic activity correlations: The difficult case of hydrogen production from ammonia”. In: *Catalysis Letters* 96.3-4 (2004), pp. 117–122.
- [21] A. Boisen, S. Dahl, J. K. Nørskov, and C. H. Christensen. “Why the optimal ammonia synthesis catalyst is not the optimal ammonia decomposition catalyst”. In: *Journal of Catalysis* 230.2 (2005), pp. 309–312.
- [22] R. Schlögl. “Ammonia Synthesis”. In: *Handbook of Heterogeneous Catalysis*. John Wiley & Sons, Ltd, 2008, pp. 2501–2575.
- [23] A. Mittasch and W. Frankenburg. “Early Studies of Multicomponent Catalysts”. In: *Advances in Catalysis* 2.C (1950), pp. 81–104.
- [24] F. Bozso, G. Ertl, M. Grunze, and M. Weiss. “Interaction of nitrogen with iron surfaces: I. Fe(100) and Fe(111)”. In: *Journal of Catalysis* 49.1 (1977), pp. 18–41.
- [25] F. Bozso, G. Ertl, and M. Weiss. “Interaction of nitrogen with iron surfaces: II. Fe(110)”. In: *Journal of Catalysis* 50.3 (1977), pp. 519–529.
- [26] B. Lu, L. Li, M. Ren, Y. Liu, Y. Zhang, X. Xu, X. Wang, and H. Qiu. “Ammonia decomposition over iron-based catalyst: Exploring the hidden active phase”. In: *Applied Catalysis B: Environmental* 314 (2022), p. 121475.
- [27] L. Wang, Y. Yi, Y. Zhao, R. Zhang, J. Zhang, and H. Guo. “NH₃ Decomposition for H₂ Generation: Effects of Cheap Metals and Supports on Plasma-Catalyst Synergy”. In: *ACS Catalysis* 5.7 (2015), pp. 4167–4174.
- [28] M. Weiss, G. Ertl, and F. Nitschké. “Adsorption and decomposition of ammonia on Fe(110)”. In: *Applications of Surface Science* 2.4 (1979), pp. 614–635.
- [29] M. Grunze, F. Bozso, G. Ertl, and M. Weiss. “Interaction of ammonia with Fe(111) and Fe(100) surfaces”. In: *Applications of Surface Science* 1.2 (1978), pp. 241–265.

- [30] Y. Wang, J. Qian, Z. Fang, M. R. Kunz, G. Yablonsky, A. Fortunelli, W. A. Goddard, and R. R. Fushimi. “Understanding Reaction Networks through Controlled Approach to Equilibrium Experiments Using Transient Methods”. In: *Journal of the American Chemical Society* 143.29 (2021), pp. 10998–11006.
- [31] N. D. Spencer, R. C. Schoonmaker, and G. A. Somorjai. “Iron single crystals as ammonia synthesis catalysts: Effect of surface structure on catalyst activity”. In: *Journal of Catalysis* 74.1 (1982), pp. 129–135.
- [32] G. A. Somorjai and N. Materer. “Surface structures in ammonia synthesis”. In: *Topics in Catalysis* 1994 1:3 1.3 (1994), pp. 215–231.
- [33] K. Homann, H. Kuhlenbeck, and H. J. Freund. “The interaction of N₂ with iron on W(110), Pd(111) and Rh(111)”. In: *Zeitschrift für Physikalische Chemie* 198.1-2 (1997), pp. 135–147.
- [34] C. M. Goodwin et al. “Operando probing of the surface chemistry during the Haber–Bosch process”. In: *Nature* 2024 625:7994 625.7994 (2024), pp. 282–286.
- [35] S. Satoh, H. Fujimoto, and H. Kobayashi. “Theoretical study of NH₃ adsorption on Fe(110) and Fe(111) surfaces”. In: *Journal of Physical Chemistry B* 110.10 (2006), pp. 4846–4852.
- [36] R. J. Lin, F. Y. Li, and H. L. Chen. “Computational investigation on adsorption and dissociation of the NH₃ molecule on the Fe(111) surface”. In: *Journal of Physical Chemistry C* 115.2 (2011), pp. 521–528.
- [37] G. Lanzani and K. Laasonen. “NH₃ adsorption and dissociation on a nanosized iron cluster”. In: *International Journal of Hydrogen Energy* 35.13 (2010), pp. 6571–6577.
- [38] X. Duan, J. Ji, G. Qian, C. Fan, Y. Zhu, X. Zhou, D. Chen, and W. Yuan. “Ammonia decomposition on Fe(1 1 0), Co(1 1 1) and Ni(1 1 1) surfaces: A density functional theory study”. In: *Journal of Molecular Catalysis A: Chemical* 357 (2012), pp. 81–86.
- [39] S. C. Yeo, S. S. Han, and H. M. Lee. “Mechanistic investigation of the catalytic decomposition of ammonia (NH₃) on an Fe(100) surface: A DFT study”. In: *Journal of Physical Chemistry C* 118.10 (2014), pp. 5309–5316.
- [40] X. Zhang, Z. Lu, D. Ma, and Z. Yang. “Adsorption and dissociation of ammonia on small iron clusters”. In: *International Journal of Hydrogen Energy* 40.1 (2015), pp. 346–352.
- [41] L. Xu, D. Kirvassilis, Y. Bai, and M. Mavrikakis. “Atomic and molecular adsorption on Fe(110)”. In: *Surface Science* 667 (2018), pp. 54–65.
- [42] J. Qian, Q. An, A. Fortunelli, R. J. Nielsen, and W. A. Goddard. “Reaction Mechanism and Kinetics for Ammonia Synthesis on the Fe(111) Surface”. In: *Journal of the American Chemical Society* 140.20 (2018), pp. 6288–6297.
- [43] J. M. P. Martirez and E. A. Carter. “First-Principles Insights into the Thermocatalytic Cracking of Ammonia-Hydrogen Blends on Fe(110): 1. Thermodynamics”. In: *The Journal of Physical Chemistry C* 126.46 (2022), pp. 19733–19744.
- [44] L. Bonati, D. Polino, C. Pizzolitto, P. Biasi, R. Eckert, S. Reitmeier, R. Schlögl, and M. Parrinello. “The role of dynamics in heterogeneous catalysis: Surface diffusivity and N₂ decomposition on Fe(111)”. In: *Proceedings of the National Academy of Sciences* 120.50 (2023), e2313023120.

- [45] J. Ji, X. Duan, G. Qian, X. Zhou, G. Tong, and W. Yuan. "Towards an efficient CoMo/ γ -Al₂O₃ catalyst using metal amine metallate as an active phase precursor: Enhanced hydrogen production by ammonia decomposition". In: *International Journal of Hydrogen Energy* 39.24 (2014), pp. 12490–12498.
- [46] S. F. Zaman, L. A. Jolaoso, S. Podila, A. A. Al-Zahrani, Y. A. Alhamed, H. Driss, M. M. Daous, and L. Petrov. "Ammonia decomposition over citric acid chelated γ -Mo₂N and Ni₂Mo₃N catalysts". In: *International Journal of Hydrogen Energy* 43.36 (2018), pp. 17252–17258.
- [47] K. G. Kirste et al. "CO_x-free hydrogen production from ammonia – mimicking the activity of Ru catalysts with unsupported Co-Re alloys". In: *Applied Catalysis B: Environmental* 280 (2021), p. 119405.
- [48] B. Lorenzut, T. Montini, M. Bevilacqua, and P. Fornasiero. "FeMo-based catalysts for H₂ production by NH₃ decomposition". In: *Applied Catalysis B: Environmental* 125 (2012), pp. 409–417.
- [49] J. Zhang, J. O. Müller, W. Zheng, D. Wang, D. Su, and R. Schlögl. "Individual Fe-Co alloy nanoparticles on carbon nanotubes: Structural and catalytic properties". In: *Nano Letters* 8.9 (2008), pp. 2738–2743.
- [50] S. B. Simonsen, D. Chakraborty, I. Chorkendorff, and S. Dahl. "Alloyed Ni-Fe nanoparticles as catalysts for NH₃ decomposition". In: *Applied Catalysis A: General* 447-448 (2012), pp. 22–31.
- [51] S. Chen et al. "Highly loaded bimetallic iron-cobalt catalysts for hydrogen release from ammonia". In: *Nature Communications* 2024 15:1 15.1 (2024), pp. 1–11.
- [52] S. Dahl, A. Logadottir, C. J. Jacobsen, and J. K. Nørskov. "Electronic factors in catalysis: the volcano curve and the effect of promotion in catalytic ammonia synthesis". In: *Applied Catalysis A: General* 222.1-2 (2001), pp. 19–29.
- [53] X. Duan, J. Ji, X. Yan, G. Qian, D. Chen, and X. Zhou. "Understanding Co-Mo Catalyzed Ammonia Decomposition: Influence of Calcination Atmosphere and Identification of Active Phase". In: *ChemCatChem* 8.5 (2016), pp. 938–945.
- [54] M. S. Spencer. "Stable and metastable metal surfaces in heterogeneous catalysis". In: *Nature* 1986 323:6090 323.6090 (1986), pp. 685–687.
- [55] Y. G. Wang, D. Mei, V. A. Glezakou, J. Li, and R. Rousseau. "Dynamic formation of single-atom catalytic active sites on ceria-supported gold nanoparticles". In: *Nature Communications* 2015 6:1 6.1 (2015), pp. 1–8.
- [56] B. Eren et al. "Activation of Cu(111) surface by decomposition into nanoclusters driven by CO adsorption". In: *Science* 351.6272 (2016), pp. 475–478.
- [57] L. Grajciar, C. J. Heard, A. A. Bondarenko, M. V. Polynski, J. Meeprasert, E. A. Pidko, and P. Nachtigall. "Towards operando computational modeling in heterogeneous catalysis". In: *Chemical Society Reviews* 47.22 (2018), pp. 8307–8348.
- [58] G. Sun and P. Sautet. "Metastable Structures in Cluster Catalysis from First-Principles: Structural Ensemble in Reaction Conditions and Metastability Triggered Reactivity". In: *Journal of the American Chemical Society* 140.8 (2018), pp. 2812–2820.

- [59] N. Gerrits and G. J. Kroes. “Curious Mechanism of the Dissociative Chemisorption of Ammonia on Ru(0001)”. In: *Journal of Physical Chemistry C* 123.46 (2019), pp. 28291–28300.
- [60] W. L. Li, C. N. Lininger, K. Chen, V. Vaissier Welborn, E. Rossomme, A. T. Bell, M. Head-Gordon, and T. Head-Gordon. “Critical Role of Thermal Fluctuations for CO Binding on Electrocatalytic Metal Surfaces”. In: *JACS Au* 1.10 (2021), pp. 1708–1718.
- [61] G. Piccini et al. “Ab initio molecular dynamics with enhanced sampling in heterogeneous catalysis”. In: *Catalysis Science & Technology* 12.1 (2022), pp. 12–37.
- [62] V. Van Speybroeck, M. Bocus, P. Cnudde, and L. Vanduyfhuys. “Operando Modeling of Zeolite-Catalyzed Reactions Using First-Principles Molecular Dynamics Simulations”. In: *ACS Catalysis* 13.17 (2023), pp. 11455–11493.
- [63] M. Yang, U. Raucci, and M. Parrinello. “Reactant-induced dynamics of lithium imide surfaces during the ammonia decomposition process”. In: *Nature Catalysis* 2023 6:9 6.9 (2023), pp. 829–836.
- [64] F. Mambretti, U. Raucci, M. Yang, and M. Parrinello. “How Does Structural Disorder Impact Heterogeneous Catalysts? The Case of Ammonia Decomposition on Nonstoichiometric Lithium Imide”. In: *ACS Catalysis* 16 (2024), pp. 1252–1256.
- [65] A. T. Gardini, U. Raucci, and M. Parrinello. “A Bulk Phase Transformation Drives Ammonia Synthesis on Barium Hydride”. In: (2024).
- [66] M. Yang, U. Raucci, and M. Parrinello. “Ammonia Decomposition on Lithium Imide Surfaces: A New Paradigm in Heterogeneous Catalysis”. In: (2022).
- [67] D. Alfè, M. J. Gillan, and G. D. Price. “The melting curve of iron at the pressures of the Earth’s core from ab initio calculations”. In: *Nature* 1999 401:6752 401.6752 (1999), pp. 462–464.
- [68] G. J. Kroes, A. Gross, E. J. Baerends, M. Scheffler, and D. A. McCormack. “Quantum theory of dissociative chemisorption on metal surfaces”. In: *Accounts of Chemical Research* 35.3 (2002), pp. 193–200.
- [69] B. G. Walker, C. Molteni, and N. Marzari. “Ab initio molecular dynamics of metal surfaces”. In: *Journal of Physics: Condensed Matter* 16.26 (2004), S2575.
- [70] L. Verlet. “Computer “Experiments” on Classical Fluids. I. Thermodynamical Properties of Lennard-Jones Molecules”. In: *Physical Review* 159.1 (1967), p. 98.
- [71] W. C. Swope, H. C. Andersen, P. H. Berens, and K. R. Wilson. “A computer simulation method for the calculation of equilibrium constants for the formation of physical clusters of molecules: Application to small water clusters”. In: *The Journal of Chemical Physics* 76.1 (1982), pp. 637–649.
- [72] D. Frenkel and B. Smit. *Understanding molecular simulation: from algorithms to applications*. Vol. 1. Elsevier, 2001.
- [73] H. C. Andersen. “Molecular dynamics simulations at constant pressure and/or temperature”. In: *The Journal of Chemical Physics* 72.4 (1980), pp. 2384–2393.
- [74] H. J. Berendsen, J. P. Postma, W. F. Van Gunsteren, A. Dinola, and J. R. Haak. “Molecular dynamics with coupling to an external bath”. In: *The Journal of Chemical Physics* 81.8 (1984), pp. 3684–3690.

- [75] S. Nosé. “A unified formulation of the constant temperature molecular dynamics methods”. In: *The Journal of Chemical Physics* 81.1 (1984), pp. 511–519.
- [76] W. G. Hoover. “Canonical dynamics: Equilibrium phase-space distributions”. In: *Physical Review A* 31.3 (1985), p. 1695.
- [77] G. J. Martyna, M. E. Tuckerman, D. J. Tobias, and M. L. Klein. “Explicit reversible integrators for extended systems dynamics”. In: *Molecular Physics* 87.5 (1996), pp. 1117–1157.
- [78] G. Bussi, D. Donadio, and M. Parrinello. “Canonical sampling through velocity rescaling”. In: *Journal of Chemical Physics* 126.1 (2007), p. 14101.
- [79] J. P. Perdew, K. Burke, and M. Ernzerhof. “Generalized Gradient Approximation Made Simple”. In: *Physical Review Letters* 77.18 (1996), p. 3865.
- [80] P. J. Stephens, F. J. Devlin, C. F. Chabalowski, and M. J. Frisch. “Ab Initio calculation of vibrational absorption and circular dichroism spectra using density functional force fields”. In: *Journal of Physical Chemistry* 98.45 (1994), pp. 11623–11627.
- [81] C. Adamo and V. Barone. “Toward reliable density functional methods without adjustable parameters: The PBE0 model”. In: *The Journal of Chemical Physics* 110.13 (1999), pp. 6158–6170.
- [82] J. Sun, A. Ruzsinszky, and J. Perdew. “Strongly Constrained and Appropriately Normed Semilocal Density Functional”. In: *Physical Review Letters* 115.3 (2015), p. 036402.
- [83] V. I. Anisimov, F. Aryasetiawan, and A. I. Lichtenstein. “First-principles calculations of the electronic structure and spectra of strongly correlated systems: the LDA+ U method”. In: *Journal of Physics: Condensed Matter* 9.4 (1997), p. 767.
- [84] V. I. Anisimov and O. Gunnarsson. “Density-functional calculation of effective Coulomb interactions in metals”. In: *Physical Review B* 43.10 (1991), p. 7570.
- [85] V. I. Anisimov, J. Zaanen, and O. K. Andersen. “Band theory and Mott insulators: Hubbard U instead of Stoner I ”. In: *Physical Review B* 44.3 (1991), p. 943.
- [86] E. K. Gross and W. Kohn. “Time-Dependent Density-Functional Theory”. In: *Advances in Quantum Chemistry* 21.C (1990), pp. 255–291.
- [87] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg. “A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H-Pu”. In: *Journal of Chemical Physics* 132.15 (2010).
- [88] S. Grimme, S. Ehrlich, and L. Goerigk. “Effect of the damping function in dispersion corrected density functional theory”. In: *Journal of Computational Chemistry* 32.7 (2011), pp. 1456–1465.
- [89] F. Chiter, V. B. Nguyen, N. Tarrat, M. Benoit, H. Tang, and C. Lacaze-Dufaure. “Effect of van der Waals corrections on DFT-computed metallic surface properties”. In: *Materials Research Express* 3.4 (2016), p. 046501.
- [90] D. Dragoni, D. Ceresoli, and N. Marzari. “Thermoelastic properties of α -iron from first-principles”. In: *PHYSICAL REVIEW B* 91 (2015), p. 104105.

- [91] J. P. Perdew and K. Schmidt. “Jacob’s ladder of density functional approximations for the exchange-correlation energy”. In: *AIP Conference Proceedings* 577.1 (2001), pp. 1–20.
- [92] W. Thiel. “Semiempirical quantum-chemical methods”. In: *Wiley Interdisciplinary Reviews: Computational Molecular Science* 4.2 (2014), pp. 145–157.
- [93] G. Seifert and J. O. Joswig. “Density-functional tight binding—an approximate density-functional theory method”. In: *Wiley Interdisciplinary Reviews: Computational Molecular Science* 2.3 (2012), pp. 456–465.
- [94] F. H. Stillinger and T. A. Weber. “Computer simulation of local order in condensed phases of silicon”. In: *Physical Review B* 31.8 (1985), p. 5262.
- [95] S. M. Foiles, M. I. Baskes, and M. S. Daw. “Embedded-atom-method functions for the fcc metals Cu, Ag, Au, Ni, Pd, Pt, and their alloys”. In: *Physical Review B* 33.12 (1986), p. 7983.
- [96] J. Huang and A. D. Mackerell. “CHARMM36 all-atom additive protein force field: Validation based on comparison to NMR data”. In: *Journal of computational chemistry* 34.25 (2013), p. 2135.
- [97] P. K. Weiner and P. A. Kollman. “AMBER: Assisted model building with energy refinement. A general program for modeling molecules and their interactions”. In: *Journal of Computational Chemistry* 2.3 (1981), pp. 287–303.
- [98] W. L. Jorgensen and J. Tirado-Rives. “The OPLS Potential Functions for Proteins. Energy Minimizations for Crystals of Cyclic Peptides and Crambin”. In: *Journal of the American Chemical Society* 110.6 (1988), pp. 1657–1666.
- [99] A. C. T. van Duin, S. Dasgupta, F. Lorant, and W. A. Goddard. “ReaxFF: A Reactive Force Field for Hydrocarbons”. In: *The Journal of Physical Chemistry A* 105.41 (2001), pp. 9396–9409.
- [100] T. B. Blank, S. D. Brown, A. W. Calhoun, and D. J. Doren. “Neural network models of potential energy surfaces”. In: *The Journal of Chemical Physics* 103.10 (1995), pp. 4129–4137.
- [101] J. Behler and M. Parrinello. “Generalized neural-network representation of high-dimensional potential-energy surfaces”. In: *Physical Review Letters* 98.14 (2007), p. 146401.
- [102] A. Stukowski. “Visualization and analysis of atomistic simulation data with OVITO—the Open Visualization Tool”. In: *Modelling and Simulation in Materials Science and Engineering* 18.1 (2009), p. 015012.
- [103] O. T. Unke, S. Chmiela, H. E. Sauceda, M. Gastegger, I. Poltavsky, K. T. Schütt, A. Tkatchenko, and K.-R. Müller. “Machine Learning Force Fields”. In: *Chemical Reviews* 121.16 (2021), pp. 10142–10186.
- [104] V. L. Deringer, M. A. Caro, and G. Csányi. “Machine Learning Interatomic Potentials as Emerging Tools for Materials Science”. In: *Advanced Materials* 31.46 (2019), p. 1902765.
- [105] L. Himanen, M. O. Jäger, E. V. Morooka, F. Federici Canova, Y. S. Ranawat, D. Z. Gao, P. Rinke, and A. S. Foster. “DScribe: Library of descriptors for machine learning in materials science”. In: *Computer Physics Communications* 247 (2020), p. 106949.

- [106] G. Wang, C. Wang, X. Zhang, Z. Li, J. Zhou, and Z. Sun. “Machine learning interatomic potential: Bridge the gap between small-scale models and realistic device-scale simulations”. In: *iScience* 27.5 (2024).
- [107] L. Zhang, J. Han, H. Wang, W. Saidi, R. Car, and W. E. “End-to-end Symmetry Preserving Inter-atomic Potential Energy Model for Finite and Extended Systems”. In: *Advances in Neural Information Processing Systems*. Vol. 31. 2018.
- [108] R. Drautz. “Atomic cluster expansion for accurate and transferable interatomic potentials”. In: *Physical Review B* 99.1 (2019), p. 014104.
- [109] J. Zeng et al. “DeePMD-kit v2: A software package for deep potential models”. In: *Journal of Chemical Physics* 159.5 (2023), p. 54801.
- [110] H. Wang, L. Zhang, J. Han, and W. E. “DeePMD-kit: A deep learning package for many-body potential energy representation and molecular dynamics”. In: *Computer Physics Communications* 228 (2018), pp. 178–184.
- [111] I. Batatia, D. P. Kovacs, G. Simm, C. Ortner, and G. Csányi. “MACE: Higher order equivariant message passing neural networks for fast and accurate force fields”. In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 11423–11436.
- [112] D. P. Kovacs, I. Batatia, E. S. Arany, and G. Csanyi. “Evaluation of the MACE Force Field Architecture: from Medicinal Chemistry to Materials Science”. In: *The Journal of Chemical Physics* 159.4 (2023), p. 44118.
- [113] A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi. “Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons”. In: *Physical review letters* 104.13 (2010), p. 136403.
- [114] J. Vandermause, S. B. Torrisi, S. Batzner, Y. Xie, L. Sun, A. M. Kolpak, and B. Kozinsky. “On-the-fly active learning of interpretable Bayesian force fields for atomistic rare events”. In: *npj Computational Materials* 6.1 (2020), p. 20.
- [115] J. Vandermause, Y. Xie, J. S. Lim, C. J. Owen, and B. Kozinsky. “Active learning of reactive Bayesian force fields applied to heterogeneous catalysis dynamics of H/Pt”. In: *Nature Communications* 13.1 (2022), p. 5183.
- [116] R. Drautz. “Atomic cluster expansion for accurate and transferable interatomic potentials”. In: *Physical Review B* 99.1 (2019), p. 014104.
- [117] C. Schran, K. Brezina, and O. Marsalek. “Committee neural network potentials control generalization errors and enable active learning”. In: *The Journal of Chemical Physics* 153.10 (2020).
- [118] Z. Zhang, B. Zandkarimi, and A. N. Alexandrova. “Ensembles of Metastable States Govern Heterogeneous Catalysis on Dynamic Interfaces”. In: *Accounts of Chemical Research* 53.2 (2020), pp. 447–458.
- [119] J. S. Smith, B. Nebgen, N. Lubbers, O. Isayev, and A. E. Roitberg. “Less is more: Sampling chemical space with active learning”. In: *The Journal of chemical physics* 148.24 (2018).
- [120] S. J. Ang, W. Wang, D. Schwalbe-Koda, S. Axelrod, and R. Gómez-Bombarelli. “Active learning accelerates ab initio molecular dynamics on reactive energy surfaces”. In: *Chem* 7.3 (2021), pp. 738–751.

- [121] L. L. Schaaf, E. Fako, S. De, A. Schäfer, and G. Csányi. “Accurate energy barriers for catalytic reaction pathways: an automatic training protocol for machine learning force fields”. In: *npj Computational Materials* 2023 9:1 9.1 (2023), pp. 1–10.
- [122] H. Jung, L. Sauerland, S. Stocker, K. Reuter, and J. T. Margraf. “Machine-learning driven global optimization of surface adsorbate geometries”. In: *npj Computational Materials* 2023 9:1 9.1 (2023), pp. 1–8.
- [123] D. Yoo, J. Jung, W. Jeong, and S. Han. “Metadynamics sampling in atomic environment space for collecting training data for machine learning potentials”. In: *npj Computational Materials* 7.1 (2021), p. 131.
- [124] J. Xu, X.-M. Cao, and P. Hu. “Accelerating metadynamics-based free-energy calculations with adaptive machine learning potentials”. In: *Journal of chemical theory and computation* 17.7 (2021), pp. 4465–4476.
- [125] L. Bonati and M. Parrinello. “Silicon Liquid Structure and Crystal Nucleation from Ab Initio Deep Metadynamics”. In: *Physical Review Letters* 121.26 (2018), p. 265701.
- [126] H. Niu, L. Bonati, P. M. Piaggi, and M. Parrinello. “Ab initio phase diagram and nucleation of gallium”. In: *Nature Communications* 11.1 (2020), pp. 1–9.
- [127] O. Abou El Kheir, L. Bonati, M. Parrinello, and M. Bernasconi. “Unraveling the crystallization kinetics of the Ge₂Sb₂Te₅ phase change compound with a machine-learned interatomic potential”. In: *npj Computational Materials* 10.1 (2024), p. 33.
- [128] M. Galib and D. T. Limmer. “Reactive uptake of N₂O₅ by atmospheric aerosol is dominated by interfacial processes”. In: *Science* 371.6532 (2021), pp. 921–925.
- [129] M. Yang, L. Bonati, D. Polino, and M. Parrinello. “Using metadynamics to build neural network potentials for reactive events: the case of urea decomposition in water”. In: *Catalysis Today* 387 (2022), pp. 143–149.
- [130] J. Zhang, O. Zhang, L. Bonati, and T. Hou. “Combining Transition Path Sampling with Data-Driven Collective Variables through a Reactivity-Biased Shooting Algorithm”. In: *Journal of Chemical Theory and Computation* 20.11 (2024). PMID: 38801759, pp. 4523–4532.
- [131] X. Guan, J. P. Heindel, T. Ko, C. Yang, and T. Head-Gordon. “Using machine learning to go beyond potential energy surface benchmarking for chemical reactivity”. In: *Nature Computational Science* 2023 3:11 3.11 (2023), pp. 965–974.
- [132] R. David, M. de la Puente, A. Gomez, O. Anton, G. Stirnemann, and D. Laage. “ArcaNN: automated enhanced sampling generation of training sets for chemically reactive machine learning interatomic potentials”. In: *arXiv:2407.07751 [PREPRINT]* (2024).
- [133] H. Zhang, V. Juraskova, and F. Duarte. “Modelling chemical processes in explicit solvents with machine learning potentials”. In: *Nature Communications* 2024 15:1 15.1 (2024), pp. 1–11.
- [134] T. A. Young, T. Johnston-Wood, H. Zhang, and F. Duarte. “Reaction dynamics of Diels–Alder reactions from machine learned potentials”. In: *Physical Chemistry Chemical Physics* 24.35 (2022), pp. 20820–20827.
- [135] Z. Zeng, F. Wodaczek, K. Liu, F. Stein, J. Hutter, J. Chen, and B. Cheng. “Mechanistic insight on water dissociation on pristine low-index TiO₂ surfaces from machine

- learning molecular dynamics simulations”. In: *Nature Communications* 2023 14:1 14.1 (2023), pp. 1–9.
- [136] R. David, I. Tuñón, and D. Laage. “Competing Reaction Mechanisms of Peptide Bond Formation in Water Revealed by Deep Potential Molecular Dynamics and Path Sampling”. In: *Journal of the American Chemical Society* 146.20 (2024), pp. 14213–14224.
- [137] Z. Benayad, R. David, and G. Stirnemann. “Prebiotic chemical reactivity in solution with quantum accuracy and microsecond sampling using neural network potentials”. In: *Proceedings of the National Academy of Sciences of the United States of America* 121.23 (2024), e2322040121.
- [138] M. Kulichenko, K. Barros, N. Lubbers, Y. W. Li, R. Messerly, S. Tretiak, J. S. Smith, and B. Nebgen. “Uncertainty-driven dynamics for active learning of interatomic potentials”. In: *Nature Computational Science* 3.3 (2023), pp. 230–239.
- [139] C. van der Oord, M. Sachs, D. P. Kovács, C. Ortner, and G. Csányi. “Hyperactive learning for data-driven interatomic potentials”. In: *npj Computational Materials* 2023 9:1 9.1 (2023), pp. 1–14.
- [140] A. R. Tan, J. C. B. Dietschreit, and R. Gomez-Bombarelli. “Enhanced sampling of robust molecular datasets with uncertainty-based collective variables”. In: (2024).
- [141] V. Zaverkin, D. Holzmüller, H. Christiansen, F. Errica, F. Alesiani, M. Takamoto, M. Niepert, and J. Kästner. “Uncertainty-biased molecular dynamics for learning uniformly accurate interatomic potentials”. In: *npj Computational Materials* 10.1 (2024), p. 83.
- [142] A. P. Thompson et al. “LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales”. In: *Computer Physics Communications* 271 (2022), p. 108171.
- [143] A. Hjorth Larsen et al. “The atomic simulation environment—a Python library for working with atoms”. In: *Journal of Physics: Condensed Matter* 29.27 (2017), p. 273002.
- [144] G. A. Tribello, M. Bonomi, D. Branduardi, C. Camilloni, and G. Bussi. “PLUMED 2: New feathers for an old bird”. In: *Computer Physics Communications* 185.2 (2014), pp. 604–613.
- [145] P. Giannozzi et al. “Advanced capabilities for materials modelling with Quantum ESPRESSO”. In: *Journal of Physics: Condensed Matter* 29.46 (2017), p. 465901.
- [146] G. M. Torrie and J. P. Valleau. “Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling”. In: *Journal of Computational Physics* 23.2 (1977), pp. 187–199.
- [147] M. Mezei. “Adaptive umbrella sampling: Self-consistent determination of the non-Boltzmann bias”. In: *Journal of Computational Physics* 68.1 (1987), pp. 237–248.
- [148] A. Laio and M. Parrinello. “Escaping free-energy minima”. In: *Proceedings of the National Academy of Sciences* 99.20 (2002), pp. 12562–12566.
- [149] A. Barducci, G. Bussi, and M. Parrinello. “Well Tempered Metadynamics: A Smoothly Converging and Tunable Free-Energy Method”. In: *Phys. Rev. Lett.* 100 (2 2008), p. 020603.

- [150] O. Valsson and M. Parrinello. “Variational Approach to Enhanced Sampling and Free Energy Calculations”. In: *Phys. Rev. Lett.* 113 (9 2014), p. 090601.
- [151] L. Bonati, Y.-Y. Zhang, and M. Parrinello. “Neural networks-based variationally enhanced sampling”. In: *Proceedings of the National Academy of Sciences* 116.36 (2019), pp. 17641–17647.
- [152] Y. Sugita and Y. Okamoto. “Replica-exchange molecular dynamics method for protein folding”. In: *Chemical Physics Letters* 314.1 (1999), pp. 141–151.
- [153] A. Lyubartsev, A. Martsinovski, S. Shevkunov, and P. Vorontsov-Velyaminov. “New approach to Monte Carlo calculation of the free energy: Method of expanded ensembles”. In: *The Journal of chemical physics* 96.3 (1992), pp. 1776–1783.
- [154] M. Invernizzi, P. M. Piaggi, and M. Parrinello. “Unified Approach to Enhanced Sampling”. In: *Phys. Rev. X* 10 (4 2020), p. 041034.
- [155] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler. “Transition path sampling: Throwing ropes over rough mountain passes, in the dark”. In: *Annual review of physical chemistry* 53.1 (2002), pp. 291–318.
- [156] J. Hénin, T. Lelièvre, M. R. Shirts, O. Valsson, and L. Delemotte. “Enhanced Sampling Methods for Molecular Dynamics Simulations”. In: *Living Journal of Computational Molecular Science* 4.1 (2022), p. 1583.
- [157] M. Invernizzi and M. Parrinello. “Rethinking Metadynamics: From Bias Potentials to Probability Distributions”. In: *Journal of Physical Chemistry Letters* 11.7 (2020), pp. 2731–2736.
- [158] M. Invernizzi and M. Parrinello. “Exploration vs Convergence Speed in Adaptive-Bias Enhanced Sampling”. In: *Journal of Chemical Theory and Computation* 18.6 (2022), pp. 3988–3996.
- [159] M. Bonomi et al. “Promoting transparency and reproducibility in enhanced molecular simulations”. In: *Nature Methods* 2019 16:8 16.8 (2019), pp. 670–673.
- [160] J. C. Dietschreit, D. J. Diestler, and C. Ochsenfeld. “How to obtain reaction free energies from free-energy profiles”. In: *The Journal of Chemical Physics* 156.11 (2022).
- [161] J. C. Dietschreit, D. J. Diestler, A. Hulm, C. Ochsenfeld, and R. Gómez-Bombarelli. “From free-energy profiles to activation free energies”. In: *The Journal of Chemical Physics* 157.8 (2022).
- [162] O. Valsson, P. Tiwary, and M. Parrinello. “Enhancing Important Fluctuations: Rare Events and Metadynamics from a Conceptual Viewpoint”. In: *Annual Review of Physical Chemistry* 67 (1 2016), pp. 159–184.
- [163] M. Invernizzi and M. Parrinello. “Rethinking metadynamics: from bias potentials to probability distributions”. In: *The Journal of Physical Chemistry Letters* 11.7 (2020), pp. 2731–2736.
- [164] N. Ansari, V. Rizzi, and M. Parrinello. “Water regulates the residence time of Benzamidine in Trypsin”. In: *Nature Communications* (2022).
- [165] D. Ray, N. Ansari, V. Rizzi, M. Invernizzi, and M. Parrinello. “Rare event kinetics from adaptive bias enhanced sampling”. In: *Journal of Chemical Theory and Computation* (2022).

- [166] H. Grubmüller. “Predicting slow structural transitions in macromolecular systems: Conformational flooding”. In: *Phys. Rev. E* 52 (3 1995), pp. 2893–2906.
- [167] J. McCarty, O. Valsson, P. Tiwary, and M. Parrinello. “Variationally optimized free-energy flooding for rate calculation”. In: *Physical review letters* 115.7 (2015), p. 070601.
- [168] A. F. Voter. “Hyperdynamics: Accelerated molecular dynamics of infrequent events”. In: *Physical Review Letters* 78.20 (1997), p. 3908.
- [169] D. Ray and M. Parrinello. “Kinetics from Metadynamics: Principles, Applications, and Outlook”. In: *Journal of Chemical Theory and Computation* 19.17 (2023). PMID: 37585703, pp. 5649–5670.
- [170] I. Marcos-Alcalde, E. López-Viñas, and P. Gómez-Puertas. “MEPSAnd: minimum energy path surface analysis over n-dimensional surfaces”. In: *Bioinformatics* 36.3 (2020), pp. 956–958.
- [171] R. F. Bader. “Atoms in Molecules”. In: *Accounts of Chemical Research* 18.1 (1985), pp. 9–15.
- [172] G. Henkelman, A. Arnaldsson, and H. Jónsson. “A fast and robust algorithm for Bader decomposition of charge density”. In: *Computational Materials Science* 36.3 (2006), pp. 354–360.
- [173] K. T. Schütt, H. E. Saucedo, P. J. Kindermans, A. Tkatchenko, and K. R. Müller. “SchNet - A deep learning architecture for molecules and materials”. In: *Journal of Chemical Physics* 148.24 (2018), p. 241722.
- [174] D. P. Kingma and J. L. Ba. “Adam: A Method for Stochastic Optimization”. In: *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings* (2014).
- [175] A. F. Voter. “A method for accelerating the molecular dynamics simulation of infrequent events”. In: *The Journal of Chemical Physics* 106.11 (1997), pp. 4665–4677.
- [176] P. Tiwary and M. Parrinello. “From metadynamics to dynamics”. In: *Physical Review Letters* 111.23 (2013), p. 230602.
- [177] E. Schubert and L. Lenssen. “Fast k-medoids Clustering in Rust and Python”. In: *Journal of Open Source Software* 7.75 (2022), p. 4183.
- [178] A. Stukowski. “Computational analysis methods in atomistic modeling of crystals”. In: *JOM* 66.3 (2014), pp. 399–407.
- [179] P. M. Piaggi and M. Parrinello. “Calculation of phase diagrams in the multithermal-multibaric ensemble”. In: *The Journal of chemical physics* 150.24 (2019).
- [180] A. P. Bartók, R. Kondor, and G. Csányi. “On representing chemical environments”. In: *Phys. Rev. B* 87 (18 2013), p. 184115.
- [181] E. J. Maginn, R. A. Messerly, D. J. Carlson, D. R. Roe, and J. R. Elliot. “Best Practices for Computing Transport Properties 1. Self-Diffusivity and Viscosity from Equilibrium Molecular Dynamics [Article v1.0]”. In: *Living Journal of Computational Molecular Science* 1.1 (2019), pp. 6324–6324.
- [182] R. Gowers et al. “MDAnalysis: A Python Package for the Rapid Analysis of Molecular Dynamics Simulations”. In: *Proceedings of the 15th Python in Science Conference* (2016), pp. 98–105.

- [183] V. Calandrini, E. Pellegrini, P. Calligari, K. Hinsén, and G. Kneller. “nMoldyn - Interfacing spectroscopic experiments, molecular dynamics simulations and models for time correlation functions”. In: *École thématique de la Société Française de la Neutronique* 12 (2011), pp. 201–232.
- [184] P. d. Buyl. “tidynamics: A tiny package to compute the dynamics of stochastic and molecular simulations”. In: *Journal of Open Source Software* 3.28 (2018), p. 877.
- [185] S. Tripathi, L. Bonati, S. Perego, and M. Parrinello. “How Poisoning Is Avoided in a Step of Relevance to the Haber–Bosch Catalysis”. In: *ACS Catalysis* (2024), pp. 4944–4950.
- [186] L. Zhang, J. Han, H. Wang, and R. Car. “Deep Potential Molecular Dynamics: A Scalable Model with the Accuracy of Quantum Mechanics”. In: *Physical Review Letters* 120 (2018), p. 143001.
- [187] G. Henkelman and H. Jónsson. “A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives”. In: *The Journal of Chemical Physics* 111.15 (1999), pp. 7010–7022.
- [188] J. Kästner and P. Sherwood. “Superlinearly converging dimer method for transition state search”. In: *The Journal of Chemical Physics* 128.1 (2008), p. 014106.
- [189] P. Giannozzi et al. “QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials”. In: *Journal of Physics: Condensed Matter* 21.39 (2009), p. 395502.
- [190] P. Giannozzi et al. “Quantum ESPRESSO toward the exascale”. In: *The Journal of Chemical Physics* 152.15 (2020), p. 154105.
- [191] A. M. Rappe, K. M. Rabe, E. Kaxiras, and J. D. Joannopoulos. “Optimized pseudopotentials”. In: *Physical Review B* 41.2 (1990), p. 1227.
- [192] N. Marzari, D. Vanderbilt, A. De Vita, and M. C. Payne. “Thermal Contraction and Disordering of the Al(110) Surface”. In: *Physical Review Letters* 82.16 (1999), p. 3296.
- [193] H. J. Monkhorst and J. D. Pack. “Special points for Brillouin-zone integrations”. In: *Physical Review B* 13.12 (1976), p. 5188.
- [194] G. Prandini, A. Marrazzo, I. E. Castelli, N. Mounet, and N. Marzari. “Precision and efficiency in solid-state pseudopotential calculations”. In: *npj Computational Materials* 2018 4:1 4.1 (2018), pp. 1–13.
- [195] I. Batatia, D. P. Kovács, G. N. C. Simm, C. Ortner, and G. Csányi. “MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields”. In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 11423–11436.
- [196] Y. Xie, J. Vandermause, L. Sun, A. Cepellotti, and B. Kozinsky. “Bayesian force fields from active learning for simulation of inter-dimensional transformation of stanene”. In: *npj Computational Materials* 2021 7:1 7.1 (2021), pp. 1–10.
- [197] L. Hammer, H. Landskron, W. Nichtl-Pecher, A. Fricke, K. Heinz, and K. Müller. “Hydrogen-induced restructuring of close-packed metal surfaces: H/Ni(111) and H/Fe(110)”. In: *Physical Review B* 47.23 (1993), p. 15969.

- [198] D. K. Escott, S. J. Pratt, and D. A. King. “Evidence for a nitrogen-induced reconstruction of Fe{1 1 1}”. In: *Surface Science* 562.1-3 (2004), pp. 226–236.
- [199] J. A. Moulijn, A. E. Van Diepen, and F. Kapteijn. “Catalyst deactivation: is it predictable?: What to do?” In: *Applied Catalysis A: General* 212.1-2 (2001), pp. 3–16.
- [200] D. Chandler. “Statistical mechanics of isomerization dynamics in liquids and the transition state approximation”. In: *The Journal of Chemical Physics* 68.6 (1978), pp. 2959–2970.
- [201] P. Kang, E. Trizio, and M. Parrinello. “Computing the committor with the committor to study the transition state ensemble”. In: *Nature Computational Science* 2024 (2024), pp. 1–10.
- [202] J. J. Mortensen, M. V. Ganduglia-Pirovano, L. B. Hansen, B. Hammer, P. Stoltze, and J. K. Nørskov. “Nitrogen adsorption on Fe(111), (100), and (110) surfaces”. In: *Surface Science* 422.1-3 (1999), pp. 8–16.
- [203] D. Liu, W. Zhao, and Q. Yuan. “Breaking the Linear Relation in the Dissociation of Nitrogen on Iron Surfaces”. In: *ChemPhysChem* 23.17 (2022).
- [204] J. Chen, Y. Chen, H. Wang, and P. Hu. “Correlation between the structural features and intrinsic activity trend of Fe surfaces for ammonia synthesis”. In: *Catalysis Science & Technology* 13.18 (2023), pp. 5237–5247.
- [205] M. Kadowaki et al. “Roles of Interstitial Nitrogen, Carbon, and Boron in Steel Corrosion: Generation of Oxyanions and Stabilization of Electronic Structure”. In: *Journal of The Electrochemical Society* 167.8 (2020), p. 081503.
- [206] C. Domain, C. S. Becquart, and J. Foct. “*Ab initio* study of foreign interstitial atom (C, N) interactions with intrinsic point defects in α -Fe”. In: *Physical Review B* 69.14 (2004), p. 144112.
- [207] Z. Lv, J. Fan, K. Guan, Z. Wu, D. Zhao, and W. Fu. “Effects of interstitial atoms (N/O) in bcc Fe from first-principle calculations”. In: *Fusion Engineering and Design* 137 (2018), pp. 22–29.
- [208] M. Purcel, S. Berendts, L. Bonati, S. Perego, A. Müller, M. Lerch, M. Parrinello, and M. Muhler. “Iron Nitride Formation and Decomposition during Ammonia Decomposition over a Wustite-Based Bulk Iron Catalyst”. In: *ACS Catalysis* (2024), pp. 13947–13957.
- [209] G. Ertl, M. Huber, and N. Thiele. “Formation and Decomposition of Nitrides on Iron Surfaces”. In: *Zeitschrift für Naturforschung - Section A Journal of Physical Sciences* 34.1 (1979), pp. 30–39.
- [210] K. Kielbasa, R. Pelka, and W. Arabczyk. “Studies of the kinetics of ammonia decomposition on promoted nanocrystalline iron using gas phases of different nitriding degree”. In: *Journal of Physical Chemistry A* 114.13 (2010), pp. 4531–4534.
- [211] R. Pelka, K. Kielbasa, and W. Arabczyk. “Catalytic ammonia decomposition during nanocrystalline iron nitriding at 475°C with NH₃/H₂ mixtures of different nitriding potentials”. In: *Journal of Physical Chemistry C* 118.12 (2014), pp. 6178–6185.

- [212] J. U. Nwalor, J. G. Goodwin, and P. Biloen. “Steady-state isotopic transient-kinetic analysis of iron-catalyzed ammonia synthesis”. In: *Journal of Catalysis* 117.1 (1989), pp. 121–134.
- [213] M. Muhler, F. Rosowski, and G. Ertl. “The dissociative adsorption of N₂ on a multiply promoted iron catalyst used for ammonia synthesis: a temperature-programmed desorption study”. In: *Catalysis Letters* 24.3-4 (1994), pp. 317–331.
- [214] G. Ertl, M. Huber, S. B. Lee, Z. Paál, and M. Weiss. “Interactions of nitrogen and hydrogen on iron surfaces”. In: *Applications of Surface Science* 8.4 (1981), pp. 373–386.
- [215] G. Ertl, S. Lee, and M. Weiss. “Kinetics of nitrogen adsorption on Fe(111)”. In: *Surface Science* 114.2-3 (1982), pp. 515–526.
- [216] M. Salmeron and R. Schlögl. “Ambient pressure photoelectron spectroscopy: A new tool for surface science and nanotechnology”. In: *Surface Science Reports* 63.4 (2008), pp. 169–199.
- [217] J. I. Falk, L. Bonati, A. Simulations, P. Novelli, M. Parrinello, and M. Pontil. “Transfer learning for atomistic simulations using GNNs and kernel mean embeddings”. In: *Advances in Neural Information Processing Systems* 36 (2023), pp. 29783–29797.
- [218] L. Chanussot et al. “Open Catalyst 2020 (OC20) Dataset and Community Challenges”. In: *ACS Catalysis* 11.10 (2021), pp. 6059–6072.
- [219] I. Batatia et al. “A foundation model for atomistic materials chemistry”. In: (2023).
- [220] S. Vandenhoute, M. Cools-Ceuppens, S. DeKeyser, T. Verstraelen, and V. Van Speybroeck. “Machine learning potentials for metal-organic frameworks using an incremental learning approach”. In: *npj Computational Materials* 2023 9:1 9.1 (2023), pp. 1–8.
- [221] C. E. Rasmussen. “Gaussian processes in machine learning”. In: *Summer school on machine learning*. Springer, 2003, pp. 63–71.
- [222] S. Perego, L. Bonati, S. Tripathi, and M. Parrinello. “How Dynamics Changes Ammonia Cracking on Iron Surfaces”. In: *ACS Catalysis* 14 (2024), pp. 14652–14664.
- [223] D. Ray, N. Ansari, V. Rizzi, M. Invernizzi, and M. Parrinello. “Rare Event Kinetics from Adaptive Bias Enhanced Sampling”. In: *Journal of Chemical Theory and Computation* 18.11 (2022), pp. 6500–6509.
- [224] U. Raucci, V. Rizzi, and M. Parrinello. “Discover, Sample, and Refine: Exploring Chemistry with Enhanced Sampling Techniques”. In: *Journal of Physical Chemistry Letters* 13.6 (2022), pp. 1424–1430.
- [225] I. Batatia et al. “The Design Space of E(3)-Equivariant Atom-Centered Interatomic Potentials”. In: (2022).
- [226] C. J. Owen et al. “Complexity of many-body interactions in transition metals via machine-learned force fields from the TM23 data set”. In: *npj Computational Materials* 10.1 (2024), p. 92.
- [227] X. Fu, Z. Wu, W. Wang, T. Xie, S. Keten, R. Gomez-Bombarelli, and T. Jaakkola. “Forces are not enough: Benchmark and critical evaluation for machine learning force fields with molecular simulations”. In: *arXiv preprint arXiv:2210.07237* (2022).

-
- [228] L. Bonati, E. Trizio, A. Rizzi, and M. Parrinello. “A unified framework for machine learning collective variables for enhanced sampling simulations: mlcolvar”. In: *The Journal of Chemical Physics* 159.1 (2023).

Acknowledgement

Completing this PhD means a great deal to me, and as I reach the end of this chapter, I feel a deep obligation to express my gratitude to many people.

First and foremost, I am profoundly grateful to Professor Michele Parrinello for making this thesis possible. His exceptional guidance and unwavering presence have taught me invaluable lessons, not just in science but in overcoming my limitations. It has been an extraordinary privilege to work under the mentorship of such a brilliant and inspiring scientist. Learning to approach research with the child's boundless curiosity, coupled with his monumental experience, has been truly invaluable.

A special thank you goes to Luigi, whose patience and insight have been crucial throughout this journey, not just scientifically but also personally. His energy, creativity, and mentorship were both inspiring and instructive. I am honored to be the first of what I am sure will be many students under his guidance.

I am also deeply grateful to Professor Marco Bernasconi for his support as my tutor during the early stages of this PhD. Without his advice, I might never have taken this path.

My sincere thanks go to Professor Erio Tosatti and Professor Daniela Polino for their meticulous review of my thesis and valuable suggestions, which greatly enriched my work. I also wish to thank Professor Vassiliki-Alexandra Glezakou for her availability and input during this process.

I am thankful to the international institutions that welcomed me during these years. I am especially grateful to Professor Annabella Selloni and Professor Roberto Car for hosting me at Princeton University. I also thank Pablo and all the CSI members for their stimulating discussions. I am profoundly grateful to Professor Beatriz Roldán Cuenya at the Fritz Haber Institute for her careful guidance and support. Special thanks go to Daniel, Thomas, and the experimental group for allowing me to engage in hands-on work and making my time in Berlin unforgettable. Finally, I thank Professor Vittorio Limongelli for hosting me at Università della Svizzera Italiana and Angelo for the many shared meals in Lugano.

I am grateful to the members of the AmmoRef project, whose multidisciplinary

collaboration and support were instrumental in this thesis. Working with such a talented and diverse team was an inspiring experience.

To my colleagues and friends at IIT, I owe immense gratitude: Ana, Andrea, Aparna, Axel, Davide, Dani, Dhiman, Enrico, Francesco, Giulia, Jintu, Luigi, Michele, Narjes, Nicolò, Manyi, Pedro, Peilin, Shivam, Sudip, Timothée, Umberto, and Valerio. You not only taught me so much but also made workdays enjoyable and meaningful. Sharing such diverse academic and cultural experiences broadened my perspectives in ways I could not have imagined. Special thanks to Umberto, Valerio, and Davide for warmly welcoming me to Genoa and guiding me through my first steps. My heartfelt gratitude also goes to Ana, Manyi, Shivam, and Enrico for creating such a wonderful atmosphere in our office; a perfect blend of care, energy, humor, and relaxation! The memories of dinners and evenings exploring Genoa with you are among the treasures I will carry from this PhD. I am also grateful to my friends and university classmates, who have always supported me despite the distance.

To my parents, Paola and Pier, and my brother Andrea, thank you for your truthful love and support. This thesis marks the conclusion of a long academic journey that started years ago with your encouragement. Thank you for nurturing my curiosity; it has brought me here. Realizing the childhood dream of becoming a scientist is as much your success as mine. My gratitude extends to my entire family, including relatives by choice, for always supporting me.

Lastly, to the love of my life, Caterina, words cannot express my gratitude. Your unwavering belief in me, positivity, and ability to find beauty in every challenge have been my anchor. This thesis is dedicated to you, for walking every step of this journey with me.