

THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Palladium Under Pressure:
Multiscale Modeling of Nanoalloys for Hydrogen Sensing

PERNILLA EKBORG-TANNER

Department of Physics
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2025

Palladium Under Pressure:
Multiscale Modeling of Nanoalloys for Hydrogen Sensing
PERNILLA EKBORG-TANNER
ISBN 978-91-8103-350-2

Acknowledgements, dedications, and similar personal statements in this thesis, reflect the author's own views.

© Pernilla Ekborg-Tanner, 2025.

Doktorsavhandlingar vid Chalmers tekniska högskola
Ny serie nr 5807
ISSN 0346-718X

Department of Physics
Chalmers University of Technology
SE-412 96 Göteborg, Sweden
Telephone +46 (0)31 772 10 00

Cover: Illustration of a Pd nanodisk array in a hydrogen environment but also the entire hydrogen economy.

Chalmers digitaltryck
Göteborg, Sweden 2025

Palladium Under Pressure: Multiscale Modeling of Nanoalloys for Hydrogen Sensing

PERNILLA EKBORG-TANNER
Department of Physics
Chalmers University of Technology

Abstract

Pd nanoalloys offer a spark-free, plasmonic, and highly tunable platform for H_2 sensing, which is a crucial aspect of a safe hydrogen economy. Owing to a fortunate combination of thermodynamic and optical properties, Pd nanoparticles rapidly absorb H_2 , which induces a measurable optical shift. Pure Pd sensors suffer, however, from issues such as hysteresis during hydrogenation and CO poisoning. Alloying with Au and Cu have been found to mitigate these issues, but introduces challenges related to long-term stability and performance.

The vast configurational space accessible through alloying and nanostructuring makes computational modeling an efficient route to understand and optimize nanoalloys for H_2 sensing. This thesis develops multiscale models to better understand the optical and thermodynamic properties of Pd nanoalloys for H_2 sensing, anchored at the atomic scale via first-principles calculations.

The effect of alloying Pd with Au and Cu on the surface composition and adsorbate coverages under different environments is studied via cluster expansion models. It is found that Pd segregates to the surface in H_2 and CO environments, due to strong adsorption, while Au segregates to the surface under vacuum conditions. Cu shows a more complex behavior, with a non-trivial preference for the subsurface layer under most conditions, and only modest presence in the top surface layer. The H–CO coadsorption behavior is primarily governed by the fabrication conditions, dictating whether Pd or Au segregates to the surface, while tuning the exact bulk alloy composition has only a minor effect. The experimentally observed role of Cu in mitigating CO poisoning must therefore go beyond adsorption thermodynamics, potentially by providing energetically feasible H absorption paths through the surface when the energetically most favorable paths are blocked by CO.

The H sensitivity of nanodisk devices is optimized by combining atomic-scale dielectric functions with continuum electrodynamic simulations of nanoalloy structures. Single disk simulations suggest that the H-induced plasmon shift is limited by an interplay between localized surface plasmonic resonances and interband transitions. In addition, a computational platform for designing optimal nanoarray-based sensors for specific targets is presented, paving the way for future efforts in multiplexed sensor design.

Keywords: Hydrogen sensing, Palladium, nanoalloys, nanoplasmonics, surface segregation, CO poisoning, surface phase diagram, coadsorption, Bayesian optimization

LIST OF PUBLICATIONS

This thesis is partly based on the author's licentiate thesis:

Pd-nanoalloys for hydrogen sensing:

Multiscale modeling of thermodynamic and optical properties

Pernilla Ekborg-Tanner (2022)

It consists of an introductory text and the following papers:

- I Construction and Sampling of Alloy Cluster Expansions — A Tutorial**
Pernilla Ekborg-Tanner, Petter Rosander, Erik Fransson, and Paul Erhart
PRX Energy **3**, 042001 (2024)
- II Hydrogen-driven Surface Segregation in Pd-alloys from Atomic Scale Simulations**
Pernilla Ekborg-Tanner and Paul Erhart
The Journal of Physical Chemistry C **125**, 17248–17260 (2021)
- III Competing Adsorption of H and CO on Pd-alloy Surfaces:
Mechanistic Insight into the Mitigating Effect of Cu on CO Poisoning**
Pernilla Ekborg-Tanner and Paul Erhart
In manuscript
- IV Computational Design of Alloy Nanostructures for Optical Sensing of Hydrogen**
Pernilla Ekborg-Tanner, J. Magnus Rahm, Victor Rosendal, Maria Bancerek,
Tuomas P. Rossi, Tomasz J. Antosiewicz, and Paul Erhart
ACS Applied Nano Materials **5**, 10225, (2022)
- V Bayesian Optimization of PdAu Metasurfaces for Hydrogen Sensing:
From First-principles Calculations to Device**
Pernilla Ekborg-Tanner, Athanasios Theodoridis, Joachim Fritzsche,
Christoph Langhammer, Andrea Baldi, and Paul Erhart
In manuscript

AUTHOR CONTRIBUTION

The author's contribution to the papers:

- I I developed the second part of the tutorial including all the calculations, analysis and presentation, assisted on the other parts and wrote the paper together with my co-authors.
- II I carried out the calculations and simulations, analyzed the results and wrote the paper with assistance from my co-author.
- III I constructed, sampled and analyzed the cluster expansions, assisted in the machine-learned interatomic potential construction, analyzed the results and wrote the paper with assistance from my co-author.
- IV I carried out the finite-difference time-domain simulations, analyzed the corresponding results and wrote the majority of the paper with assistance from my co-authors.
- V I developed the Bayesian optimization framework, carried out the optimization and all simulations, selected samples for the experimental verification, analyzed the results and wrote the paper with assistance from my co-authors.

PUBLICATIONS NOT INCLUDED IN THESIS

The following publications have been written during the doctoral studies but are outside the scope of this thesis:

Machine Learning-Based Interpretation of Optical Properties of Colloidal Gold with Convolutional Neural Networks

Frida Bilén, Pernilla Ekborg-Tanner, Antoine Balzano, Michaël Ughetto, Robson Rosa da Silva, Hannes Schomaker, Paul Erhart, Kasper Moth-Poulsen, and Romain Bordes
The Journal of Physical Chemistry C **128**, 33, 13909–13916 (2024)

Accelerating Plasmonic Hydrogen Sensors for Inert Gas Environments by Transformer-Based Deep Learning

Viktor Martvall, Henrik Klein Moberg, Athanasios Theodoridis, David Tomeček, Pernilla Ekborg-Tanner, Sara Nilsson, Giovanni Volpe, Paul Erhart, and Christoph Langhammer
ACS Sensors **10**, 1, 376–386 (2025)

Contents

List of abbreviations	xiii
1 Introduction	1
1.1 Pd-based nanoplasmonic hydrogen sensing	2
1.2 Alloying to improve material properties	2
1.3 Objective and structure of thesis	4
I Thermodynamics of Pd-alloy surfaces	5
2 Thermodynamical principles for bulk and surfaces	7
2.1 Bulk thermodynamics	8
2.1.1 The Pd–H system	10
2.1.2 The Pd–Au system	10
2.1.3 The Pd–Cu system	11
2.2 Surface thermodynamics	11
2.3 Gas–surface interactions	12
3 Atomistic modeling	15
3.1 Density functional theory	16
3.1.1 The Hohenberg–Kohn theorems	16
3.1.2 The Kohn–Sham ansatz	17
3.1.3 Solving the Kohn–Sham equations	18
3.1.4 The exchange–correlation functional	20
3.2 Machine-learned interatomic potentials	21
3.2.1 Message passing atomic cluster expansion	21
3.3 Cluster expansions	22
3.3.1 Definitions and formalism	23
3.3.2 Cluster expansion construction	24
3.3.3 Improving cluster expansions with physical intuition	25
3.3.4 Training data generation	26

3.4	Monte Carlo simulations	27
3.4.1	The canonical ensemble	28
3.4.2	The semi-grand canonical ensemble	28
4	Segregation and adsorption on Pd-alloy surfaces	31
4.1	Adsorption and segregation in simple systems	31
4.1.1	From chemical potential to pressure	33
4.2	Surface segregation of Pd-alloys in H_2	34
4.3	Competing H and CO adsorption on Pd-alloys	36
II	Optical properties of Pd nanoalloys	39
5	Plasmonic resonances of metal nanoparticles	41
5.1	The dielectric function	41
5.2	Localized surface plasmon resonances	44
5.2.1	The quasistatic approximation	44
5.2.2	The modified long-wavelength approximation	46
5.3	Surface lattice resonances	47
6	Electrodynamic simulations	49
6.1	Maxwell's equations	50
6.2	Finite-difference time-domain method	50
7	Bayesian optimization	53
7.1	Bayes' theorem	54
7.2	Gaussian processes	55
7.3	Maximizing the acquisition function	56
8	Optical properties of Pd nanoalloys for H_2 sensing	57
8.1	Dielectric functions of the Pd–Au–H system	57
8.2	H_2 sensitivity in Pd-alloy single disks	58
8.3	Optimized nanoalloy arrays for H_2 sensing	60
III	Final remarks	63
9	Conclusions	65
9.1	Main findings	65
9.1.1	Thermodynamics of Pd-alloy surfaces	65
9.1.2	Optical properties of Pd nanoalloys	66
9.2	Limitations	68

9.3 Outlook	69
Acknowledgments	71
Bibliography	73
Papers I–V	85

List of abbreviations

- BO** Bayesian optimization. vi, 3, 49, 53–55, 60, 61, 66, 67
- CE** cluster expansion. vi, 3, 11, 15, 16, 22–28, 33, 65, 68
- DF** dielectric function. 41–45, 49, 50, 52, 57–61, 66–69
- DFT** density-functional theory. 3, 8, 15–17, 19, 21–23, 26, 32, 34, 43, 49, 53, 65, 68
- ECI** effective cluster interaction. 24–26
- FCC** face-centered cubic. 25, 32
- FDTD** finite-difference time-domain. vi, 3, 49–53, 58, 61, 66, 68, 69
- FoM** figure of merit. 61, 62, 67
- FWHM** full width at half maximum. 61, 67
- GGA** generalized-gradient approximation. 20
- GP** Gaussian process. 54–56
- KS** Kohn-Sham. 17–19
- LDA** local density approximation. 20
- LSPR** localized surface plasmon resonance. 2, 41, 43–48, 60, 67
- MACE** message passing atomic cluster expansion. 16, 21, 22, 32, 38
- MC** Monte Carlo. 3, 8, 15, 16, 21, 27–29, 33, 34, 36, 53
- ML** machine learning. 16, 21, 26

- MLIP** machine-learned interatomic potential. vi, 3, 15, 16, 21, 22, 68, 69
- MLWA** modified long-wavelength approximation. 46–48
- NEB** nudged elastic band. 32, 38, 69
- NP** nanoparticle. 2, 11, 41, 44–49
- PAW** projector augmented wave. 19, 20
- RA** Rayleigh anomaly. 47, 48
- RBF** radial basis function. 55
- SGC** semi-grand canonical. 28, 29
- SLR** surface lattice resonance. 41, 47, 48, 60, 67
- SPD** surface phase diagram. 14, 36–38
- TDDFT** time-dependent density-functional theory. 3, 43, 49, 50, 57, 58, 61
- UHV** ultra-high vacuum. 35, 36, 66
- XC** exchange-correlation. 18, 20

Introduction

Here we stand or here we fall, history won't care at all

Queen

The hydrogen (H) economy envisions an energy system with H_2 as the primary energy carrier, enabling the transition toward a carbon-free energy system. H_2 possesses the highest gravimetric energy density of any fuel and can be produced from renewable sources [1], making it an attractive alternative to fossil fuels. Fuel cells convert the chemical energy of H_2 into electricity and operate efficiently across a wide range of power scales, from portable units on the order of 1 W to stationary plants approaching 1 MW [2]. In recent years, H_2 -powered fuel cells have gained attention for heavy vehicles (~ 100 kW) such as trucks [3], buses [4], and trains [5], where battery-powered electric solutions are limited by energy density and material abundance. Despite these advances, widespread adoption of the hydrogen economy remains constrained by challenges in efficient production [1, 6], compact storage [1, 7], and detection of leaks [8].

This thesis is related to the safety issues arising from the high flammability of H_2 mixed with air. Because leaks cannot be entirely prevented, reliable H detection is essential for the safe handling and storage of H_2 . H_2 sensing is an active research field with approaches that exploit changes in, for instance, the optical [9–21] or electrical [22–24] properties of a material upon interaction with H, leading to measurable shifts in optical spectra or electrical resistance, that can be correlated to the H_2 partial pressure of the surrounding environment. Specifically, this thesis concerns plasmonic H_2 sensing based on Pd-alloy nanostructures. This sensing concept was first demonstrated by Langhammer *et al.* in 2007 [11] and has since been widely explored [12–20, 25–27]. However, remaining issues related to, for instance, poor kinetics, long-term stability and cross-sensitivity to other species motivates further investigations.

1.1 Pd-based nanoplasmonic hydrogen sensing

When a Pd nanoparticle (NP) is subjected to an electric field, the conduction electrons can start to collectively oscillate, forming a localized surface plasmon resonance (LSPR) [28, 29]. Here, localized refers to the confinement of the plasmon to the nanoscaled particle, in contrast to bulk plasmons in extended materials. The excitation of this resonance produces a characteristic peak in the optical extinction spectrum at the resonance frequency. Upon H absorption, the electronic structure of Pd changes, causing a shift in the plasmon resonance, typically toward lower frequencies. Pd is furthermore well known for its ability to reversibly absorb large amounts of H [11, 30]. The combination of this fast H uptake with the H-induced optical shift forms the basis of Pd-based nanoplasmonic H₂ sensing.

By monitoring the plasmon peak position over time and calibrating it against known H₂ pressures, the sensor can quantify the surrounding H₂ concentration. Formally, the H sensitivity can be defined as [31]

$$S = \frac{df_{\text{opt}}}{dp_{\text{H}_2}} = \frac{df_{\text{opt}}}{dc_{\text{H}}} \times \frac{dc_{\text{H}}}{dp_{\text{H}_2}} = S_{\text{opt}} \times S_{\text{th}}, \quad (1.1)$$

where f_{opt} is the optical readout (for example, the plasmonic peak position), p_{H_2} is the H₂ partial pressure and c_{H} is the absorbed H concentration. The separate optical S_{opt} and thermodynamic S_{th} contribution to the sensitivity demonstrates that a comprehensive understanding of Pd-based H₂ sensors requires consideration of both their thermodynamic and optical properties.

1.2 Alloying to improve material properties

Alloying is a key strategy for tuning material properties which has proven effective in Pd-based nanoplasmonic H₂ sensing. Pure Pd NPs exhibit hysteresis in their H absorption–desorption isotherms due to a miscibility gap between the metal α - and hydride β -phases, preventing accurate pressure determination within this region. Alloying with, for instance, Au can suppress this phase separation and eliminate the hysteresis [15, 27, 32].

Pd NPs also suffer from CO poisoning, where adsorbed CO molecules block surface sites and hinder H uptake. Two main strategies have been suggested to mitigate this effect. The first involves coating the nanostructures with selective barrier materials that permit H₂ diffusion while blocking larger molecules [18, 33, 34]. The second, less obvious strategy, relies on alloying with Cu. At Cu concentrations as low as 5 %, significant reduction of CO poisoning has been observed [35, 36]. The underlying mechanism of the mitigating effect of Cu on CO is, however, not fully understood.

Reliable and long-term stable H₂ sensors require minimization of both hysteresis and CO poisoning. Recent experiments show that alloy compositions around 25 % Au and

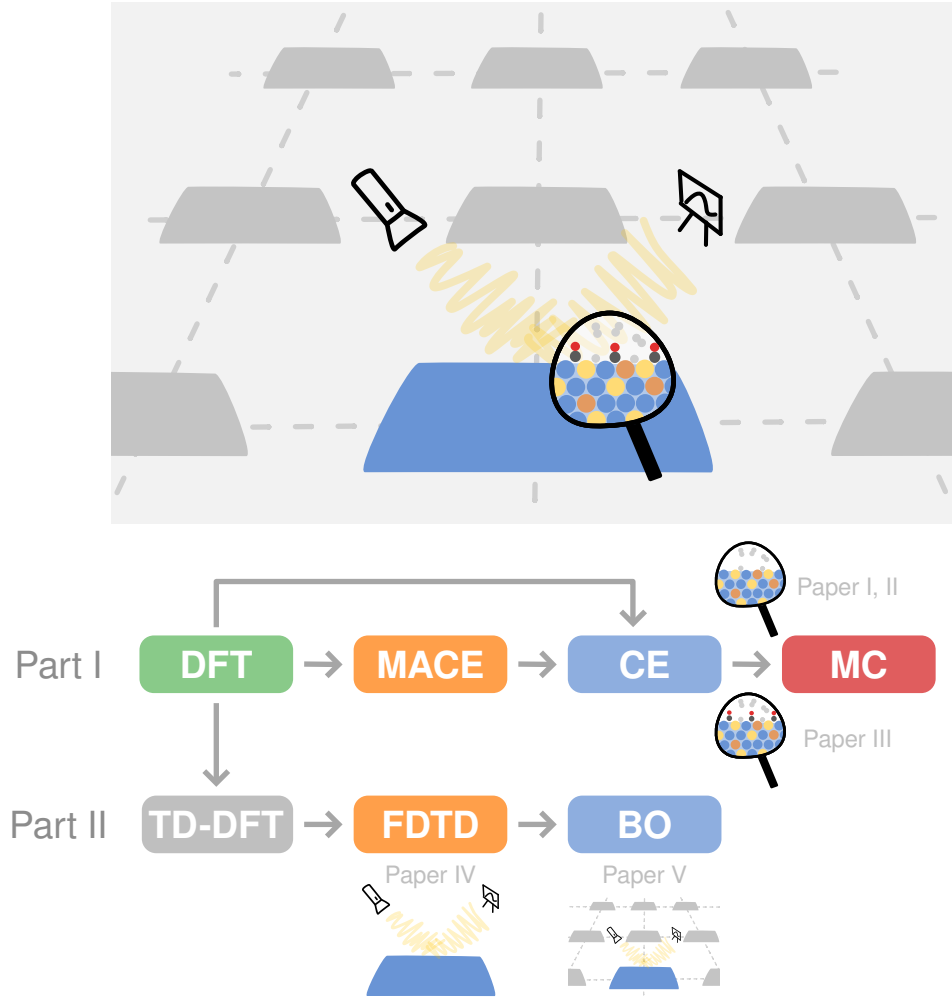


Figure 1.1: Overview of the system studied in this thesis and the employed modeling framework. The first part of the thesis examines the surface state of alloy–adsorbate systems via a multiscale approach starting from first-principles density-functional theory (DFT) calculations, via machine-learned interatomic potentials (MLIPs) and cluster expansions (CEs), to Monte Carlo (MC) simulations of thermodynamic averages. The second part studies the optical response of nanoalloys via time-dependent density-functional theory (TDDFT)-informed finite-difference time-domain (FDTD) simulations, and utilizes Bayesian optimization (BO) to optimize the nanostructures with respect to geometry and composition.

10 % Cu simultaneously address both issues [35, 36]. At the same time, alloying increases the compositional complexity and makes prediction of various properties more challenging. Atomic-scale chemical ordering can introduce effects such as phase separation and surface segregation. The equilibrium surface state is, in general, environment-dependent, meaning that the conditions during fabrication, operation, and storage can affect the alloy composition of the surface region and, as a consequence, the adsorption tendency of different species. These phenomena affect the ability to absorb H, and in turn the performance of devices. Therefore, knowledge about the segregation and adsorption behavior is key for making informed design decisions.

1.3 Objective and structure of thesis

Alloying and nanoengineering enable systematic tuning of thermodynamic and optical properties by varying geometry, composition, and alloying elements, giving rise to a vast configuration space. Computational modeling provides an efficient route to explore this space, reducing the time and cost required for experimental screening. The objective of this thesis is to develop and apply multiscale modeling approaches to address key challenges in the design, fabrication, and long-term stability of Pd-based hydrogen sensors. These approaches bridge multiple length scales, from first-principles calculations of electronic properties to thermodynamic modeling of the atomic configuration or continuum descriptions of the optical response at the mesoscale (Fig. 1.1). Specifically, the aim of this thesis is to answer the following questions:

- What is the surface segregation behavior of Pd-alloys under realistic conditions?
- How does alloying with Cu mitigate CO poisoning?
- How can computational modeling be used to design optimal H sensors?

Although the thermodynamic and optical aspects of Pd-based nanoplasmonic sensing are closely linked, they are investigated using different methodological approaches (Fig. 1.1). For this reason, this thesis is organized into two main parts: Part I, which addresses thermodynamics, and Part II, which focuses on optical properties. Each part comprises three main sections: a review of the theoretical background, a description of the computational methods (spanning two chapters in Part II), and a presentation of the principal results. Finally, Part III summarizes the main findings and limitations, and outlines directions for future work.

Part I

Thermodynamics of Pd-alloy surfaces

Thermodynamical principles for bulk and surfaces

Insanity laughs, under pressure we're breaking

Queen

Pd is a core component in many H technologies due to its unique ability to quickly absorb and desorb large quantities of H, in relation to the amount of H₂ in the immediate environment, at ambient conditions. This is possible due to the thermodynamics of the Pd–H system. Alloying is a common strategy for material optimization in various applications as it introduces tunability of the material properties. This strategy requires knowledge about how the introduction of additional species affects the thermodynamics of the system. In this section, the thermodynamics of hydrogenated Pd-alloys are discussed with emphasis on the effects of alloying on surface segregation and the interaction with gas environments.

According to the laws of thermodynamics, a system in constant temperature and pressure environments strives towards the state with lowest Gibbs free energy [37],

$$G = U + pV - TS. \quad (2.1)$$

The Gibbs free energy is an example of a thermodynamic potential, and consists of the internal energy U of the material as well as the work pV associated with taking up space corresponding to its volume V and the entropy contribution $-TS$ corresponding to the heat that can be extracted from its environment of temperature T . Other common thermodynamic potentials, relevant for other conditions, include the enthalpy

$$H = U + pV,$$

and Helmholtz free energy

$$F = U - TS,$$

as well as the grand potential, closely related to F , presented in Sect. 3.4.1.

For many solids, including the alloy surfaces considered in this thesis, the change in the pV term between two relaxed atomic configurations is small compared to the other contributions and can be neglected [38]. In such cases, the Gibbs and Helmholtz free energies can be used interchangeably. The Helmholtz free energy is dominated by the internal electronic energy E_{el} , the vibrational free energy F_{vib} and the configurational entropy associated with the atomic configuration S_{conf} ,

$$F(T, V) \simeq E_{\text{el}} + F_{\text{vib}} - TS_{\text{conf}},$$

for non-magnetic materials, neglecting electronic excitations.

In ab-initio thermodynamics, the energy contributions are obtained from ab-initio (first-principles) calculations, often based on density-functional theory (DFT) [38]. DFT calculations are typically used to determine the zero-temperature and zero-pressure ground state electronic energy at a fixed volume, often corresponding to the relaxed volume V_0 obtained via structural relaxation. The energy of this state is

$$F(0, V_0) = E_{\text{DFT}} + E_{\text{ZPE}},$$

where E_{DFT} is the DFT electronic energy and E_{ZPE} the zero-point vibrational energy.

At finite temperatures, the vibrational free energy $F_{\text{vib}}(T)$ (including E_{ZPE}) and configurational entropy S_{conf} contribute to the free energy,

$$F(T, V_0) \simeq E_{\text{DFT}} + F_{\text{vib}}(T) - TS_{\text{conf}}.$$

neglecting any temperature dependence of the electronic energy. In most cases, the absolute free energy is not of importance but rather the difference in free energy between two states. For the Pd-based alloys considered in this thesis, the change in vibrational energy between atomic configurations can typically be neglected [39], leaving only the DFT energy and the configurational entropy which can be obtained via Monte Carlo (MC) simulations (Sect. 3.4). In the following, it is demonstrated how the Gibbs free energy dictates the bulk, surface and adsorption properties of a system.

2.1 Bulk thermodynamics

In a bulk multicomponent system, the average atomic configuration will be the one minimizing the Gibbs free energy. For a binary system $A_{1-x}B_x$, the free energy is

$$G(T, p, x) = (1 - x)G_A + xG_B + \Delta G_{\text{mix}}$$

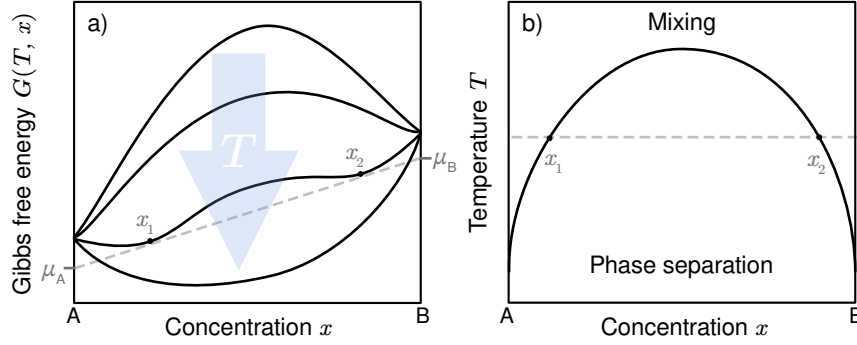


Figure 2.1: (a) Schematic illustration of the Gibbs free energy for a mixed system consisting of atoms A and B. If any part of the Gibbs free energy curve is concave-down, the system will separate into two phases with compositions x_1 , x_2 in this region following the tangent line below the curve. (b) Phase diagram obtained by following $x_1(T)$ and $x_2(T)$.

where G_X are the free energies of the separated systems and

$$\Delta G_{\text{mix}} = \Delta H_{\text{mix}} - T \Delta S_{\text{mix}}$$

is the change in free energy upon mixing [37, Chapter 5]. If ΔG_{mix} is negative, the system is miscible. For a regular solution, with random mixing and a single parameter Ω for interatomic interactions, the changes in enthalpy and entropy are given by

$$\begin{aligned} \Delta S_{\text{mix}} &= -Nk_B [x \ln x + (1 - x) \ln (1 - x)] \\ \Delta H_{\text{mix}} &= \Omega x (1 - x), \end{aligned}$$

where N is the number of atoms and k_B is the Boltzmann constant. The entropy contribution will always favor mixing, while the enthalpy favors mixing if $\Omega < 0$ and phase separation otherwise.

Figure 2.1a shows the Gibbs free energy of a regular solution with $\Omega < 0$ as a function of the concentration x_B for different temperatures. At high temperatures, the curve is concave-up everywhere ($\Delta G_{\text{mix}} < 0$) which indicates that the system is miscible over the entire concentration range. As the temperature decreases, the curve will form a concave-down region in the middle, in which the system will separate into phases with concentrations x_1 and x_2 corresponding to the interception with a tangent below the curve. As temperature decreases further, the interception points move toward the edges, until eventually the energetically favorable state is separation into the two pure A and B phases. By tracking the interception points, a phase diagram can be constructed (Fig. 2.1b). In the following, the bulk thermodynamic properties of the Pd–Au–Cu–H system are discussed.

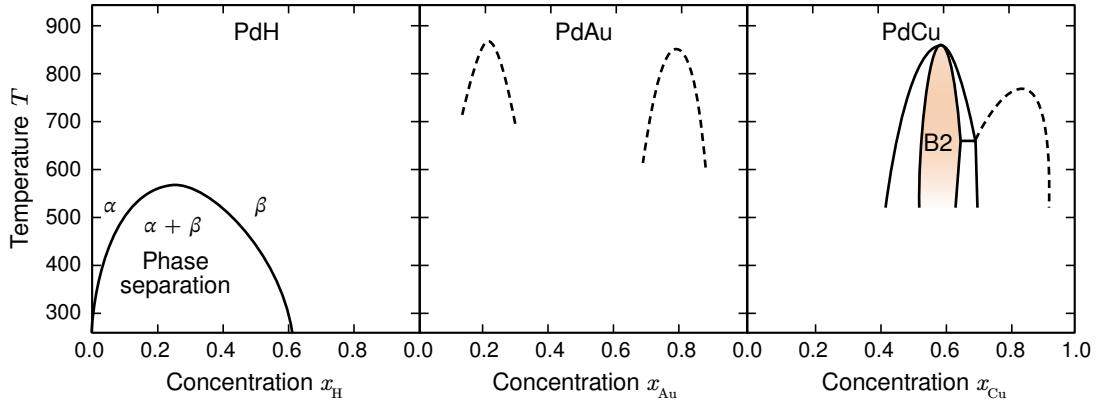


Figure 2.2: Phase diagrams for Pd–H [30], Pd–Au [40], and Pd–Cu [41]. Pd–H exhibits two phases with the same crystal structure, α and β , with phase separation at temperatures below approximately 600 K. Pd–Au is miscible in the entire region with a few ordered structures reported, indicated by the dashed lines. Pd–Cu is mostly miscible but have several ordered phases in the Cu-rich end, causing phase separation (white areas next to the orange area).

2.1.1 The Pd–H system

H absorption is associated with an expansion of the Pd lattice. At moderate temperatures (≤ 600 K), this causes phase separation between the H-poor α -phase and the H-rich β -phase as the H content increases (Fig. 2.2). The relation between the H uptake and the H_2 pressure is associated with a jump from low to high H content, at the so-called plateau pressure, and hysteresis caused by the energy barriers of forming β in α and vice versa [30] (Fig. 2.3). These phenomena are an issue in the context of H_2 sensing, since the H content is not a well-defined function of H_2 pressure in this region.

2.1.2 The Pd–Au system

PdAu alloys are miscible over the entire composition range (Fig. 2.2). Some ordered phases have been observed, but these do not cause significant phase separation [40]. Alloying Pd with Au is often beneficial in H applications since the miscibility gap and hysteresis associated with the Pd–H system vanish at approximately 20 % Au [27, 32, 42] (Fig. 2.3). This effect is attributed to the fact that Au atoms are larger than Pd atoms, as indicated by their respective lattice parameters 4.1 Å and 3.9 Å (Fig. 4.1c), causing a lattice expansion similar to the expansion needed to absorb large quantities of H [43, Chapter 3]. At the same time, the H solubility decreases with Au content [32].

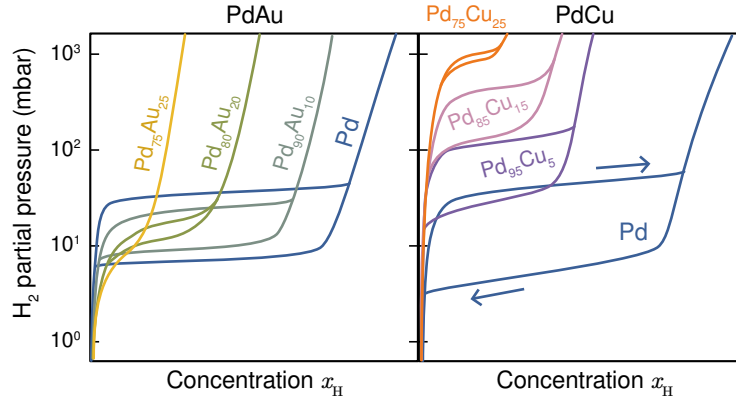


Figure 2.3: Hysteresis in PdAu [27] and PdCu [35] NPs during H₂ absorption and desorption at ambient temperature.

2.1.3 The Pd–Cu system

PdCu alloys are completely miscible for compositions with more than 60 % Pd (Fig. 2.2), which are the primary focus of this thesis. For higher Cu content, however, several ordered phases can be found which cause phase separation. This includes the B2-phase close to 50 % Cu, which exhibits a different crystal structure (CsCl structure) than the solid solution. This difference in crystal structure causes difficulties in modeling, in particular using lattice-based methods such as cluster expansion (CE). Since the Cu atom is smaller than Pd, with lattice parameter 3.6 Å compared to 3.9 Å (Fig. 4.1c), alloying with Cu does not have the same benefits on the H absorption thermodynamics as alloying with Au. Similar to alloying with Au, the H solubility decreases with increasing Cu content [44], but at the same time the plateau pressure increases significantly (Fig. 2.3) limiting the potential for H₂ sensing.

2.2 Surface thermodynamics

The energy landscape of an atom close to a surface differs from an atom in the bulk due to the change in coordination number, surface relaxation, and potential interaction with adsorbates. The energy difference gives rise to the concept of surface free energy,

$$\gamma(T, p, N) = \frac{1}{2A} [G^{\text{slab}}(T, p, N) - N g^{\text{bulk}}(T, p)] \quad (2.2)$$

where N is the number of atoms in the surface slab and g^{bulk} is the energy per atom for the corresponding bulk structure. The surface energy is the energy per area required to form a surface by making a cut in a bulk material. Surface orientations with lower energy are more stable, and therefore more likely to naturally occur for surface slabs and nanoparticle (NP) [45].

For an alloy, the different energy landscapes generally result in differences in the chemical ordering of atoms close to the surface compared to the bulk. The most stable state will be the one that minimizes the surface free energy which for a binary alloy $A_{1-x}B_x$ (where x is the bulk composition) reads [46]

$$\gamma(T, p, N, x, \Gamma) = \frac{1}{2A} \left[G^{\text{slab}}(T, p, N, x, \Gamma) - N g^{\text{bulk}}(T, p, x) - \Delta\mu \Gamma \right] \quad (2.3)$$

where Γ is the surplus of B atoms in the slab compared to the corresponding bulk material and $\Delta\mu = \mu_B - \mu_A$ is the difference in the bulk-consistent chemical potential satisfying $(1-x)\mu_A + x\mu_B = g^{\text{bulk}}(x)$ (Fig. 2.1). The extension for systems with more than two components is straightforward.

The strive towards the minimal surface energy drive phenomena such as surface segregation and surface ordering. These processes strongly influence surface properties and are critical for performance in applications like catalysis and sensing, where the composition of the outermost atomic layers plays a decisive role. The surface segregation behavior of an arbitrary alloy is generally difficult to predict since it requires knowledge about the energetics of all possible atomic configurations. There are, however, some general segregation indicators that can be identified [47–50]:

- **Pure metal surface energy:** The species with the lower surface energy tends to segregate towards the surface. This is generally considered the strongest indication of surface segregation.
- **Size-mismatch:** If there is a large size difference between solute and host, the solute tends to segregate to the surface [48], in particular if it is larger [50, 51].
- **Adsorption energies** (if adsorbates are involved): The species with the lower adsorption energy tends to segregate to the surface [52].

A more quantitative approach is to calculate the segregation energy E^{seg} from dilute alloy surface slabs consisting of a single A (solute) atom in B (host),

$$E_n^{\text{seg}} = E_{A \text{ in } n} - E_{A \text{ in bulk}}, \quad (2.4)$$

where the energies on the right-hand side are calculated for a single A atom in layer n and for a single A atom in bulk.

2.3 Gas–surface interactions

If a surface is placed in a gas environment, the gas species X can adsorb with an associated adsorption energy

$$E_X^{\text{ads}} = E_{\text{slab:X}} - E_{\text{slab}} - E_X^{\text{gas}}, \quad (2.5)$$

where $E_{\text{slab:X}}$ is the energy of the slab with the adsorbed species, E_{slab} is the energy of the slab alone, and E_X^{gas} is the energy of the species in the gas phase. Similar to the previous section (Eq. (2.2)), the stable (average) state is the one that minimizes the surface free energy [38]

$$\gamma(T, p, N, N_X) = \frac{1}{2A} \left[G^{\text{slab}}(T, p, N, N_X) - N g^{\text{bulk}}(T, p) - N_X \mu_X(T, p) \right] \quad (2.6)$$

where N_X is the number of adsorbates with chemical potential μ_X . This expression can be readily extended for alloy surfaces according to Eq. (2.3), multiple adsorbed species and H absorption¹.

According to the ideal gas law, the chemical potential of a species in the gas phase is

$$\mu(T, p) = g(T, p) = \mu^0(T, p^0) + k_B T \ln \frac{p}{p^0} \quad (2.7)$$

given in relation to a reference pressure p^0 , where μ^0 is the temperature-dependent reference chemical potential

$$\mu^0(T, p^0) = \mu^0(0, p^0) + g^0(T) = E_{\text{DFT}} + E_{\text{ZPE}} + h^0(T) - Ts^0(T). \quad (2.8)$$

Here, E_{DFT} and E_{ZPE} are the ground state electronic and zero-point energy of the species in the gas phase. The temperature-dependence is captured by the free energy of the reference state, $g^0(T) = h^0(T) - Ts^0(T)$, which can be calculated or taken from thermodynamic tables based on experimental data [53].

The treatment of the vibrational free energy should be reconsidered in the context of modeling open surface–gas systems. Since the number of gas species can vary (Eq. (2.6)), relying on cancellation between configurations is no longer appropriate. Instead, a lack of treatment needs to be motivated by overall small vibrational contributions of the adsorbate. If the vibrational contribution of the adsorbate can be approximated as independent of the configuration, it can be treated as a correction to G^{slab} that is incorporated in μ_X . This will be discussed further in the context of pressure conversion in Sect. 4.1.1.

¹This thesis is limited to studies of surface adsorption and H absorption is thus not taken into account.

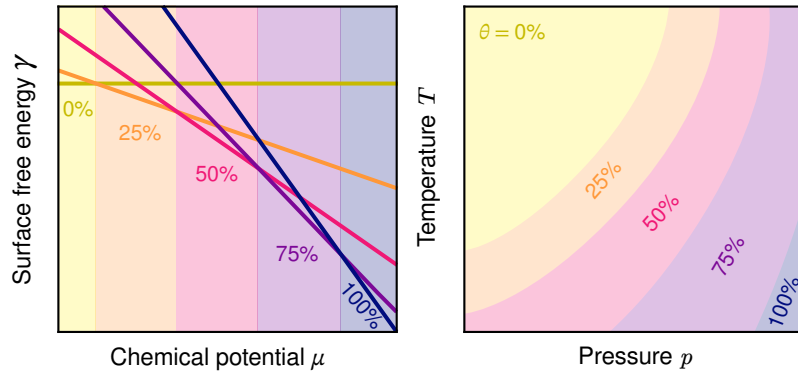


Figure 2.4: Constructing a SPD from the surface free energy as a function of the adsorbate chemical potential. The adsorbate coverage θ is indicated by the colors.

If the surface free energy (Eq. (2.6)) as a function of adsorbate chemical potential is known, one can construct a surface phase diagram (SPD) that indicates the adsorbate coverage at a given chemical potential, or, if an appropriate reference chemical potential (Eq. (2.8)) can be established, as a function of temperature and pressure (Fig. 2.4).

Atomistic modeling

A kind of magic

Queen

From the previous chapter, it is clear that accurately and efficiently describing the energetics of the alloy–adsorbate system is essential for studying its thermodynamics. An atomistic modeling framework for the quantum-mechanical properties of interacting many-body systems is therefore required, which is a common objective in many fields of materials science.

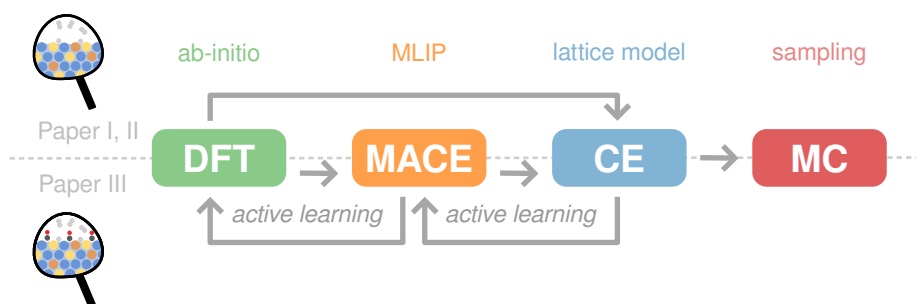


Figure 3.1: Schematic illustration of the atomistic modeling framework, starting from the ab-initio level with DFT, via MLIP and CE models to MC sampling of thermodynamic averages.

The gold standard for such calculations is density-functional theory (DFT), which provides first-principles energies and forces for systems with up to about 1000 atoms with high accuracy. However, DFT remains computationally demanding and becomes infeasible for larger systems. To address this limitation, a range of computational methods has been developed, including interatomic potentials that predict energies and forces

from atomic positions, and lattice-based models that predict the energetics of systems mapped onto an ideal lattice. With recent advances in performance and accessibility in the field of machine learning (ML), DFT-based machine-learned interatomic potentials (MLIPs) have become increasingly popular, since they approach DFT-level accuracy while being many orders of magnitude faster (see comparison in Paper III). Even with the most efficient MLIPs, however, accessible timescales are typically limited to a few nanoseconds which is far shorter than those relevant for phase separation, surface segregation, and H absorption. In such cases, lattice-based Monte Carlo (MC) simulations are better suited as they circumvent kinetic barriers by only considering the energy difference between two configurations. For the alloy–adsorbate systems considered in this thesis, lattice-based models, such as cluster expansions (CEs), remain the best option since the adsorbates can be mapped to an adsorbate–vacancy lattice. In contrast, defining meaningful and sampling-efficient MC trial moves for adsorbates within interatomic potential frameworks is not straightforward.

Figure 3.1 summarizes the atomistic modeling approaches used in this thesis. In Paper I and Paper II, CEs are constructed from DFT training data for binary {111}-alloy surfaces with adsorbed H to study H-driven surface segregation via MC simulations. In Paper III, this framework is extended to ternary alloys and multiple surface orientations with H and CO adsorbed to investigate coadsorption. This expansion in scope requires a substantially larger training set (by 1 to 2 orders of magnitude), which is enabled by developing a MACE MLIP to serve as DFT-emulator. The following sections describe the computational methods in more detail.

3.1 Density functional theory

Density-functional theory (DFT) is a widely used method for calculating the quantum-mechanical ground state of many-body systems containing up to thousands of atoms. In this thesis, the Vienna Ab initio Simulation Package (VASP) [54, 55] is used for all DFT calculations. This section provides a brief overview of the theoretical framework of DFT, following Ref. [56].

3.1.1 The Hohenberg–Kohn theorems

Consider a many-body system consisting of electrons and nuclei. Because nuclei are much heavier than electrons, their characteristic timescales of motion are much longer. From the perspective of the nuclei, the electrons appear to adjust instantaneously to nuclear motion, whereas the electrons perceive the nuclei as fixed. This separation of timescales is the basis of the Born–Oppenheimer approximation [57], under which the wavefunctions of the electrons and the nuclei are separated. The nuclei are then treated

as stationary when solving for the electronic ground state, and their contributions are typically handled separately using classical methods.

The electronic Hamiltonian for a system with fixed nuclei, expressed in Hartree atomic units ($\hbar = m_e = e = 4\pi\epsilon_0 = 1$), is [56, Chapter 6]

$$\hat{H} = -\frac{1}{2} \sum_i \nabla_i^2 + \sum_i V_{\text{ext}}(\mathbf{r}_i) + \frac{1}{2} \sum_{i \neq j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}, \quad (3.1)$$

where the first term is the kinetic energy, the second represents the external potential $V_{\text{ext}}(\mathbf{r})$ (including electron–nucleus interactions), and the last term is the electron–electron Coulomb interaction. For most systems, solving the corresponding Schrödinger equation directly is not feasible, but the problem can be reformulated via DFT.

DFT is based on the Hohenberg–Kohn theorems [58]. The first theorem states that the external potential $V_{\text{ext}}(\mathbf{r})$, and in turn all properties of the many-body system, are uniquely determined by the ground state electron density $n_0(\mathbf{r})$. The second theorem states that a total energy functional $E[n]$ of the electron density $n(\mathbf{r})$ can be defined for any external potential, and that the density minimizing this functional is the exact ground state density $n_0(\mathbf{r})$. Consequently, all ground state properties can, in principle, be obtained by finding and minimizing the functional $E[n]$.

3.1.2 The Kohn–Sham ansatz

A general expression for the total energy functional of an interacting many-body system is [56, Chapter 6]

$$E_{\text{HK}}[n] = T[n] + E_{\text{int}}[n] + \int d\mathbf{r} V_{\text{ext}}(\mathbf{r})n(\mathbf{r}) + E_{II}, \quad (3.2)$$

where T and E_{int} are the electronic kinetic and potential energies, respectively, and E_{II} is the ion–ion interaction energy.

To simplify the problem, the Kohn–Sham (KS) ansatz [59] can be used, which replaces the many-body problem by an auxiliary non-interacting problem of independent electrons. The underlying assumption is that the independent-particle problem can be defined such that it reproduces the ground state density of the interacting system. Rewriting the functional Eq. (3.2) in terms of the non-interacting system yields [56, Chapter 7]

$$E_{\text{KS}}[n] = T_{\text{KS}}[n] + E_{\text{Hart}}[n] + E_{\text{xc}}[n] + \int d\mathbf{r} V_{\text{ext}}(\mathbf{r})n(\mathbf{r}) + E_{II}, \quad (3.3)$$

where T_{KS} is the non-interacting kinetic energy and $E_{\text{Hart}}[n]$ is the Hartree energy,

$$E_{\text{Hart}}[n] = \frac{1}{2} \int d\mathbf{r} d\mathbf{r}' \frac{n(\mathbf{r})n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|},$$

which represents the classical Coulomb interaction in terms of the electron density. All other many-body effects are collected in the exchange-correlation (XC) functional $E_{\text{xc}}[n]$, which encodes the difference between the interacting and non-interacting kinetic and internal energies. If E_{xc} was known, solving the auxiliary problem would yield the exact ground state energy and density.

Solving the auxiliary problem is equivalent to minimizing E_{KS} (Eq. (3.3)) under the constraint of a fixed number of electrons. This can be done using the Lagrange multiplier method and leads to the Schrödinger-like KS equations (Eq. (3.4)–(3.6)) [56, Chapter 7]

$$(\hat{H}_{\text{KS}} - \epsilon_i) \psi_i(\mathbf{r}) = 0, \quad (3.4)$$

with the effective Hamiltonian

$$\hat{H}_{\text{KS}} = -\frac{1}{2}\nabla^2 + V_{\text{KS}}(\mathbf{r}), \quad (3.5)$$

and the effective KS potential

$$V_{\text{KS}}(\mathbf{r}) = V_{\text{ext}}(\mathbf{r}) + \frac{\delta E_{\text{Hart}}}{\delta n(\mathbf{r})} + \frac{\delta E_{\text{xc}}}{\delta n(\mathbf{r})} = V_{\text{ext}}(\mathbf{r}) + V_{\text{Hart}}(\mathbf{r}) + V_{\text{xc}}(\mathbf{r}). \quad (3.6)$$

Here, $\psi_i(\mathbf{r})$ and ϵ_i are the KS wavefunctions and eigenvalues of the non-interacting system. The corresponding ground state density for a system of N electrons is

$$n_0(\mathbf{r}) = \sum_{i=1}^N |\psi_i(\mathbf{r})|^2, \quad (3.7)$$

where the sum runs over the N lowest eigenstates.

3.1.3 Solving the Kohn-Sham equations

Solving the KS equations requires the effective potential V_{KS} (Eq. (3.6)), which depends on the electron density, but the density itself is only known once the problem is solved. This circular dependence is resolved through a self-consistent procedure consisting of the following steps [56, Chapter 9]:

1. Choose an initial guess for the ground state density.
2. Calculate the effective potential (Eq. (3.6)).
3. Solve the KS equation (Eq. (3.4)).
4. Calculate the new density (Eq. (3.7)).

5. If the new density is consistent with the previous one, the problem is solved. Otherwise, repeat from step 2 with the updated density.

Since the KS equations are set up to minimize E_{KS} , the self-consistent density corresponds to the ground state. The ground state energy (as well as other properties such as forces and stresses) of the interacting system can then be calculated via the total energy functional Eq. (3.3).

In each iteration, the density is calculated based on a set of KS wave functions. For crystalline materials with periodic boundary conditions, the wave functions are typically expanded in plane waves restricted to the first Brillouin zone [56, Chapter 12]. For systems without periodic boundary conditions in all directions, such as surfaces and molecules, plane waves can still be employed by introducing sufficiently large vacuum regions in the non-periodic directions.

The wave functions have the form

$$\psi_{i,\mathbf{k}}(\mathbf{r}) = \sum_{\mathbf{G}} c_{i,\mathbf{G}}(\mathbf{k}) e^{i(\mathbf{k}+\mathbf{G})\cdot\mathbf{r}}, \quad (3.8)$$

where $c_{i,\mathbf{G}}(\mathbf{k})$ are Fourier coefficients, \mathbf{k} is a wave vector and \mathbf{G} are reciprocal lattice vectors. The effective KS potential has the periodicity of the lattice and can be expressed as a sum of Fourier components

$$V_{\text{KS}}(\mathbf{r}) = \sum_{\mathbf{G}} V_{\text{KS}}(\mathbf{G}) e^{i\mathbf{G}\cdot\mathbf{r}}. \quad (3.9)$$

Lastly, the KS equation (Eq. (3.4)) can be expressed in \mathbf{k} -space as

$$\sum_{\mathbf{G}'} H_{\mathbf{G},\mathbf{G}'}(\mathbf{k}) c_{i,\mathbf{G}'}(\mathbf{k}) = \epsilon_i(\mathbf{k}) c_{i,\mathbf{G}}(\mathbf{k}), \quad (3.10)$$

where

$$H_{\mathbf{G},\mathbf{G}'}(\mathbf{k}) = \frac{1}{2} |\mathbf{k} + \mathbf{G}|^2 \delta_{\mathbf{G},\mathbf{G}'} + V_{\text{KS}}(\mathbf{G} - \mathbf{G}'). \quad (3.11)$$

This expansion is truncated using an energy cutoff associated with the wave vector $|\mathbf{k} + \mathbf{G}|$.

A plane wave basis efficiently captures the periodicity of the system but poorly represents the rapid variations of the wave function close to the atomic cores (in contrast to, for example, an atomic orbital basis set). As a result, high energy cutoffs are required for accurate descriptions, which yields a high computational cost. To address this, pseudopotentials can be introduced, which are effective potentials representing the joint effect of nuclei and core electrons on the valence electrons [56, Chapter 11]. Pseudopotentials are constructed for each atomic species prior to the DFT calculation and yield smooth valence electron wave functions, reducing the computational cost. A related, more complex approach is the projector augmented wave (PAW) method [60, 61], which

defines a linear transformation between the all-electron wave function and the smooth valence wave function, allowing core and valence states to be treated at different levels. The PAW method is typically combined with the frozen-core approximation [62], in which core electron wave functions remain fixed during the self-consistency cycle.

3.1.4 The exchange–correlation functional

A crucial remaining step is determining the XC functional $E_{xc}[n]$. The equivalence between the total energy functional of the interacting and non-interacting systems holds only if E_{xc} is known exactly [56, Chapter 7], but in practice this functional needs to be approximated. A wide range of XC functionals has been developed, each with specific advantages and limitations, and the choice depends on the system, the properties of interest, and the available computational resources.

Since the kinetic and long-range Hartree energy contributions are treated explicitly, E_{xc} can be approximated as a local or semi-local functional of the density [56, Chapter 8]. The simplest such approximation is the local density approximation (LDA),

$$E_{xc}^{LDA}[n] = \int d\mathbf{r} n(\mathbf{r}) \epsilon_{xc}^{LDA}(n(\mathbf{r})), \quad (3.12)$$

where ϵ_{xc}^{LDA} is the exchange–correlation energy density of a homogeneous electron gas with density $n(\mathbf{r})$ [56, 59]. This approach is motivated by the fact that valence electrons of metals with nearly free electrons often behave similarly to a homogeneous electron gas.

Another family of XC functionals is based on the generalized-gradient approximation (GGA)

$$E_{xc}^{GGA}[n] = \int d\mathbf{r} n(\mathbf{r}) \epsilon_{xc}^{GGA}(n(\mathbf{r}), |\nabla n(\mathbf{r})|), \quad (3.13)$$

which introduces a dependence on the density gradient, accounting for variations in the local environment, and can thus be classified as semi-local [56, Chapter 8]. Several such functionals exist, relying on different approximations and parametrizations of ϵ_{xc}^{GGA} [63–65].

To include van der Waals interactions, a non-local XC functional is required. Such functionals belong to the family of van der Waals density functionals [66–69] and have the general form

$$E_{xc}^{vdW-DF}[n] = E_x^{GGA}[n] + E_c^{LDA}[n] + E_c^{nl}[n]. \quad (3.14)$$

with an exchange term from GGA as well as two correlation terms corresponding to LDA and non-local contributions. The latter contribution has the form

$$E_c^{nl}[n] = \frac{1}{2} \int d\mathbf{r} d\mathbf{r}' n(\mathbf{r}) K(\mathbf{r}, \mathbf{r}') n(\mathbf{r}'), \quad (3.15)$$

with a kernel function $K(\mathbf{r}, \mathbf{r}')$ describing the interaction strength between the density at two points [69, 70]. In this thesis, the vdW-DF-cx functional [68] is used due to its excellent performance for non-magnetic transition metals in general [71] and for the surface properties of Pd, Au, and Cu in particular [72].

3.2 Machine-learned interatomic potentials

An interatomic potential is a function that, given a set of atoms i with chemical species z_i and positions \mathbf{r}_i , returns the total energy E of the system. Traditional interatomic potentials rely on physics-based models and are typically computationally efficient but fail to approach DFT-level accuracy. Over the past two decades, machine-learned interatomic potentials (MLIPs) have emerged, using ML models trained on DFT data to achieve high accuracy at a greatly reduced cost [73, 74]. MLIPs provide the potential energy surface of the system, as well as the forces via differentiation, and can be used in molecular dynamics and MC simulations or as general-purpose DFT emulators.

Most MLIPs operate by transforming atomic positions and species into learnable atomic descriptors, which are then fed into an ML model that outputs the total energy. MLIPs can be broadly categorized as local or graph-based [74]. Local MLIPs use descriptors that describe the local environment of each atom and serve as inputs to models such as a linear regressors, kernel functions, or neural networks. Graph MLIPs are based on a graph representation where each atom is represented by a node and interactions by an edge, and information is exchanged via message passing between nodes. The message passing allows for learning the local environment of an atom without an explicit description. Graph MLIPs are generally slower than local ones, particularly when many-body interactions are important, which requires many message passing steps [74].

3.2.1 Message passing atomic cluster expansion

The message passing atomic cluster expansion (MACE)¹ model combines elements of both local and graph MLIPs [73–75]. The system is represented as a graph, with information exchanged between nodes through message passing, while the ACE descriptors capture many-body interactions through polynomial basis functions that efficiently encode the local environment.

In MACE, the state of node i at message-passing step t is represented by a tuple [75],

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

¹Note that atomic cluster expansion and cluster expansion are two different frameworks.

where \mathbf{r}_i and z_i denote the position and species of atom i , and $\mathbf{h}_i^{(t)}$ its learnable features. Exchanging information between nodes consists of three steps; message construction, feature update and readout.

During message construction, a message is created for each node i by summing over its neighbors j ,

$$\mathbf{M}_i^{(t)} = \sum_j \mathbf{u}_1(\sigma_i^{(t)}; \sigma_j^{(t)}) + \sum_{j_1, j_2} \mathbf{u}_2(\sigma_i^{(t)}; \sigma_{j_1}^{(t)}; \sigma_{j_2}^{(t)}) + \dots + \sum_{j_1, \dots, j_\nu} \mathbf{u}_\nu(\sigma_i^{(t)}; \sigma_{j_1}^{(t)}; \dots; \sigma_{j_\nu}^{(t)}), \quad (3.16)$$

where \mathbf{u} are learnable ACE-based message functions and ν is the maximum correlation order, effectively including interactions with body order $\nu + 1$. Next, the messages are used to update the features,

$$\mathbf{h}_i^{(t+1)} = \mathbf{U}_t(\sigma_i^{(t)}, \mathbf{M}_i^{(t)})$$

where \mathbf{U}_t is a learnable update function. This process of constructing messages and updating features can be repeated multiple times, and each repetition creates a new layer in the graph neural network. Multiple layers extend the receptive field of the nodes. With one layer, the node only sees its immediate neighbors, i.e., other nodes within a cutoff radius. With two layers, the node is affected also by the neighbors of its neighbors, and so on, effectively increasing the receptive field. After T message-passing steps, learnable readout functions \mathbf{R}_t map the node states to the total energy,

$$E = \sum_i \sum_{t=1}^T \mathbf{R}_t(\sigma_i^{(t)}).$$

Standard graph MLIPs incorporate only two-body interactions through edges. Capturing higher body order effects requires many message-passing layers to transmit information across the graph. By constructing messages with ACE descriptors, MACE directly encodes high body order interactions with a low number of layers, resulting in more efficient models [75].

3.3 Cluster expansions

For many crystalline systems, configurational energy and entropy are the most important contributions to the free energy. Exhaustively sampling all possible atomic configurations using DFT or even a MLIP is practically impossible for relevant system sizes in most cases. Instead, one has to choose between studying a few selected atomic configurations or employ models that enable efficient sampling of the configuration space, such as lattice-based CEs.

A CE is a generalized Ising model that predicts the energy as a function of the atomic configuration on a fixed lattice [76], based on reference data typically obtained from DFT calculations. In this thesis, alloy CEs [76] are used to model surface segregation and adsorption behavior of Pd-alloys in H_2 and CO environments, implemented via the open-source Python package ICET [77].

3.3.1 Definitions and formalism

In a CE, the atomic configuration is represented by a configuration vector

$$\sigma = \{\sigma_1, \sigma_2, \dots, \sigma_N\},$$

where σ_i specifies the occupation of lattice site i and N is the total number of sites. For a binary system with species A and B , σ_i is commonly set to 0 for species A and 1 for B .

A cluster is defined as a set of k lattice sites. Clusters are classified by their order k and their geometric size, which can be defined based on the distances between sites. Clusters obey the symmetry of the underlying lattice, and a group of symmetrically equivalent clusters form an orbit.

For each site i , a set of M point functions $\Theta_n(\sigma_i)$ is defined as

$$\Theta_n(\sigma_i) = \begin{cases} 1, & n = 0 \\ -\cos\left(\frac{\pi(n+1)\sigma_i}{M}\right), & n \text{ odd} \\ -\sin\left(\frac{\pi n \sigma_i}{M}\right), & n \text{ even} \end{cases}$$

where M is the number of allowed species on the site and $n = 0, 1, \dots, M-1$ is the point function index. The point functions form a complete, orthogonal basis of the configuration space [78] with basis functions

$$\Pi_{\alpha}(\sigma) = \Theta_{n_1}(\sigma_1)\Theta_{n_2}(\sigma_2) \cdots \Theta_{n_l}(\sigma_l),$$

where $\alpha = \{n_1, n_2, \dots, n_N\}$ assigns a point function index to each site. For binary alloys, n_i takes the value 1 for all sites belonging to the cluster and 0 elsewhere, ensuring that only sites in the cluster contribute. For multicomponent systems with M species, n_i can take values $1, 2, \dots, M-1$ for sites included in the cluster. Different index permutations produce multiple basis functions $\Pi_{\alpha}^{(s)}(\sigma)$ associated with the same cluster α , where (s) indicates the corresponding point function index.

Any function f of the configuration (most often the energy) can be expanded exactly as

$$f(\sigma) = f_0 + \sum_{\alpha, s} f_{\alpha}^{(s)} \Pi_{\alpha}^{(s)}(\sigma),$$

where $f_{\alpha}^{(s)}$ are the contributions from non-zero order clusters and f_0 is the configuration independent zeroth order contribution [76]. In practice, the number of included orbits

has to be truncated by imposing upper limits for the cluster order and size. Using lattice symmetry, the sum over all clusters can be reduced to a sum over all orbits, each represented by one of its clusters α , and the averaged basis function over all clusters α' in the orbit,

$$f(\sigma) = J_0 + \sum_{\alpha,s} m_{\alpha} J_{\alpha}^{(s)} \langle \Pi_{\alpha'}^{(s)}(\sigma) \rangle_{\alpha}. \quad (3.17)$$

Here, m_{α} is the multiplicity of the orbit containing α and the energetic contributions have been replaced by the effective cluster interactions (ECIs) $J_{\alpha}^{(s)}$.

For systems where the allowed species differ across sites, a CE consisting of multiple sublattices can be constructed. For the alloy–adsorbate systems studied in this thesis, this corresponds to one sublattice for the metal slab and one for the adsorbates, where the latter can be occupied by adsorbates or vacancies to represent different coverages.

3.3.2 Cluster expansion construction

A CE is constructed by determining the optimal ECIs based on reference data $f = [f_1(\sigma_1), f_2(\sigma_2), \dots]^T$ for a set of selected atomic configurations $\{\sigma_1, \sigma_2, \dots\}$. The reference data is typically obtained from first-principles calculations. Although the CE models an ideal lattice, the reference structures often have fully relaxed atomic positions and cell sizes. The structures are then rescaled and mapped onto the ideal lattice, effectively incorporating relaxation effects into the lattice model.

To find the optimal ECIs, represented by the vector $J = [J_0, J_1, J_2, \dots]^T$ (the index s is omitted for brevity), Eq. (3.17) is rewritten as a linear problem

$$f = \Pi J. \quad (3.18)$$

Each row Π_i of the sensing matrix Π contains the averaged basis functions for all orbits multiplied by their multiplicities, and is referred to as a cluster vector,

$$\Pi_i = [1, m_{\alpha_1} \langle \Pi_{\alpha'_1}(\sigma_i) \rangle_{\alpha_1}, m_{\alpha_2} \langle \Pi_{\alpha'_2}(\sigma_i) \rangle_{\alpha_2}, \dots]. \quad (3.19)$$

Solving the linear problem is equivalent to finding

$$J_{\text{opt}} = \underset{J}{\operatorname{argmin}} \{ \|\Pi J - f\|_2^2 + J^T \Lambda J + \|MJ\|_1 \}, \quad (3.20)$$

which includes standard ℓ_2 (Λ) and ℓ_1 (M) regularization terms. Different choices of Λ and M correspond to different regression methods. The unregularized ordinary least-squares solution is obtained by setting $\Lambda = M = 0$. Setting $\Lambda = 0$ and $M = \alpha I$ yields the least absolute shrinkage and selection operator method [79]. Conversely, choosing

$\Lambda \neq 0$ and $M = 0$ corresponds to different variations of ridge regression, with the standard ridge approach obtained by $\Lambda = \alpha I$, Bayesian ridge by $\Lambda_{ij} = \delta_{ij} \lambda_i$ and automatic relevance detection regression by extending Bayesian ridge regression with hyperpriors for λ_i and subsequent pruning [80]. Regularization generally produces sparser ECIs and is often essential for obtaining stable CEs, as further discussed in Paper I.

3.3.3 Improving cluster expansions with physical intuition

For low symmetry systems, such as surfaces, the number of ECIs can become very large, making data-efficient fitting challenging. Paper I introduces several strategies for addressing this issue, including coupling of similar orbits and applying weights and constraints to improve the prediction of specific properties. Below, two approaches for incorporating physical intuition into the CE construction process are presented.

3.3.3.1 Exploiting local symmetries

For a bulk binary alloy with a simple face-centered cubic (FCC) structure, the primitive cell consists of a single site. Consequently, the list of orbits includes one singlet, one nearest neighbor-pair, and so forth. For a surface slab of the same alloy, however, the primitive cell consists of one site per atomic layer. This creates several symmetrically inequivalent singlets and nearest neighbor-pairs. For example, the nearest neighbor-pair between layers i and $i+1$ is not equivalent to the pair between layers $i+1$ and $i+2$, and as a result they will be represented by different ECIs. However, if the slab is sufficiently thick and layer i lies far from the surface, the behavior of these pairs should nevertheless be similar and approach that of the single nearest neighbor-pair in the bulk.

This observation motivates the introduction of local symmetries that couple orbits with matching order and geometric radii that are expected to behave similarly despite not being strictly symmetrically equivalent. By merging coupled orbits, the number of ECIs is reduced. In Paper I – III, this approach is successfully applied to Pd-alloy surfaces with adsorbates. In Paper III, the approach is further extended by training bulk ECIs separately for a bulk system, enabling a consistent bulk description across multiple CEs representing different surface orientations.

3.3.3.2 Bayesian cluster expansion

A more general framework for incorporating physical intuition is through a Bayesian CE, where prior beliefs about the system are encoded via Bayesian priors [81–83]. Below, this approach is introduced based on Ref. 81.

Assuming Gaussian priors for the ECIs yields the posterior distribution

$$P(\mathbf{w}|X) \propto \prod_{\alpha} e^{-w_{\alpha}^2/2\sigma_{\alpha}^2} \prod_{\alpha, \beta \neq \alpha} e^{-(w_{\alpha} - w_{\beta})^2/2\sigma_{\alpha\beta}^2}.$$

The first product controls the magnitude of the ECIs via σ_α , the prior standard deviation, which can be parametrized to, for example, yield decreasing ECI with increasing cluster size or order. The second product introduces coupling between orbits through $\sigma_{\alpha\beta}$, the inverse coupling strength between orbits α and β . By decreasing $\sigma_{\alpha\beta}$ for orbits that are believed to be similar, the effect will be similar to that of merging.

The priors enter the linear regression problem (Eq. (3.20)) via the regularization matrix Λ , extending the standard Bayesian ridge formulation. The resulting matrix has diagonal elements

$$\Lambda_{\alpha\alpha} = \frac{\sigma^2}{\sigma_\alpha^2} + \sum_{\beta \neq \alpha} \frac{\sigma^2}{\sigma_{\alpha\beta}^2},$$

and off-diagonal elements

$$\Lambda_{\alpha\beta} = \Lambda_{\beta\alpha} = -\frac{\sigma^2}{\sigma_{\alpha\beta}^2},$$

where σ is a scaling constant.

3.3.4 Training data generation

The predictive accuracy of a CE depends strongly on the atomic configurations included in the training set. Configurations that are poorly represented during training will be associated with a larger error in the predicted energy. Increasing the size of the training set typically improves accuracy and reduces the risk of overfitting, but doing so requires additional costly DFT calculations. It is therefore important to select structures with care.

Many different approaches to structure selection for CEs have been proposed [83–90]. Random generation is conceptually simple but often inefficient, as many randomly drawn configurations share similar cluster vectors (Eq. (3.19)), limiting the information gained and leading to poorly conditioned regression problems (Eq. (3.18)). For systems with small unit cells, all symmetrically distinct structures up to a given size can be enumerated [85, 86], but this becomes infeasible for larger unit cells since the number of structures quickly becomes unmanageable. In the following, two approaches to structure generation are presented.

3.3.4.1 Active learning

Active learning refers to the iterative process of training intermediate models used to expand the training set by identifying structures that are of particular importance for the problem. That could, for example, include structures close to the ground state [87, 89], structures that minimize the variance of the predicted property [84, 88, 89], or structures associated with large model uncertainty [90]. Active learning is widely used in ML as it provides a pragmatic and flexible strategy for improving model accuracy.

3.3.4.2 Orthogonalization of cluster vectors

A compelling idea is to generate a training set where each structure contributes with as much new information as possible, by aiming for structures with orthogonal cluster vectors [83]. In Ref. [83], the following iterative algorithm is proposed

1. Generate a random cluster vector.
2. Orthogonalize this vector with respect to the cluster vectors of all previously selected structures.
3. Identify the structure whose cluster vector most closely matches the orthogonalized vector and add it to the training set.

When bench-marked against other methods for the bulk system in Paper I, this approach was outperformed. This is likely due to correlations between the elements of the cluster vectors making it impossible to obtain a sensing matrix with perfectly orthogonal rows. For the surface systems modeled in Paper II and Paper III, however, this approach was adopted for most of the structures (Paper II) or as a starting point for further structure generation (Paper III) with good results.

3.4 Monte Carlo simulations

CEs are often sampled with MC simulations to obtain thermodynamic information about the system. MC methods form a broad class of algorithms that use random sampling to evaluate deterministic properties [91, Chapter 10]. For alloy modeling, MC simulations, and specifically the Metropolis algorithm [92], provide an efficient way to sample the vast configuration space and obtain thermodynamic averages such as free energies, site occupancies, or order parameters.

The probability of a specific atomic configuration σ is

$$p(\sigma) = \frac{1}{Z} \exp \left[-\frac{\Phi(\sigma)}{k_B T} \right], \quad (3.21)$$

where Z is the partition function, Φ is the thermodynamic potential associated with the chosen thermodynamic ensemble (introduced in the following sections), k_B is the Boltzmann constant, and T is the temperature. A simulation starts from an initial supercell with an arbitrary atomic configuration. At each step, the system is subjected to a trial step where the atomic configuration undergoes some random change, for instance, swapping the atomic species of two lattice sites. The probabilities (Eq. (3.21)) of the atomic configuration before (σ_1) and after (σ_2) the swap are compared. If the ratio

between the latter and the former is larger than a uniform random number $\xi \in (0, 1)$,

$$\frac{p(\sigma_2)}{p(\sigma_1)} > \xi, \quad (3.22)$$

the new configuration (σ_2) is accepted, otherwise the system remains in the old configuration (σ_1). Note that the partition function Z never has to be calculated explicitly. With this algorithm, lower energy configurations are always accepted, while higher energy configurations are accepted only sometimes, which enables exploration of low energy configurations without becoming trapped in local minima. Quantities of interest are recorded during the simulation, and because configurations are sampled according to Eq. (3.21), the resulting averages correspond to ensemble averages.

The choice of thermodynamic ensemble, which determines the thermodynamic potential in Eq. (3.21), depends on the properties of the system. For example, an isolated system with a constant number of particles, volume, and energy is described by the microcanonical ensemble. In the microcanonical ensemble, a macrostate is defined by the values of the constant macroscopic variables (N, V, E). A macrostate generally corresponds to multiple microstates with varying probabilities. During a MC simulation, the system changes between different microstates while remaining in the same macrostate. In the following, two thermodynamic ensembles relevant to this thesis are presented.

3.4.1 The canonical ensemble

The canonical ensemble represents a system at constant volume V with a fixed number of atoms of each species N_i (where i is the atomic species) in contact with an infinite thermal reservoir at constant temperature T . A macrostate in the canonical ensemble is thus described by (N_i, V, T) . The thermodynamic potential, which determines the probability of a microstate, corresponds to the internal energy.

In this thesis, the internal energy is represented by the configurational energy predicted by the CEs while other contributions, such as vibrations, are neglected. The microstate is then determined by the atomic configuration σ and the probability is

$$p_C(\sigma) = \frac{1}{Z} \exp \left[-\frac{E(\sigma)}{k_B T} \right], \quad (3.23)$$

where E is the configurational energy obtained, for example, from a CE. Since the numbers of particles of each species are constant, the atomic concentrations are fixed, and a suitable trial step consists of swapping the species of two sites.

3.4.2 The semi-grand canonical ensemble

In the semi-grand canonical (SGC) ensemble, the total number of atoms N , the volume V and the temperature T are constant. The difference from the canonical ensemble is

that the number of atoms per species, N_i , can vary (under the condition that $N = \sum_i N_i$) while the chemical potential difference $\Delta\mu_i = \mu_1 - \mu_i$ (for each species $i > 1$ in relation to the first species) is constant [93]. Fixing the chemical potential difference effectively controls the average concentrations during a MC simulation. The SGC ensemble thus represents a system in contact with an infinite reservoir with fixed temperature and chemical potential differences, such that both energy and atoms can be exchanged, and a macrostate is described by $(N, \Delta\mu_i, V, T)$.

The probability of a microstate in the SGC ensemble is defined as

$$p_{\text{SGC}}(\sigma) = \frac{1}{Z} \exp \left[-\frac{E(\sigma) + \sum_{i>1} \Delta\mu_i N_i}{k_B T} \right], \quad (3.24)$$

where E is the configurational energy and the second term is associated to the potential energy related to the chemical potentials. A trial step consists of changing the occupancy of one site.

The SGC ensemble has two main advantages over the canonical ensemble. First, the control of $\Delta\mu_i$ allows for efficient sampling over the entire concentration range and the free energy can be obtained by integration of the thermodynamic relation

$$\Delta\mu_i = -\frac{1}{N} \frac{\partial F}{\partial c_i}, \quad (3.25)$$

which is convenient for, for example, phase diagram construction. Second, for surface-adsorbate systems, knowledge of the chemical potential $\mu_i(T, p_i)$ provides a link to the corresponding partial pressure p_i (Eq. (2.7)).

Segregation and adsorption on Pd-alloy surfaces

I want to break free

Queen

Plasmonic H_2 sensing begins with H adsorbing to the surface. Therefore, understanding how the atomic configuration and the presence of competing gas species affect H adsorption is essential. Here, the modeling framework described in Chapter 3 is applied to Pd-alloy surfaces exposed to different gas environments to study surface segregation and adsorption. The first objective is to determine how the environment influences surface segregation tendencies, an aspect that is crucial for optimizing fabrication conditions and ensuring long-term stability. The second objective is to quantify how different alloy surfaces affect the CO–H coadsorption behavior, which is central to understanding CO poisoning.

4.1 Adsorption and segregation in simple systems

Qualitative insights regarding the surface segregation tendencies of alloys can be obtained by studying pure metal and dilute alloy surfaces. Figure 4.1a shows the segregation energies associated with a single Au or Cu atom positioned in different atomic layers of a Pd surface slab with a 2×2 surface supercell and 10 layers. In vacuum, Au shows a strong preference for the surface. This is in line with the general segregation predictors presented in Sect. 2.2, due to the lower surface energy and larger lattice parameter (Fig. 4.1c) of Au compared to Pd. For the case of Cu in Pd in vacuum, the segregation

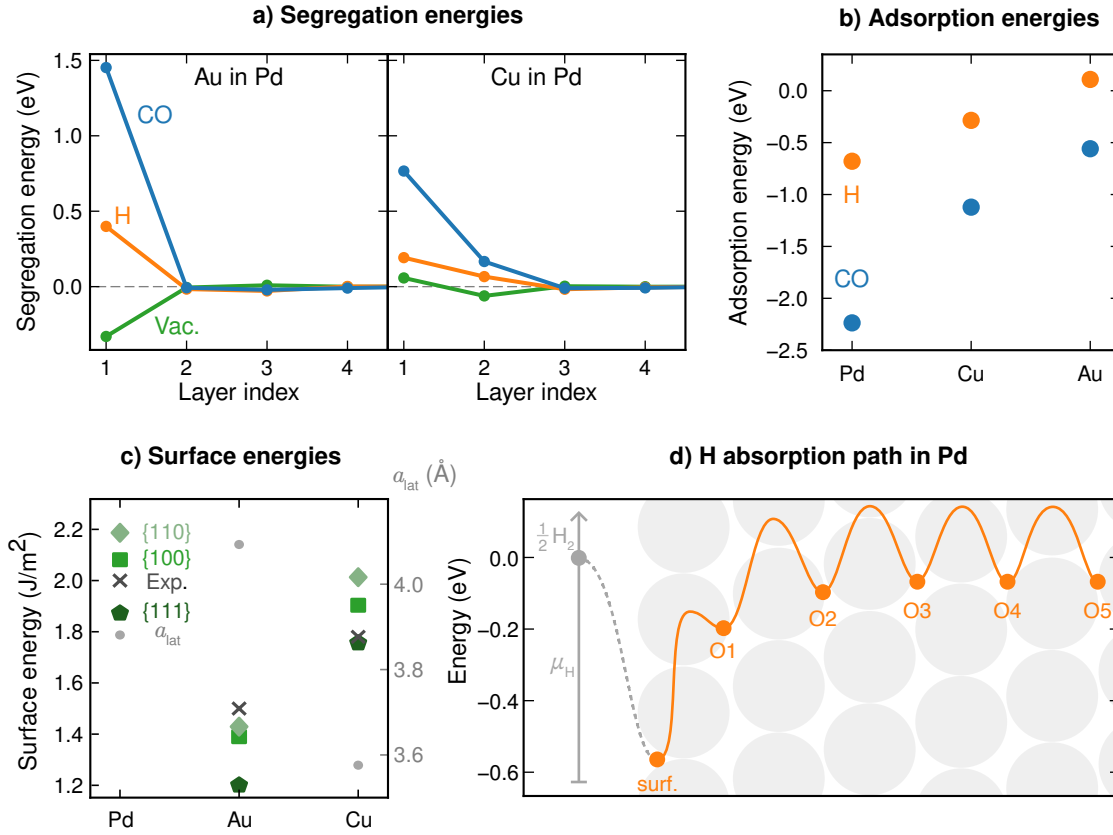


Figure 4.1: Surface properties relevant for adsorption and segregation. (a) Segregation energies for Pd- $\{111\}$ alloyed with Au or Cu in the dilute limit, in vacuum or with a monolayer of H or CO adsorbed to the FCC surface site, obtained from DFT in Paper III. (b) Adsorption energies for 25 % surface coverage of H or CO at the energetically favorable $\{111\}$ -FCC site on Pd, Cu, and Au, obtained from DFT in Paper III. (c) Surface energies and lattice parameters a_{lat} from Ref. 72 calculated with density-functional theory (DFT) using the same functional and general settings as in the work of this thesis, and experimentally measured surface energies from Ref. 94. (d) H absorption path in Pd- $\{111\}$ with relaxed site energies and (NEB) kinetic barriers calculated with the MACE model in Paper III.

tendencies are notably weaker than for Au with a small preference for Pd at the surface and Cu in the second atomic layer. With CO or H adsorbed to the surface, Pd is strongly preferred in the surface layer, also in line with the discussion in Sect. 2.2, due to the strong adsorption between these species and Pd (Fig. 4.1b).

Figure 4.1d shows the energetics of a H atom adsorbing to a Pd surface and migrating into the bulk. The site energies and kinetic barriers are obtained from structural relaxation and nudged elastic band (NEB) [95–97] calculations via the message passing atomic cluster expansion (MACE) model in Paper III. Here, the H chemical potential

(shown to the left) can be viewed as the controller of the H absorption process. Increasing the chemical potential (i.e., increasing the pressure) will first fill the adsorption sites at the surface. Upon further increase, the first octahedral (subsurface) site will be filled, and lastly the subsequent octahedral sites all the way into the bulk. Although this thesis is limited to the adsorption of H, this picture is useful to have in mind in the following discussion.

4.1.1 From chemical potential to pressure

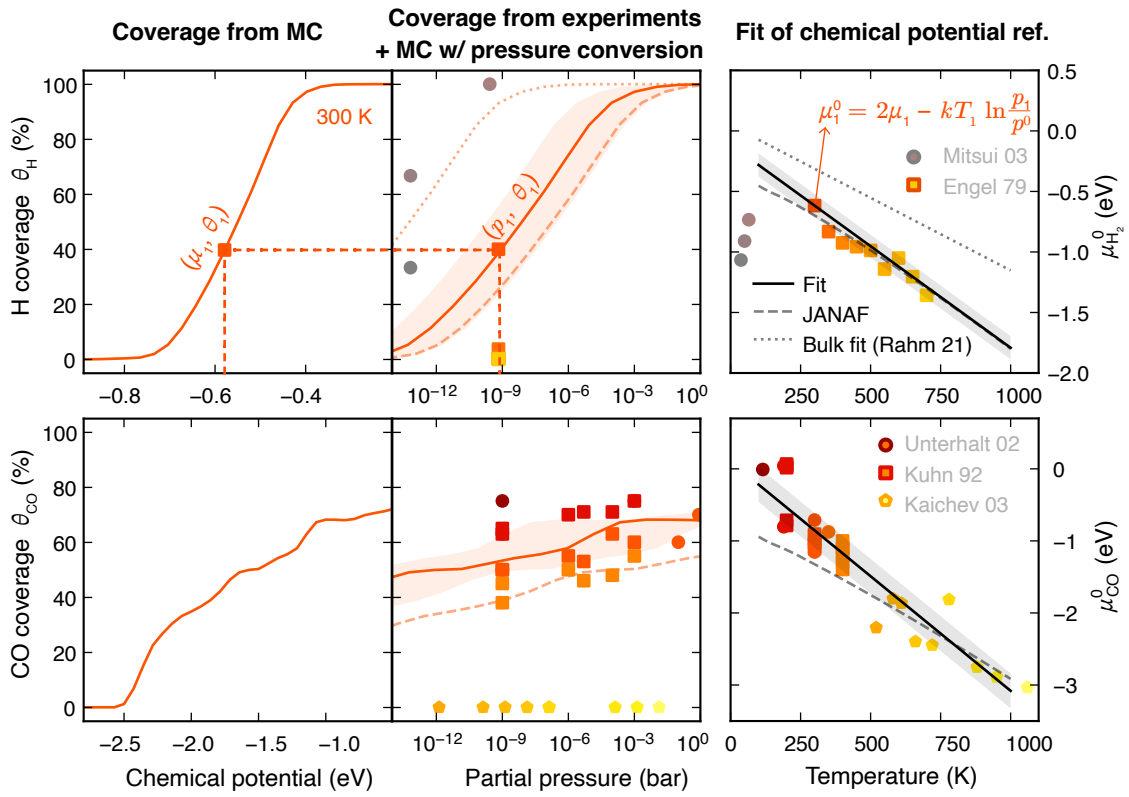


Figure 4.2: Converting the chemical potential to partial pressure for Pd with H (upper panel) and CO (lower panel) adsorbed to the surface. The first column shows the adsorbate coverages as a function of the chemical potential based on Monte Carlo (MC) simulations at 300 K. The second column shows experimental records of adsorbate coverages as a function of the partial pressure [98–106], as well as the MC simulated coverages after converting to partial pressure based on the reference chemical potential μ^0 shown in the third panel. The reference chemical potential μ^0 is either from thermodynamical tables (JANAF) [53], modeling of the bulk Pd–H system [107] (Bulk fit, for H specifically) or from a fit of the data from the first two columns (Fit).

The cluster expansion (CE)-MC framework of this thesis can provide adsorbate cov-

erages as a function of the involved chemical potentials. In order to relate such results to reality, a conversion between chemical potential and partial pressure is necessary. A conceptually simple approach is to use JANAF thermodynamical tables [53] to calculate the gas phase free energy entering in the reference chemical potential μ^0 (Eq. (2.8)). This is one of the approaches considered for H and CO in Paper III, with an additional constant correction factor to account for the typical error in adsorption energy compared to experimental measurements (Fig. 4.2). A limitation of this approach is the notorious difficulty in achieving quantitative agreement between first-principles calculations and experiments [108], as well as the lack of vibrational free energy contributions from the adsorbates. For both H and CO on Pd, the resulting pressure-converted MC-simulated coverages are underestimated compared to the experimental records, when using the JANAF data (Fig. 4.2).

Another approach is to use experimental data of the observable of interest as a reference point. For example, one can use information about the surface coverage (Fig. 4.2), in which case any relevant vibrational contributions should be included in μ^0 . This approach relies on the availability of such experimental data, and for Pd:H, in particular, few observations are available at the relevant coverages, temperatures and pressures. Furthermore, the plateau pressure of the Pd–H bulk system can be considered, for which a lot more experimental data is available [107]. While such an approach accurately recovers the plateau pressure for the Pd–H bulk system, surface-specific adsorbate effects are not taken into account. Compared to the other two approaches, the μ^0 obtained in this fashion shifts the H uptake to lower pressures. Based on further analysis of the results below (Sect. 4.2), it appears that this last approach yields the most reliable conversion for the case of H.

This comparison illustrates that converting the chemical potential to partial pressure is inherently difficult and is associated with large uncertainty. Already an error of 0.1 eV in the reference chemical potential, which is comparable with the typical uncertainty of adsorption energy calculations at the DFT level, leads to several orders of magnitude uncertainty in the pressure. Still, pressure conversion is a crucial aspect in order to relate simulated results to reality.

4.2 Surface segregation of Pd-alloys in H₂

In Paper II, the surface segregation behavior of PdAu- and PdCu-alloys in different H₂ environments is studied. The atomic configuration, both alloyants and adsorbates, are allowed to fully equilibrate to study the surface state in equilibrium with its environment. This corresponds to the effect of long-term storage in a constant environment, or annealing at high enough temperatures to overcome kinetic barriers on reasonable timescales.

The Pd–Au system exhibits strong segregation of the first atomic layer, shifting



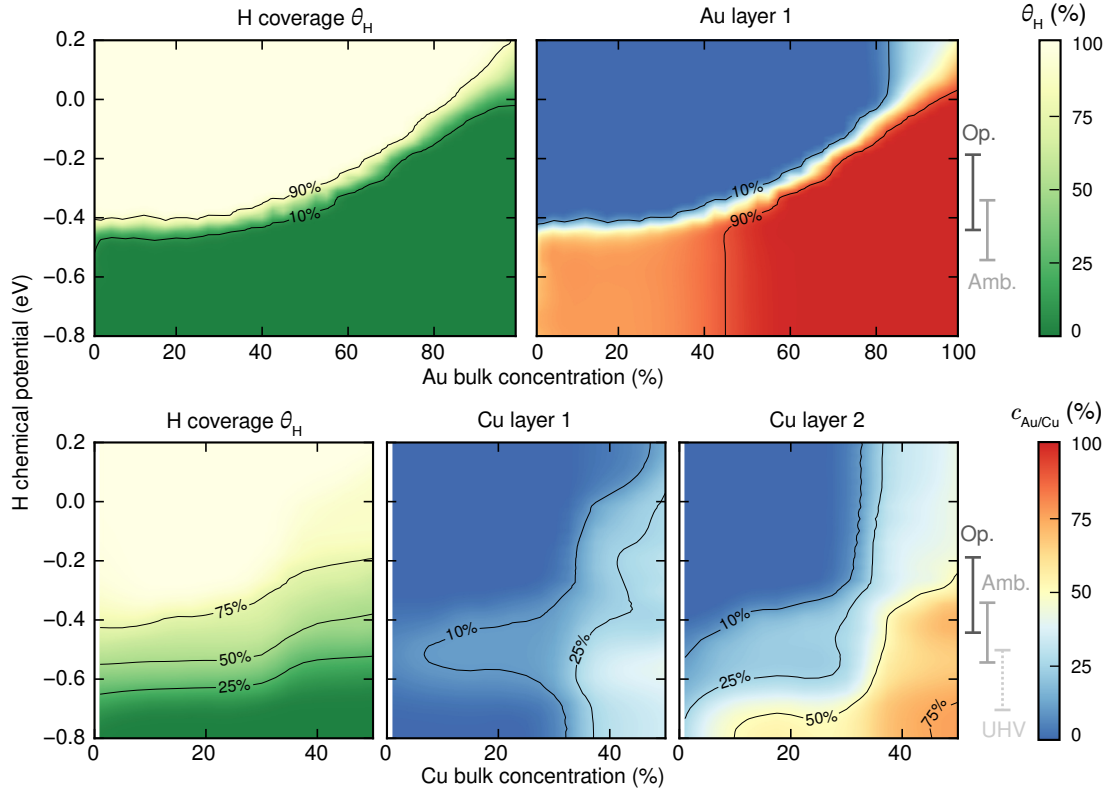


Figure 4.3: H-induced segregation at 300 K for PdAu (upper panel) and PdCu (lower panel) {111} surfaces. Indicated to the right are the approximate pressure regions corresponding to ambient (6×10^{-5} bar [109]), operating (1 mbar–100 mbar H₂ [110, 111]) and UHV ($\sim 10^{-12}$ bar) conditions where the spread stems from the different μ^0 values in Fig. 4.2.

abruptly from at least 75 % Au and no H adsorption to 100 % Pd and 100 % H coverage at a threshold H chemical potential. These findings are in line with the trends expected based on the principles described above, as well as the literature for PdAu surfaces in vacuum [39, 112–119] and in H₂ [117, 119–121].

Because of the uncertainty associated with the pressure conversion (Sect. 4.1.1), the estimated pressure intervals corresponding to operating and ambient conditions both span over the shift in segregation behavior, in the limit of high Pd bulk concentration. According to Fig. 4.3, however, H adsorbs on PdAu at approximately $\mu_H = -0.4$ eV and the onset for H absorption is expected a few tenths of an eV above this value (Fig. 4.1d). This suggests that the upper ends of the estimated pressure ranges are more realistic, since that places the operating range (where H clearly absorbs) at $\mu_H = -0.2$ eV. That would place the system at ambient conditions right above the transition to Pd-rich and H-covered.

For the Pd–Cu system, the shift in segregation tendencies is more gradual. Three

separate regions can be identified. For high H coverages, the first and second atomic layers are dominated by Pd, similar to the case of Pd–Au. At moderate H coverages, however, a very slight excess of Cu in the first and second layers is observed in the Pd-rich bulk composition limit. With no H adsorbed, the surface is again dominated by Pd while the second atomic layer sees a strong surplus of Cu.

This non-trivial behavior is interesting in the light of an apparent discrepancy in the literature regarding segregation of the Pd–Cu system in vacuum. Many experimental studies find Pd enrichment in the surface region (~ 6 layers), but Cu enrichment in the topmost surface layer(s) [122–124]. While one computational study supports this [125], the majority of theory-based work reports the opposite behavior with a slight enrichment of Pd in the topmost surface layer [39, 112, 113, 119], at least for low Cu concentrations [117, 126]. For the Pd–Cu system in a H_2 environment, the literature includes findings of a Pd enrichment at the surface [117, 127] and Pd depletion at the surface for high Cu concentrations [128, 129].

The discrepancy of the segregation behavior in vacuum is primarily between theoretical and experimental work, and could be attributed to lacking resolution of experimental techniques and systematic errors in the computational methods. The findings of Paper II suggest an alternative explanation, where the experimental ultra-high vacuum (UHV) conditions correspond to the region with Cu surplus at the surface, while computational work considers “true” vacuum, which is associated with a Pd surplus at the surface.

4.3 Competing H and CO adsorption on Pd-alloys

In Paper III, the coadsorption of H and CO on PdAuCu surfaces under realistic operating conditions is studied, meaning timescales that allow for equilibration of the adsorbate coverage while leaving the alloy configuration unchanged. The alloy configurations are obtained from MC simulations at elevated temperatures with varying H coverage, modeling the experimental procedure of annealing in H_2 [35, 36]. Subsequently, the configuration is kept fixed while running MC on the adsorbate sublattice at 300 K. This procedure results in surface phase diagrams (SPDs) of the H and CO coverage as a continuous function of their respective chemical potentials, in analogy to the one-dimensional discrete SPD of Fig. 2.4.

Figure 4.4 shows SPDs for various annealing conditions and alloy compositions. For $Pd_{75}Au_{25}$, the amount of CO and H adsorbed increases with the H coverage during annealing (Fig. 4.4a–c), approaching the limit of the pure Pd surface (Fig. 4.4d). This is a direct consequence of an increase in the amount of Pd at the surface with increasing H coverage during annealing, which favors CO and H adsorption.

When introducing Cu and varying the alloy composition (Fig. 4.4e–g), the SPD remains almost identical to that of $Pd_{75}Au_{25}$ (Fig. 4.4c). In Paper III it is shown that this

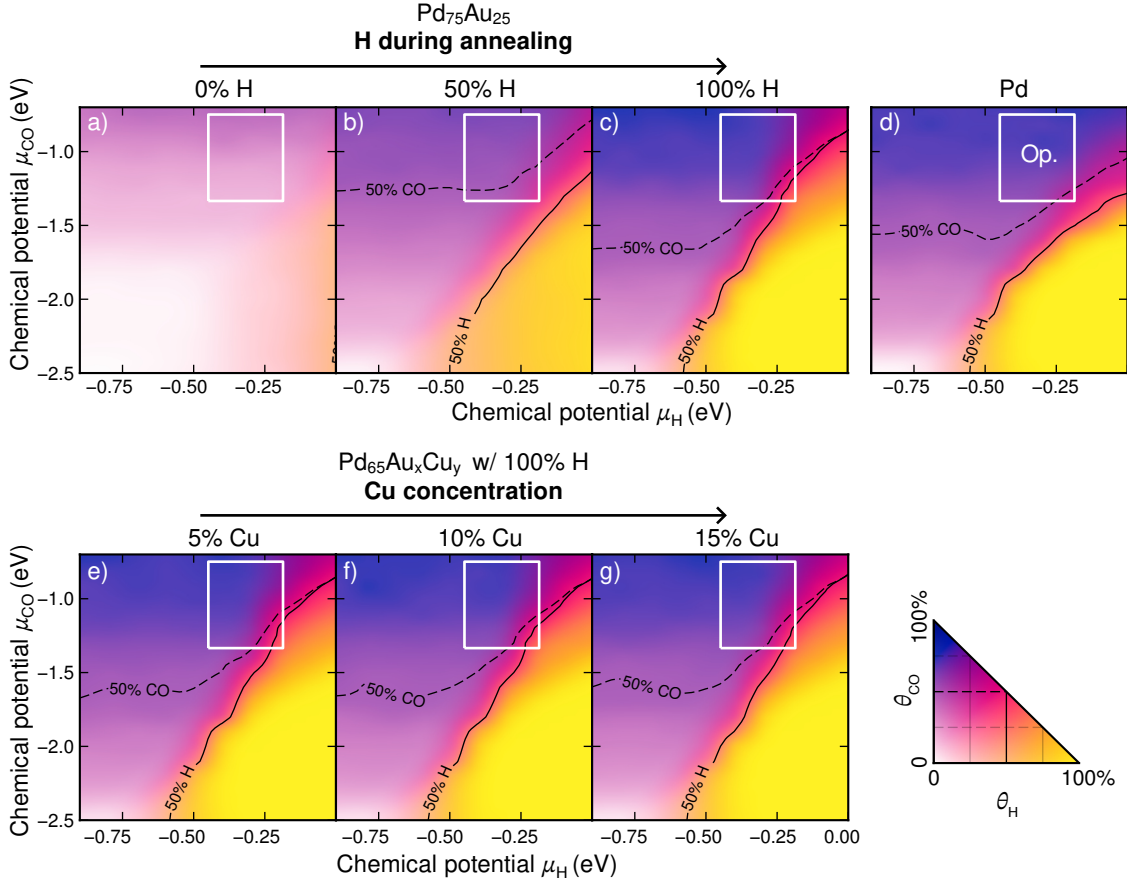


Figure 4.4: Surface phase diagram (SPD) of the CO and H coverage on PdAuCu-alloy {111} surfaces, where the alloy configuration is obtained from annealing at 773 K before exposing the system to H and CO at varying chemical potentials and 300 K while keeping the alloy configuration frozen. The coverage is represented by the colormap, as well as solid and dashed contour lines for θ_{H} and θ_{CO} , respectively. The upper panel shows the variation of the SPD for $\text{Pd}_{75}\text{Au}_{25}$ with H coverage during annealing. The lower panel shows the (lack of) variation with the Cu concentration for alloys annealed with 100 % H.

is true for all considered alloy compositions and annealing conditions. Alloying with Au and Cu is, however, found to increase the amount of H adsorbed in the coadsorption limit, as can be seen in Fig. 4.4 in the lower right corner of the operating region for the alloys annealed in 100 % H (Fig. 4.4c, e–g) compared to pure Pd (Fig. 4.4d) as well as in Fig. 4.5a. This observation could explain why alloying in general is successful in reducing CO poisoning [35, 36].

The experimentally observed mitigating effect of Cu, specifically, on CO [35, 36] is not apparent from adsorption thermodynamics alone (Fig. 4.5a). The measured optical response in a mixture of CO and H_2 reveals that CO poisoning greatly affects the H

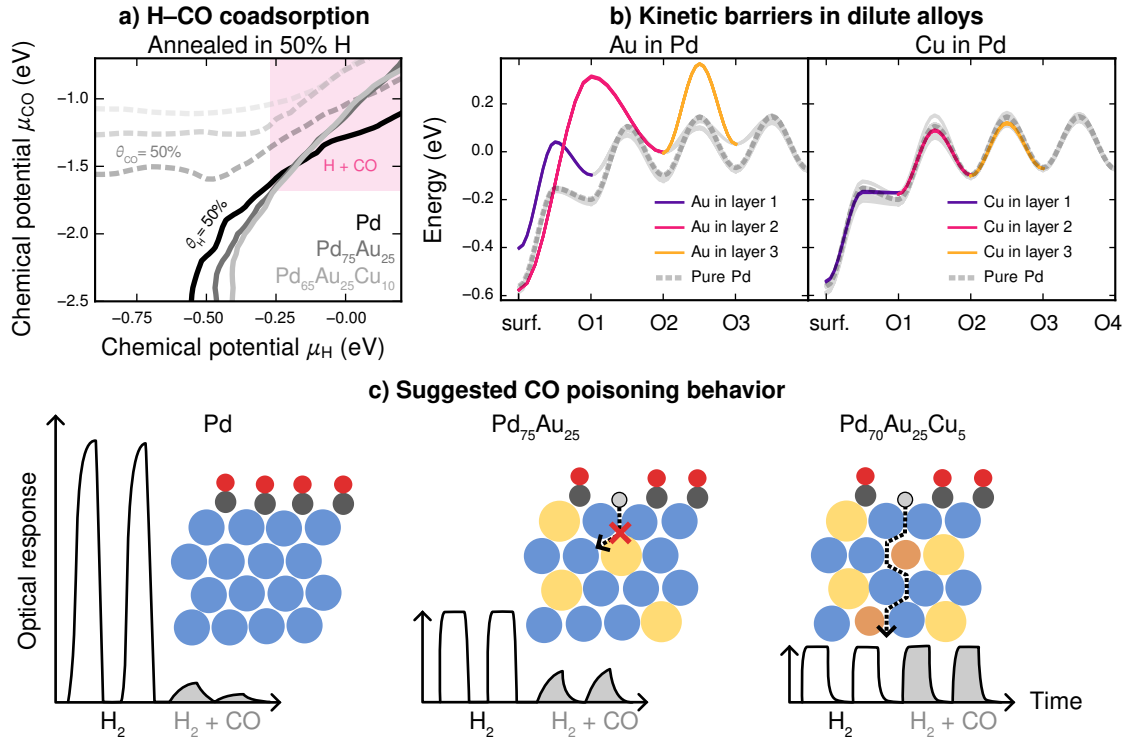


Figure 4.5: (a) The 50 % coverage contour lines from the SPDs for Pd, Pd₇₅Au₂₅, and Pd₆₅Au₂₅Cu₁₀, with the H–CO coadsorption region highlighted in pink. (b) Change in site energies and kinetic barriers when introducing a single Au or Cu atom in atomic layers 1, 2, or 3 of a Pd slab with a 3×4 surface supercell, compared to pure Pd, from NEB calculations using the MACE model in Paper III. The highlighted paths pass right next to a Au or Cu atom. (c) Suggested mechanism of CO poisoning (mitigation) for pure Pd, PdAu, and PdAuCu, as well as the measured optical shifts in H₂ and H₂ + CO adapted from [35].

absorption kinetics without Cu, motivating studies of the associated kinetic barriers. Figure 4.5b shows the energy landscape for H absorption and migration for a 10-layer Pd slab with a 3×4 surface supercell, with a single Au or Cu atom placed in atomic layers 1, 2, or 3, respectively. Clearly, passing a Au atom is associated with a large increase in the initial and final site energy as well as the associated kinetic barrier, while a single Cu atom has a minimal effect on the H migration energetics, compared to pure Pd.

Based on the results for coadsorption behavior and kinetic barriers, the following hypothesis can be formulated (Fig. 4.5c). For pure Pd, the surface is completely blocked by CO under operating conditions. For the alloys, the surface is only partially blocked, allowing some H to adsorb. H still needs to absorb through the surface, and the barriers associated with this process will depend on the local chemical environment. It is thus suggested that the presence of Cu in the surface region provides viable paths for H absorption when the most favorable paths are blocked by CO.

Part II

Optical properties of Pd nanoalloys

Plasmonic resonances of metal nanoparticles

Mmm num ba de

Dum bum ba be

Doo buh dum ba beh beh

Queen

H₂ sensing relies not only on efficient H uptake, but also on the H-induced shift in optical response. In this chapter, the optical response of metal nanoparticles (NPs) is described within the framework of classical electrodynamics. Starting from the dielectric function (DF), the mechanisms behind the formation of localized surface plasmon resonances (LSPRs) and surface lattice resonances (SLRs) are explained.

5.1 The dielectric function

The intrinsic optical properties of a material are encoded in its dielectric function (DF). This includes electronic and ionic excitations, as well as various types of electron scattering [28, Chapter 9]. If the DF is known, it can be applied in classical frameworks such as Mie theory or electrodynamic simulations to study macroscopic optical properties based on the atomistic quantum mechanical properties of the material. To qualitatively connect the DF to the underlying physical phenomena, the classical Drude and Lorentz models provide a useful foundation.




Figure 5.1: DFs based on the Drude, Lorentz and Drude-Lorentz models.

The Drude model [28, Chapter 9],

$$\varepsilon(\omega) = 1 - \frac{\omega_p^2}{\omega^2 + i\gamma\omega}, \quad (5.1)$$

describes the optical properties of a free-electron system (Fig. 5.1). Despite its simplicity, it captures key physical phenomena. The plasma frequency, $\omega_p^2 = \frac{ne^2}{m^*\varepsilon_0}$, depends on the free-electron density n and the effective mass m^* (with the elementary charge e and the vacuum permittivity ε_0 as constants), and corresponds to the bulk plasmon frequency. A Drude metal is highly reflective at frequencies below ω_p , typically covering the visible spectrum. Above ω_p , the DF approaches the permittivity of vacuum, and the metal becomes transparent. The damping constant, $\gamma = \tau^{-1}$, is the inverse of the average time τ between collisions of electrons with other electrons, defects, or phonons, with phonon scattering dominating at ambient temperatures. Together, ω_p and γ determine the electrical conductivity, $\sigma_{DC} = \varepsilon_0\omega_p^2\gamma^{-1}$.

The Lorentz model [28, Chapter 9],

$$\varepsilon(\omega) = 1 + \frac{\omega_p^2}{\omega_0^2 - \omega^2 - i\gamma\omega},$$

describes the optical response of a harmonic oscillator with resonance frequency ω_0 , where the oscillator represents a bound electron or ion (Fig. 5.1). At resonance, $\text{Im } \varepsilon$ increases, leading to higher absorption, followed by an increase in reflectance at slightly higher frequencies due to a dip in $\text{Re } \varepsilon$. Far from resonance, the material behaves like vacuum. Note that setting the resonance frequency, or equivalently the spring constant, to zero reduces the model to the Drude case (Eq. (5.1)).




Figure 5.2: DFs for Pd, Au and Ag obtained from time-domain DFT calculations in [130]. The gray line indicates $\text{Re } \epsilon(\omega) = -2$ corresponding to the LSPR for small spheres according to the quasistatic approximation.

Neither the Drude nor the Lorentz model alone fully capture the optical response of real metals. In most cases, the optical properties depend on both free and localized electrons, with the latter represented by multiple oscillators of different frequencies and amplitudes. The DF can then be modeled by a combination of the Drude and Lorentz models, the Drude–Lorentz model [28, Chapter 9],

$$\epsilon(\omega) = 1 - \frac{\omega_p^2}{\omega^2 + i\gamma\omega} + \sum_j \frac{\omega_{p,j}^2}{\omega_{0,j}^2 - \omega^2 - i\gamma_j\omega}, \quad (5.2)$$

where the sum runs over all relevant Lorentz-type contributions (Fig. 5.1). In such metals, the optical response depends on how the bound-electron resonances are spectrally positioned relative to the free-electron behavior. This interplay gives rise to the characteristic colors of different metals.

Figure 5.2 shows the dielectric functions of Pd, Au, and Ag calculated using time-dependent density-functional theory (TDDFT) in Ref. 130. Au and Ag exhibit similar DFs, characterized by a sharp Drude peak below 1 eV and an increase in the imaginary part near 2 eV for Au and 4 eV for Ag, corresponding to the d-band onset, i.e., the distance between the upper edge of the d-band and the Fermi level [130]. For Pd, the imaginary part of the DF is qualitatively different because the d-band edge lies slightly above the Fermi level [130], enabling d-band transitions even at thermal energies. As a result, the imaginary part remains large over a broad energy range, with no distinct separation between the Drude peak and d-band transitions.

Differences in the imaginary part of the DF is crucial for LSPRs, which, as will be discussed in Sect. 5.2, occur when $\text{Re } \epsilon$ is negative (specifically below -2 for oblate ellipsoids). A large $\text{Im } \epsilon$ at the LSPR frequency leads to strong absorption losses causing

broadening of the resonance. Consequently, Ag and Au generally exhibit stronger and narrower LSPRs than Pd.

Metals can also exhibit spectrally localized interband transitions, corresponding to single oscillators that produce pronounced Lorentz-like features in the DF [131, 132]. These transitions cause an increase primarily in absorption, resembling a LSPR, but differ in that they are bulk phenomena independent of the environment or geometry. When an interband transition lies spectrally close to a LSPR, the two can couple, leading to avoided-crossing behavior [131, 132].

5.2 Localized surface plasmon resonances

A plasmon is a collective oscillation of free electrons driven by an external electric field. In a NP, spatial confinement gives rise to a distinct plasmon mode, the localized surface plasmon resonance (LSPR), in which free electrons oscillate collectively between opposite sides of the particle, producing an internal oscillating field. The spectral position and width of the LSPR depend on the material, size, shape, and surrounding medium.

The LSPR can be described within classical electrodynamics. The polarizability $\alpha(\omega)$ of a material is defined as

$$\alpha(\omega) = \frac{1}{\epsilon_m} \frac{\mathbf{p}(\omega)}{E(\omega)},$$

where E is the total electric field, ϵ_m is the relative permittivity of the surrounding medium, and \mathbf{p} is the induced dipole moment. If the polarizability is known, the absorption, scattering, and extinction cross sections follow as

$$\begin{aligned}\sigma_{\text{abs}}(\omega) &= k \operatorname{Im} \alpha(\omega) \\ \sigma_{\text{sca}}(\omega) &= \frac{k^4}{6\pi} |\alpha(\omega)|^2 \\ \sigma_{\text{ext}}(\omega) &= \sigma_{\text{abs}}(\omega) + \sigma_{\text{sca}}(\omega),\end{aligned}\tag{5.3}$$

where k is the wavenumber of the electric field.

Many NP shapes can be approximated as ellipsoids, for which the polarizability can be calculated exactly [28, Chapter 5]. Although this approach has many benefits, its complex form limits intuitive interpretation, can be computationally demanding, and cannot be easily extended to interacting particles. In the following, two simplistic models based on Mie theory are presented, which yield good agreement with the full theory for small particles and facilitate building physical intuition about the nature of LSPRs.

5.2.1 The quasistatic approximation

The quasistatic approximation of Mie theory assumes that for a NP much smaller than the wavelength, the electromagnetic field is effectively static and homogeneous across

the NP at any given time. Under this approximation, the polarizability of a sphere is given by [28, Chapter 5]

$$\alpha_{\text{QS}}(\omega) = 3V \frac{\varepsilon(\omega) - \varepsilon_m}{\varepsilon(\omega) + 2\varepsilon_m}, \quad (5.4)$$

where $\varepsilon(\omega)$ is the DF of the sphere and ε_m is the relative permittivity of the surrounding medium. The LSPR corresponds to the maximum in extinction, or equivalently, the maximum in polarizability (Eq. (5.3)). If $\text{Im } \varepsilon(\omega)$ is small or at least slowly varying, the resonance condition becomes

$$\text{Re } \varepsilon(\omega) = -2\varepsilon_m. \quad (5.5)$$

A similar expression can be derived for small ellipsoids. For an ellipsoid with semiaxes b_x , b_y , and b_z , the polarizability along a principal axis b_i is given by [28, Chapter 5]

$$\alpha_{\text{QS},i}(\omega) = V \frac{\varepsilon(\omega) - \varepsilon_m}{L_i(\varepsilon(\omega) - \varepsilon_m) + \varepsilon_m}, \quad (5.6)$$

where L_i is the geometrical factor

$$L_i = \frac{1}{2} b_x b_y b_z \int_0^\infty dq \left[(b_i^2 + q) \sqrt{(b_x^2 + q)(b_y^2 + q)(b_z^2 + q)} \right]^{-1},$$

which satisfies the relation

$$L_x + L_y + L_z = 1. \quad (5.7)$$

Maximizing the polarizability yields the resonance condition for an ellipsoid

$$\text{Re } \varepsilon(\omega) = \varepsilon_m \left(1 - \frac{1}{L_i} \right). \quad (5.8)$$

Note that for a sphere, all geometrical factors are equal, $L_i = \frac{1}{3}$, and Eq. (5.8) reduces to Eq. (5.5).

For the purposes of this thesis, nanodisks are of interest which are characterized by their radius r and height h . These can be approximated as oblate spheroids ($b_x = b_y = r > b_z = h/2$) with the incident field polarized parallel to the xy -plane. Figure 5.3 shows extinction spectra for such ellipsoids. The resonance condition (Eq. (5.8)) indicates that the resonance frequency depends on the magnitude of the appropriate geometrical factor L_i , which in turn depends on the relative sizes of the semiaxes. In this case, the relevant factor is $L_x = L_y = L_{xy}$. As the radius r increases, L_{xy} decreases, meaning that for an oblate ellipsoid $L_{xy} < 1/3$ (Eq. (5.7)). Consequently, the resonance condition (Eq. (5.8)) shifts to more negative values of $\text{Re } \varepsilon(\omega)$ as the aspect ratio $b_{xy}/b_z = 2r/h$ increases. This corresponds to lower photon energy, or equivalently, a longer wavelength (Fig. 5.2), consistent with the results in Fig. 5.3.




Figure 5.3: LSPRs for Pd and Ag oblate ellipsoids with height 20 nm, described by the quasistatic approximation (transparent) and MLWA (opaque). The extinction is normalized with the projected disk area ($\propto r^2$) and the quasistatic approximation for Ag is further scaled down according to the legend.

Although the discussion above is not strictly valid for many real nanoplasmonic systems due to larger particle sizes, complex geometries, and non-negligible $\text{Im } \epsilon(\omega)$, they provide a useful foundation for understanding the general behavior of the LSPR. In particular, they show that the LSPR peak shifts to longer wavelengths as the aspect ratio increases.

5.2.2 The modified long-wavelength approximation

For large NPs, the quasistatic approximation becomes inaccurate because the internal electric field within the particle becomes inhomogeneous in response to the external field. This effect is captured by the modified long-wavelength approximation (MLWA), which extends the quasistatic approximation by including two correction terms that account for dynamic depolarization and radiation damping [133].

According to the MLWA, the polarizability is given by

$$\alpha_{\text{MLWA}} = \frac{\alpha_{\text{QS}}}{1 - \frac{k^3 i}{6\pi} \alpha_{\text{QS}} - \frac{k^2}{4\pi b} \alpha_{\text{QS}}}, \quad (5.9)$$

where k is the wavenumber in the medium and b is the semiaxis of the particle parallel to the incident electric field. The k^3 term represents radiation damping, arising from the electromagnetic radiation emitted by oscillating dipoles, while the k^2 term accounts for dynamic depolarization caused by phase differences between dipoles at different parts of the particle [134].

To illustrate the effect of these corrections, Fig. 5.3 compares results from both the quasistatic approximation and the MLWA for Pd and Ag NPs. For the smallest NP considered, a disk-like ellipsoid with a radius of 25 nm, both methods produce similar results. For larger NPs, however, the quasistatic approximation underestimates the size-dependent peak shift and overestimates the amplitude. This is particularly severe in the case of Ag, leading to unphysically sharp peaks that must be scaled down by one to two orders of magnitude to appear on the same y-scale as the MLWA spectra.

5.3 Surface lattice resonances

When NPs are arranged in an ordered array, they can produce constructive in-plane diffraction of light. This occurs at Rayleigh anomalies (RAs), which are determined by the array geometry. For a two-dimensional square array illuminated at normal incidence, the RAs occur at wavelengths [135]

$$\lambda_{\text{RA}}^{(i,j)} = \frac{na}{\sqrt{i^2 + j^2}},$$

where n is the refractive index of the surrounding medium, a is the array periodicity and i, j are indices corresponding to the RA order. When a RA and a LSPR are spectrally close, they can couple to form a surface lattice resonance (SLR) [29, 133]. Unlike LSPRs, which are localized to individual particles, SLRs are associated with enhanced field intensities between particles and sharp features in the extinction spectra.

SLRs can be described using the coupled dipole approximation. In this framework, each NP is represented as an electric dipole subject to both the incident field and the fields generated by all other dipoles, resulting in a modified polarizability [29]

$$\alpha_{\text{SLR}} = \frac{1}{\alpha_{\text{LSPR}}^{-1} - S},$$

for normally incident light polarized along one of the symmetry axes of the array. Here, S is a geometry factor given by

$$S = \frac{1}{4\pi} \sum_j \exp(ikr_j) \left[\frac{(1 - ikr_j)(3 \cos^2 \theta_j - 1)}{r_j^3} + \frac{k^2 \sin^2 \theta_j}{r_j} \right],$$

where r_j is the distance from the central particle to particle j , and θ_j is the angle between \mathbf{r}_j and the dipole moment of particle j . The single-particle polarizability α_{LSPR} can, for instance, be taken as α_{MLWA} (Eq. (5.9)).

Figure 5.4 illustrates how SLRs emerge from the coupling between LSPRs and RAs. The RAs introduce indents in the extinction spectra. When a RA lies close to the LSPR,




Figure 5.4: SLRs for an arrays with pitch $a = 600$ nm of Pd or Ag oblate ellipsoids with height 20 nm in vacuum and the underlying RAs (vertical gray lines) and LSPRs (described with MLWA, transparent colored lines). The extinction is normalized with the projected disk area ($\propto r^2$).

their coupling produces a sharp SLR feature at a slightly longer wavelength. This effect is significantly stronger for Ag than for Pd, due to the previously discussed lossy nature of Pd (Sect. 5.1). If the LSPR is pushed away from the RA by modifying the NP geometry, the associated sharpening and amplification largely disappear.

Electrodynamic simulations

Thunderbolts and lightning, very, very frightening

Queen

The approximative methods outlined in Chapter 5 provide qualitative insight into plasmonic phenomena. To achieve quantitative agreement for realistic systems with nanoparticle (NP) sizes on the order of 100 nm as well as more complex geometries, including substrates and coatings, continuum electrodynamic simulations at the macroscopic scale are required. The atomistic material properties enter through the dielectric function (DF) (Fig. 6.1), as discussed in Sect. 5.1. This chapter presents the finite-difference time-domain (FDTD) framework for continuum electrodynamic simulations, beginning from Maxwell's equations.

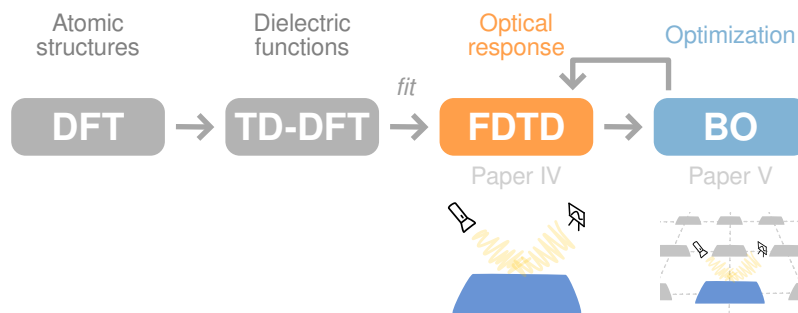


Figure 6.1: Schematic illustration of the modeling of optical properties. Density-functional theory (DFT) and TDDFT calculations (outside the scope of this thesis) provide the intrinsic optical materials properties used in FDTD simulations to study, and optimize via Bayesian optimization (BO) (introduced in Chapter 7), the optical response of alloy nanostructures.

6.1 Maxwell's equations

The time-dependent macroscopic Maxwell's equations in differential form, excluding electric and magnetic current sources, are given by [136, Chapter 3]

$$\frac{\partial \mathbf{D}}{\partial t} = \nabla \times \mathbf{H} - \mathbf{J} \quad (6.1)$$

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E} - \mathbf{M} \quad (6.2)$$

$$\nabla \cdot \mathbf{D} = 0 \quad (6.3)$$

$$\nabla \cdot \mathbf{B} = 0 \quad (6.4)$$

where \mathbf{E} is the electric field, \mathbf{D} the electric displacement field, \mathbf{H} the magnetic field, \mathbf{B} the magnetic flux density, \mathbf{J} the electric current density and \mathbf{M} the magnetic current density. These are connected through the constitutive relations

$$\mathbf{D} = \varepsilon \mathbf{E} \quad (6.5)$$

$$\mathbf{B} = \mu \mathbf{H}, \quad (6.6)$$

where ε and μ are the electric permittivity and magnetic permeability, respectively. The current densities can include both independent sources (\mathbf{J}_s and \mathbf{M}_s) and nondispersive losses represented by the electric and magnetic conductivities (σ and σ^*),

$$\mathbf{J} = \mathbf{J}_s + \sigma \mathbf{E} \quad (6.7)$$

$$\mathbf{M} = \mathbf{M}_s + \sigma^* \mathbf{H}. \quad (6.8)$$

The material properties enter through ε , μ , σ , and σ^* . In this thesis, only non-magnetic materials are considered ($\mu = 1$, $\sigma^* = 0$), and the relevant material-specific property is the permittivity, or DF, ε . The DF of an arbitrary alloy is generally unknown but can be calculated using time-dependent density-functional theory (TDDFT), as outlined in Ref. [130].

6.2 Finite-difference time-domain method

The finite-difference time-domain (FDTD) method solves the six coupled partial differential equations given by Eq. (6.1) and Eq. (6.2) by explicit time evolution. The method is typically implemented on a Yee grid [137], where the \mathbf{E} and \mathbf{H} components are positioned on a staggered three-dimensional grid offset by half a grid cell. Each \mathbf{E} component is then surrounded by four perpendicular \mathbf{H} components and vice versa, as illustrated in Fig. 6.2. This arrangement enables central-difference derivatives with second-order accuracy, enforces continuity of tangential \mathbf{E} and normal \mathbf{H} at interfaces, and implicitly




Figure 6.2: Illustration of the Yee grid.

satisfies Eq. (6.3) and Eq. (6.4) [136, Chapter 3]. The Yee grid is typically cubic, with spatial resolution $\Delta x = \Delta y = \Delta z = \Delta$ much smaller than the wavelength [137].

Similarly, the evolution in time follows a leapfrog scheme in which the \mathbf{E} and \mathbf{H} components are offset by half a time step. Consequently, the time evolution is fully explicit, and the time derivatives are approximated using central differences with second-order accuracy. Numerical stability requires that the time step satisfies the Courant condition [136, Chapter 4]

$$\Delta t < \frac{\Delta}{c\sqrt{3}} \quad (\text{in 3D}), \quad (6.9)$$

where c is the speed of light.

A FDTD simulation is performed in a finite computational cell bounded by suitable boundary conditions. Systems periodic in one or more directions naturally use periodic boundary conditions in those directions. For non-periodic boundaries, a perfectly matched layer is commonly applied. The perfectly matched layer is a fictitious material designed to absorb all incident fields without reflection [136, Chapter 7], effectively simulating an unbounded domain.

In the following, the implementation of a FDTD method on a Yee grid is outlined based on Ref. [136, Chapter 3]. This is first done for materials with non-dispersive optical properties and later extended to dispersive media.

The notation is defined as follows. Any function u of space and time is written as

$$u(i\Delta x, j\Delta y, k\Delta z, n\Delta t) = u_{i,j,k}^n$$

with spatial and temporal derivatives approximated by central differences

$$\begin{aligned} \frac{\partial u_{i,j,k}^n}{\partial x} &= \frac{u_{i+1/2,j,k}^n - u_{i-1/2,j,k}^n}{\Delta x} + \mathcal{O}[(\Delta x)^2] \\ \frac{\partial u_{i,j,k}^n}{\partial t} &= \frac{u_{i,j,k}^{n+1/2} - u_{i,j,k}^{n-1/2}}{\Delta t} + \mathcal{O}[(\Delta t)^2]. \end{aligned}$$

The first of the six partial differential equations, Eq. (6.1) in x -direction, is

$$\frac{\partial E_x}{\partial t} = \frac{1}{\varepsilon} \left[\frac{\partial H_z}{\partial y} - \frac{\partial H_y}{\partial z} - J_x \right], \quad (6.10)$$

where Eq. (6.5) has been used to express D_x in terms of E_x . Applying central differences gives

$$\begin{aligned} \frac{E_x|_{i,j+1/2,k+1/2}^{n+1/2} - E_x|_{i,j+1/2,k+1/2}^{n-1/2}}{\Delta t} = \\ = \frac{1}{\varepsilon} \left[\frac{H_z|_{i,j+1,k+1/2}^n - H_z|_{i,j,k+1/2}^n}{\Delta y} - \frac{H_y|_{i,j+1/2,k+1}^n - H_y|_{i,j+1/2,k}^n}{\Delta z} - J_x|_{i,j+1/2,k+1/2}^n \right]. \end{aligned}$$

At time step $n + 1/2$, the term $E_x|_{i,j+1/2,k+1/2}^{n+1/2}$ can be updated since all other terms correspond to earlier time steps. The same procedure applies to the remaining five equations, yielding \mathbf{E} at half-integer time steps ($n + 1/2$) and \mathbf{H} at integer time steps ($n + 1$).

For dispersive materials, i.e., materials with a frequency-dependent permittivity (DF), the FDTD algorithm must be modified. This can be achieved by the auxiliary differential equation method [136, Chapter 9]. The electric displacement field can be written as

$$\mathbf{D} = \varepsilon \mathbf{E} = \varepsilon_0 \mathbf{E} + \mathbf{P}, \quad (6.11)$$

where \mathbf{P} is the induced polarization. Given an analytical form of the DF $\varepsilon(\omega)$, such as the Drude–Lorentz model (Sect. 5.1), the polarization and electric fields are related in frequency space by

$$\mathbf{P}(\omega) = [\varepsilon(\omega) - \varepsilon_0] \mathbf{E}(\omega). \quad (6.12)$$

Applying an inverse Fourier transform yields a differential equation for the time evolution of \mathbf{P} , which is solved numerically with central differences alongside Maxwell's equations. Lastly, Eq. (6.10) becomes

$$\frac{\partial E_x}{\partial t} = \frac{1}{\varepsilon_0} \left[\frac{\partial H_z}{\partial y} - \frac{\partial H_y}{\partial z} - \frac{\partial P_x}{\partial t} - J_x \right], \quad (6.13)$$

where P_x is the x -component of the polarization obtained from the auxiliary differential equation, typically involving additional E_x terms.

Bayesian optimization

Whatever happens, I'll leave it all to chance

Queen

In materials design, coupling physical models, such as the electromagnetic simulations of Chapter 6, with efficient optimization methods accelerates the discovery of structures with targeted properties. Formally, optimization is the task of finding the parameters \mathbf{x} that maximize (or minimize) an objective (or loss or cost) function $f(\mathbf{x})$,

$$\max_{\mathbf{x} \in D} f(\mathbf{x}),$$

where \mathbf{x} is constrained to the domain D . Different optimization methods are suited for different problems, depending mainly on whether $f(\mathbf{x})$ is known, differentiable, and/or computationally expensive to evaluate.

When $f(\mathbf{x})$ is available in analytical form or through differentiable simulations, gradient-based methods are common. These methods are also applicable when $f(\mathbf{x})$ is expensive to evaluate but differentiable, as in structural relaxation. If the function is unknown or non-differentiable but inexpensive to compute, grid-search or, for larger domains, stochastic methods such Monte Carlo (MC) sampling (Sect. 3.4) are often suitable.

In many scientific and engineering applications, the objective function relies on computationally expensive methods such as finite-difference time-domain (FDTD) (Sect. 6.2) or density-functional theory (DFT) (Sect. 3.1). In such cases, it is essential to minimize the number of function evaluations while still exploring the design space effectively. A class of methods designed for this regime is Bayesian optimization (BO).

BO iteratively builds a probabilistic surrogate model of the objective function, typically a Gaussian process (GP), which captures both the estimated mean value and uncertainty of $f(x)$ (Fig. 7.1). A prior encodes initial beliefs about the behavior of the function, which is combined with the observed data to predict the mean and variance over the entire domain. An acquisition function guides the selection of the next data point by balancing exploration of regions with large uncertainty and exploitation of areas with high probability of finding a new maximum. As a result, BO typically reaches convergence with far fewer evaluations than conventional optimization methods. It is well suited for objective functions that are expensive to evaluate and/or unknown but accessible through (possibly noisy) observations over complex design spaces. In the following, the BO framework is presented following Ref. 138.

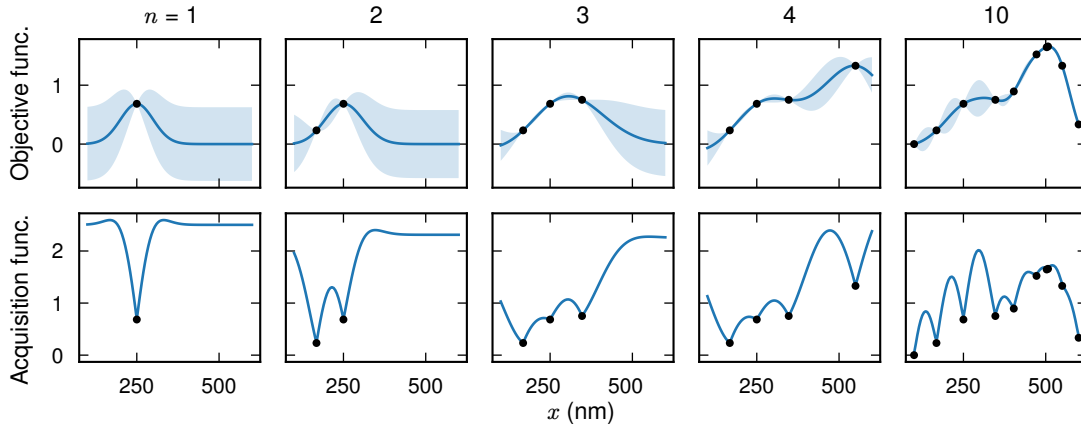


Figure 7.1: BO in one dimension. In this example, the sensor performance (objective function $f(x)$) of a nanodisk array is optimized as a function of the array pitch (x). New data points are iteratively added based on the maxima of the (upper confidence bound) acquisition function in the previous step, effectively exploring the entire interval.

7.1 Bayes' theorem

The mathematical basis of BO is Bayes' theorem,

$$P(M | E) \propto P(E | M) P(M),$$

which states that the posterior probability of a model M given evidence E is proportional to the likelihood of E given M multiplied by the prior probability of M . Applied to the optimization problem, this becomes

$$P(f | y_{1:n}, x_{1:n}) \propto P(y_{1:n}, x_{1:n} | f) P(f), \quad (7.1)$$

where f is the unknown objective function, and $\mathbf{y}_{1:n} = \{y_1, \dots, y_n\}$ are the observations at the sampled points $\mathbf{x}_{1:n} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. Here,

- $P(f)$ is the prior probability of a possible objective function, encoding beliefs about the function space such as smoothness, periodicity, and scale.
- $P(\mathbf{y}_{1:n}, \mathbf{x}_{1:n} | f)$ is the likelihood, quantifying how probable the observed data are given the function f .
- $P(f | \mathbf{y}_{1:n}, \mathbf{x}_{1:n})$ is the posterior probability of f given the observed data, representing updated beliefs about f after combining the prior and observed data via the likelihood.

In practice, this is often implemented via GP surrogate functions.

7.2 Gaussian processes

A Gaussian process (GP) is an extension of the Gaussian distribution from random variables to random functions. A Gaussian distribution is a distribution over a random variable x through its mean and variance,

$$x \sim \mathcal{N}(\mu, \sigma^2).$$

Analogously, a GP is a distribution over functions, specified by a mean function $m(\mathbf{x})$ and a covariance (kernel) function $k(\mathbf{x}, \mathbf{x}')$,

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')).$$

In BO, the GP serves as a probabilistic surrogate for the true objective function. Rather than returning a single value $f(\mathbf{x})$, it yields the mean and variance that describe the posterior distribution of $f(\mathbf{x})$ given the observed data.

In most cases, it is sufficient to set the prior mean function to zero, $m(\mathbf{x}) = 0$. The covariance function, however, determines how function values at different points are correlated and must be chosen carefully. A common choice is the radial basis function (RBF) or squared exponential kernel,

$$k_{\text{RBF}}(\mathbf{x}_i, \mathbf{x}_j) = \sigma_{\text{RBF}}^2 \exp\left(-\frac{1}{2l^2} \|\mathbf{x}_i - \mathbf{x}_j\|^2\right),$$

parametrized by the length scale l controlling how rapidly correlations decay with distance, and the variance σ_{RBF}^2 setting the amplitude of the function.

In each BO iteration, the GP predicts the mean and variance of the objective function across the domain (Fig. 7.1). This is achieved at iteration $n + 1$ based on the following derivations. An observation y_i is defined as

$$y_i = f(\mathbf{x}_i) + \epsilon_i,$$

where the difference between the true objective function and the observation is represented by Gaussian noise, $\epsilon_i \sim \mathcal{N}(0, \sigma_{\text{noise}}^2)$. Based on the n observations $y_{1:n}$ at locations $\mathbf{x}_{1:n}$, the kernel matrix is calculated

$$\mathbf{K} = \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & \cdots & k(\mathbf{x}_1, \mathbf{x}_n) \\ \vdots & \ddots & \vdots \\ k(\mathbf{x}_n, \mathbf{x}_1) & \cdots & k(\mathbf{x}_n, \mathbf{x}_n) \end{bmatrix} + \sigma_{\text{noise}}^2 \mathbf{I}.$$

From the properties of GPs it follows that the previously observed values $y_{1:n}$ and the predicted f_{n+1} at a new point \mathbf{x}_{n+1} are jointly Gaussian [139],

$$\begin{bmatrix} y_{1:n} \\ f_{n+1} \end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} \mathbf{K} & \mathbf{k} \\ \mathbf{k}^T & k(\mathbf{x}_{n+1}, \mathbf{x}_{n+1}) \end{bmatrix}\right),$$

with

$$\mathbf{k} = [k(\mathbf{x}_{n+1}, \mathbf{x}_1) \quad k(\mathbf{x}_{n+1}, \mathbf{x}_2) \quad \cdots \quad k(\mathbf{x}_{n+1}, \mathbf{x}_n)].$$

The predictive posterior distribution can then be derived as [139]

$$P(f_{n+1} | y_{1:n}, \mathbf{x}_{1:n+1}) = \mathcal{N}(\mu_n(\mathbf{x}_{n+1}), \sigma_n^2(\mathbf{x}_{n+1})),$$

where

$$\begin{aligned} \mu_n(\mathbf{x}_{n+1}) &= \mathbf{k}^T \mathbf{K}^{-1} y_{1:n}, \\ \sigma_n^2(\mathbf{x}_{n+1}) &= k(\mathbf{x}_{n+1}, \mathbf{x}_{n+1}) - \mathbf{k}^T \mathbf{K}^{-1} \mathbf{k}. \end{aligned} \tag{7.2}$$

In the sense of Bayes' theorem (Eq. (7.1)), the prior is the structure of the GP and the likelihood arises by adding the data and noise.

7.3 Maximizing the acquisition function

Computable expressions are now available for the predicted mean and variance at any candidate point \mathbf{x}_{n+1} (Eq. (7.2)), but it is still not clear how to select \mathbf{x}_{n+1} . To address this, an acquisition function $u(\mathbf{x})$ is introduced, designed to adopt high values in regions with high uncertainty and/or large predicted mean. The acquisition function depends on the GP predictions $\mu_n(\mathbf{x})$ and $\sigma_n(\mathbf{x})$ and must be inexpensive to evaluate, in contrast to the objective function, so that it can be efficiently maximized.

A simple and widely used choice is the upper confidence bound (Fig. 7.1),

$$u_{\text{UCB}}(\mathbf{x}) = \mu_n(\mathbf{x}) + \beta \sigma_n(\mathbf{x}),$$

which balances the predicted mean and uncertainty via the hyperparameter β . By tuning β , the balance between exploitation of regions with high predicted mean (low β) and exploration of regions with high uncertainty (high β) can be controlled. Maximizing the acquisition function with a suitable optimization method identifies the next observation point

$$\mathbf{x}_{n+1} = \arg \max_{\mathbf{x} \in D} u(\mathbf{x}).$$

Optical properties of Pd nanoalloys for H₂ sensing

One shaft of light that shows the way

Queen

The H-induced optical shift of Pd nanoalloys is the mechanism behind plasmonic H₂ sensing. Building on the theoretical background in Chapter 5 and the computational framework introduced in Chapter 6 and Chapter 7, this chapter presents the main results concerning the intrinsic optical properties of the Pd–Au–H system and their manifestation in single nanodisks and periodic nanodisk arrays.

8.1 Dielectric functions of the Pd–Au–H system

In order to simulate the optical response of Pd-based nanoalloys during hydrogenation, the DF of each relevant composition needs to be known. While both experimental and theoretical records of the DFs of Pd [140], PdH [141–143] and PdAu [130] are available in the literature, previous reports for the full Pd–Au–H system were not found.

In Paper IV, DFs for the Pd–Au–H system were therefore calculated using TDDFT (Fig. 8.1a). It is found that adding H or Au to Pd has a similar effect on the DF, as can be rationalized based on the electronic configurations. The main electronic difference between Pd and Au is that Pd belongs to an earlier column of the periodic table, resulting in one less valence electron compared to Au, which leaves its d-band partially unfilled, as discussed in Sect. 5.1. By absorbing H, Pd gains an electron which effectively makes it more similar to Au.

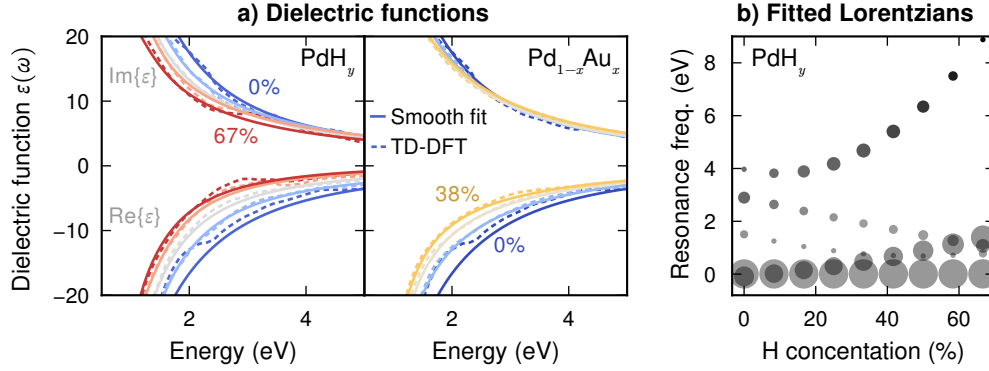


Figure 8.1: (a) Dielectric functions (DFs) of the Pd–H and Pd–Au systems calculated from time-dependent density-functional theory (TDDFT) and fitted using a composition-parametrized (smooth) Lorentzian representation. (b) The individual Lorentzians of the smooth DFs for the Pd–H system with the amplitude qualitatively represented by the marker size and the broadening by the marker opacity (with lower opacity for larger broadening).

The calculated DFs contain notable fine structure (Fig. 8.1a). Generally, finite-difference time-domain (FDTD) simulations require fitting the DF to some functional form. In Paper IV, each DF is individually fitted to a Drude-Lorentz representation (Eq. (5.2)), which preserves most of the fine structure. This set of DFs will be referred to as the as-calculated DFs. Paper V also employs a Drude-Lorentz representation, but here each feature is parametrized by the composition. This approach allows for a continuous representation of the DF over the relevant composition range, which removes most of the fine structure (Fig. 8.1a), and will therefore be referred to as the smooth DFs. Note that the resulting features (illustrated in Fig. 8.1b for PdH_y) do not necessarily represent the underlying physical mechanisms of the Drude and Lorentz features.

8.2 H₂ sensitivity in Pd-alloy single disks

In Paper IV, the optical H sensitivity of Pd-alloy single disk is studied based on the as-calculated DFs. Figure 8.2a shows single disk extinction spectra, for two nanodisk geometries and varying H concentration. From such simulations, the optical sensitivity (Eq. (1.1)) is calculated from the slope of the (typically linear) shift in peak position as a function of H concentration. Fig. 8.2b shows the resulting optical sensitivity for Pd and Pd₇₅Au₂₅ for a wide range of nanodisk geometries. It is found that the optical sensitivity primarily depends on the disk diameter, while the Au concentration and disk height have a limited effect in the intervals considered. Two distinct regimes can be identified. For disks with diameters below approximately 100 nm, the optical sensitivity decreases linearly with the diameter while for larger disks, the optical sensitivity remains constant at approximately 200 nm/c_H.




Figure 8.2: (a) Extinction spectra for PdH (with varying H concentration as indicated by the color) and Pd₇₅Au₂₅ (black dashed line) for two disk geometries based on as-calculated DFs. (b) Optical sensitivity for Pd and Pd₇₅Au₂₅ as a function of the disk diameter for different disk heights h based on as-calculated (opaque) and smooth (transparent) DFs, as well as the corresponding experimentally measured values (plus-signs) [15]. (c) Comparison of extinction spectra obtained using the as-calculated and smooth DFs.

The two regimes are the result of a double peak phenomenon prominent for smaller nanodisks, as illustrated in the lower panel of Fig. 8.2a. Studying the optical response of Pd vs. PdH nanodisks as a function of the disk aspect ratio provides further insight (Fig. 8.3). For Pd, the (single) feature shifts in energy with the aspect ratio, as is char-




Figure 8.3: Double peak phenomenon observed for PdH (compared to the single peak of Pd) using the as-calculated DFs and the associated avoided crossing behavior following the peak positions as a function of the aspect ratio. The high-energy feature in PdH corresponds to a feature in the DF, shown to the right. The plus and minus signs indicates the sign of the optical sensitivity in that region.

acteristic for a localized surface plasmon resonance (LSPR). For PdH, the low-energy feature shifts in a similar way for large aspect ratios, while the high-energy feature remains constant in energy, indicating plasmonic nature for the former while the latter has the characteristics of an interband transition. This is further supported by the fact that the high-energy feature is visible in the DF. As the two features become spectrally close, the trends change and the high-energy feature begins shifting to higher energies while the low-energy feature flattens out. This is indicative of coupling between the features causing an avoided crossing [131, 132]. In Fig. 8.3b the obtained signs of the optical sensitivity in particular regions is indicated by plus and minus signs. For large aspect ratios, before the onset of the avoided crossing, the low-energy feature dominates and the sensitivity is positive and constant. As the aspect ratio decreases, the high-energy feature of PdH will become dominant, causing a sudden drop to a negative sensitivity. Further decrease of the aspect ratio eventually causes the features to cross, resulting in zero sensitivity, followed by an increasing positive sensitivity.

In contrast, experimental measurements on random nanodisk arrays exhibit a monotonic increase in optical sensitivity as a function of the (average) disk diameter, approaching the constant sensitivity trend as the upper limit. The difference between the experimental and simulated sensitivities can be attributed to several broadening channels present in the experimental measurements, such as uneven hydrogenation [144, 145], size distribution, grain boundaries [146] and other defects. Such effects are hard to explicitly account for in simulations, but by recalculating the optical sensitivity with the smooth DFs, one obtains a trend that is very similar to the experimentally observed one (Fig. 8.2b-c). This can only be explained by the overall smoothing and broadening of the extinction spectra upon removal of most of the fine structure from the DFs. It is thus probable that the imperfections of the real samples leads to similar smoothing, which appears to lead to an effective averaging between the two trends obtained from simulations of pristine single disks using the as-calculated DFs.

8.3 Optimized nanoalloy arrays for H₂ sensing

The findings in the previous section indicate that alloying with Au does not significantly alter the optical sensitivity. However, from the discussion in Part I, it is clear that Au has a huge impact on the thermodynamical sensitivity (Eq. (1.1)). Therefore, studies of the optical sensitivity alone have limited significance.

In Paper V, the optical response of periodic nanodisk arrays is optimized using simulations based on the smooth DFs and subsequently verified by experimental measurements. The nanodisks are coated with poly(methyl methacrylate), matching the index of the silica substrate, to support surface lattice resonance (SLR) formation. A Bayesian optimization (BO) framework is developed to optimize the sensor geometry and composition at predefined target pressures (1, 10 or 100 mbar). The thermodynamic sensitivity

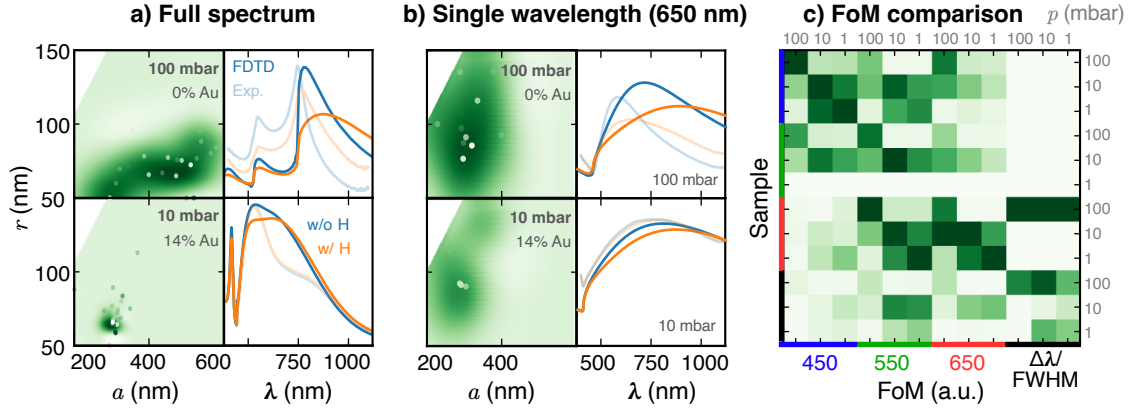


Figure 8.4: Optimizing geometries and compositions for (a) full spectrum ($\Delta\lambda/\text{FWHM}$) and (b) single wavelength (extinction change at 650 nm) sensors. Each panel shows the predicted figure of merit (FoM) landscape close to the optima to the left and the corresponding simulated (opaque) and measured (transparent) spectra for target pressures of 100 mbar (top) and 10 mbar (bottom). (c) Comparison of all considered figure of merits (FoMs) for all samples, where each column corresponds to one FoM at a specific target pressure and the largest value of that column is assigned the darkest color. Successful optimization is thus indicated by dark colors on the diagonal.

is explicitly included by relating the target pressure to a specific H concentration based on computational data from Ref. [107]. The optical sensitivity is extended such that it quantifies either the peak shift, referred to as a full spectrum sensor, or the shift in extinction at a single wavelength (450, 550 or 650 nm). In the former case, the BO maximizes the peak shift $\Delta\lambda$ over full width at half maximum (FWHM), known as a quality factor, and in the latter, the extinction shift at the selected wavelength. These quantities are the FoMs, or objective functions, of the BO.

Examples of predicted FoM landscapes as well as extinction spectra corresponding to the optimized geometries for the full spectrum (Fig. 8.4a) and single wavelength (650 nm, Fig. 8.4b) show that the BO navigates the five dimensional search space well and is able to find both wide and narrow optima. The experimentally measured extinction spectra are in qualitative agreement with the simulated spectra, but there are notable shifts which affect the measured FoM and ultimately limits the success of this approach. Qualitative rather than quantitative agreement between FDTD and experimental measurement is often observed, due to two main sources of discrepancy. First, the intrinsic optical properties represented by the DFs will generally differ between the TDDFT calculations of pristine bulk material and the real samples due to approximations in the calculations as well as grain-boundaries and other defects in the samples. While the discussion in the previous section indicates that the smooth DFs improve the qualitative agreement with experiments, resulting in similar trends for the optical sen-

sitivity, the quantitative agreement for a specific sample is not evaluated. Second, the geometrical description of the samples is associated with some level of uncertainty inherent to the fabrication process and due to the limited resolution in the simulations. The latter is particularly severe in this work due to annealing required to form alloys, which is associated with a noticeable change in the geometry, not yet accounted for.

Cross-referencing all considered FoMs on all samples shows that the majority of samples perform best at their indented target (Fig. 8.4c). In particular, all single wavelength sensors perform well. The magnitude of the change in extinction primarily depends on the magnitude of the extinction, which generally favors low packing densities. Such geometries are associated with broad features, and as a result, the predicted optima are less sensitive to differences between experiment and simulation. The full spectrum sensors benefit from sharp features which naturally results in more narrow optima. The full spectrum sensors for the two lower target pressures, which exhibit very narrow optima, completely miss their targets in the experimental verification. To reach the full potential of this approach, better agreement between experiments and simulations is thus necessary.