Alexander Egger, BSc.

Surface Polymorph Prediction using Bayesian Learning: To Monolayer Coverage and Beyond

MASTERARBEIT

zur Erlangung des akademischen Grades

Diplom-Ingenieur

Masterstudium Technische Physik

eingereicht an der

Technischen Universität Graz

Betreuer

Dr.techn. Oliver T. Hofmann

Institut für Festkörperphysik

Graz, Februar 2019

EIDESSTATTLICHE ERKLÄRUNG

Ich erkläre an Eides statt, dass ich die vorliegende Arbeit selbstständig verfasst, andere als die angegebenen Quellen/Hilfsmittel nicht benutzt, und die den benutzten Quellen wörtlich und inhaltlich entnommenen Stellen als solche kenntlich gemacht habe. Das in TUGRAZonline hochgeladene Textdokument ist mit der vorliegenden Masterarbeit identisch.

Datum

Unterschrift

Danksagung

Ich möchte an dieser Stelle allen danken, die diese Arbeit möglich gemacht haben. Insbesondere möchte ich natürlich meinem Betreuer, Oliver Hofmann, danken der stets ein offenes Ohr für meine Fragen hatte und gleichzeitig sicherstellte, dass diese Arbeit auch zu einem Ende gelangt.

Weiters möchte ich mich bei meinen Kollegen Andi, Lukas, Michael und Simon bedanken. Sie haben viele Grundsteine und Werkzeuge erschaffen, auf denen ich meine Arbeit aufbauen durfte. Aber vor allem haben sie diese Arbeit durch zahlreiche konstruktive Diskussionen und schlicht durch eine wunderbare Arbeitsatmosphäre ermöglicht. Selbiger Dank gilt auch allen weiteren Kollegen und Kolleginnen des Advanced Materials Modeling Teams mit welchen ich eine wundervolle Zeit verbringen durfte.

Für ihre unermüdliche Rechenleistung zu jeder Tages und Nachtzeit möchte ich mich bei den CPUs des Vienna Scientific Clusters sowie des Argonne National Laboratory (MIRA) bedanken. In diesem Zusammenhang sei auch den Entwicklern und Entwicklerinnen etlicher open-source Projekte gedankt. Ohne tools wie Python [1], Numpy [2], matplotlib [3], ASE [4], git, Ubuntu, LaTeX [5] u.v.m. wäre ein effizientes Arbeiten kaum denkbar.

Eine Diplomarbeit ist bekanntlich nicht möglich ohne Diplomstudium und dieses wiederum wäre kaum möglich gewesen ohne jene Freunde, die mich durch das Studium begleitet haben und mir in dieser Zeit so viele schöne Momente geschenkt haben. Schlussendlich möchte ich mich bei meiner Familie bedanken die mir dieses Studium ermöglicht hat und ohne deren Unterstützung diese Arbeit nie zustande gekommen wäre.

Danke!

Abstract

Surface Polymorph Prediction using Bayesian Learning: To Monolayer Coverage and Beyond

Alexander Egger
Institute of Solid State Physics, Graz University of Technology

Finding and synthesizing materials with tailor-made properties are central goals of material science. Interfaces are an especially interesting area of research. In their vicinity, materials can form vastly different polymorphs than in an extended bulk material. With those unique polymorphs, exceptional physical properties arise. To find out which of these possible polymorphs actually form, one needs to calculate the Gibbs free energy of each with high accuracy. This level of accuracy necessitates computationally expensive dispersion-corrected density functional theory (DFT). At the same time, the number of possible polymorphs is basically infinite. Consequently, conducting one DFT simulation for each possible polymorph is impossible.

SAMPLE is a machine-learning-based approach to solve this problem for commensurate phases. By combining physically motivated coarse graining of the search space with Bayesian linear regression we can predict the Gibbs free energy of an exhaustive set of polymorphs at reasonable computational cost while preserving the level of accuracy of dispersion-corrected DFT.

In this work, the SAMPLE approach is applied to close-packed polymorphs of increasing coverage in order to study potential structural transitions. Above the full monolayer coverage of flat-lying (face-on oriented) molecules, the question arises whether adsorbate molecules either switch to an upright-standing orientation (edge-on) or if they start to form a second layer.

To resolve this question, we first implement and benchmark multiple design of experiment schemes in order to optimize the training set selection such that the desired level of accuracy can be achieved with as few DFT calculations as possible. We then employ SAMPLE to find all polymorphs of the first adsorbate layer up to high packing density and arrange them in a surface phase diagram. Based on this information, selected bilayer polymorphs are built. We find that densely-packed monolayers of upright-standing molecules have a lower Gibbs free energy than bilayers with the same total coverage. Therefore, we predict a coverage-driven phase transition from flat-lying to upright-standing adsorbate molecules in the first monolayer.

Finally, we calculate the vibrational spectrum and molecular orbital density of states (MODOS) of predicted polymorphs. We find that – in contrast to current knowledge – no charge transfer to the second adsorbate layer occurs. Consequently, the experimentally-observed singly-charged species does not correspond to TCNE molecules in the second adsorbate layer but instead to upright-standing TCNE molecules in the first monolayer.

Kurzfassung

Struktursuche auf Oberflächen mit Bayes'schem Lernen

Alexander Egger
Institut für Festkörperphysik, Technische Universität Graz

Die Herstellung maßgeschneiderter Materialien ist eines der großen Ziele der Materialwissenschaft. Ein besonders interessantes Fachgebiet sind hierbei Oberflächen und Grenzschichten: In deren Nähe ist es möglich, dass Materialien vollkommen andere Polymorphe bilden als in einem ausgedehnten Festkörper. Diese Oberflächenpolymorphe können wiederum außergewöhnliche Materialeigenschaften mit sich bringen. Um vorherzusagen, welche Polymorphe sich bilden, muss die Gibbs-Energie jedes möglichen Polymorphs mit hoher Genauigkeit bestimmt werden. Allerdings erfordert diese Genauigkeit die Nutzung ressourcenintensive Rechenmethoden und gleichzeitig ist die Anzahl der möglichen Polymorphe im Grunde unendlich weshalb eine direkte Simulation aller möglichen Polymorphe unmöglich ist.

SAMPLE ist ein Ansatz, um dieses Problem für kommensurable Phasen mittels maschinellen Lernens zu lösen. Die Kombination von physikalisch motivierter Diskretisierung des Suchraums mit Bayes'scher Linearer Regression ermöglicht uns die Gibbs-Energie eines umfassenden Sets von Polymorphen mit vertretbaren Berechnungskosten und unter Erhaltung der Genauigkeit von dispersionkorrigierter Dichtefunktionaltheorie (DFT) zu bestimmen.

In dieser Arbeit wird SAMPLE auf dicht gepackte Polymorphe mit ansteigendem Bedeckungsgrad angewandt, um mögliche strukturelle Phasenübergänge zu untersuchen. Sobald Bedeckungsgrade über dem einer vollen Monolage liegender Molekülen erreicht werden, stellt sich die Frage, ob Adsorbatmoleküle entweder zu einer stehenden Orientierung wechseln oder, ob sie beginnen eine zweite Lage zu bilden.

Um diese Frage zu beantworten, implementieren und vergleichen wir zuerst mehrere Strategien aus dem Bereich der statistischen Versuchsplanung. Dies ermöglicht uns die notwendige Genauigkeit mit einer möglichst geringen Anzahl an Trainingsdaten – und damit möglichst geringen Berechnungskosten – zu erreichen. Anschließend bestimmen wir mit SAMPLE alle Polymorphe der ersten Adsorbatlage bis zu hohen Packungsdichten und visualisieren diese Polymorphe in einem Phasendiagramm. Basierend darauf, werden ausgewählte Doppelschichten erstellt und anhand der Gibbs-Energie mit dicht gepackten Monolagen verglichen. Dieser Vergleich zeigt, dass dicht gepackte Monolagen bestehend aus aufrecht stehenden TCNE-Molekülen energetisch günstiger sind als Doppelschichten des selben Bedeckungsgrades. Dementsprechend ist bei steigender Bedeckungsdichte ein Phasenübergang von liegenden zu stehenden TCNE-Molekülen in der ersten Adsorbatlage zu erwarten.

Abschließend simulieren wir die Schwingungsspektren und die Zustandsdichte der Molekülorbitale (MODOS) der vorhergesagten Phasen. Wir stellen fest, dass – im Widerspruch zum bisherigen Wissensstand – kein Ladungstransfer in die zweite Adsorbatlage stattfindet. Dementsprechend können auch die experimentell beobachteten Vibrationsfrequenzen einfach geladener TCNE-Moleküle nicht mit Molekülen in der zweiten Adsorbatlage erklärt werden. Vielmehr sind diese Vibrationsfrequenzen auf tatsächlich einfach geladene, stehende TCNE-Moleküle in der ersten Adsorbatlage zurückzuführen.

Preamble

As often in science, every single contribution is built upon the fundament of prior work. Therefore, I would like to clarify which parts of this study present techniques that were invented by my colleagues and other scientists and which parts present my own work.

Chapter 1: Introduction

This chapter mainly summarizes general scientific knowledge. Used resources are noted in the text. An exception to that is chapter 1.2.1 *Calculating Gibb's Free Energy for an Adsorbate Layer*. The model that is presented in this chapter is based on Ref. [6] as well as Ref. [7] but it is specifically tailored to the studied system by Lukas Hörmann and me.

Chapter 2: SAMPLE – Surface Adsorbate Prediction With Little Effort

While I contribute to the current implementation of SAMPLE – which consists of several modules of python code – the foundations of this tool were laid by my colleagues before I started this thesis. Consequently, the first publications [8, 9] concerning this approach were done without contributions from my side, while the latest one already contains contributions done by me [10]. The studied system of TCNE on Cu(111) was previously described by Veronika Obersteiner [11] and in a proof-of-concept of the SAMPLE approach by Michael Scherbela [12]. Therefore, I start from his *local adsorption geometries*. However, I additionally optimized these geometries, as different DFT settings are used in this study to ensure fully converged calculations.

Chapter 3: Increasing the Packing Density

This is the point where the core of my contribution starts. The test system, as well as all evaluations for the benchmarks, were created by me. The training set selections based on the A-, E-, G-, and V-optimality criteria are implemented solely by me, while the D-optimality criterion was already previously implemented by Michael Scherbela [12] but has been improved by me.

Chapter 4: Results: Predicted Polymorphs

Work and results presented in this chapter are done entirely by me.

Chapter 5: Comparison to Experiment

Apart from the presented experiment [13] itself, this chapter is based solely on my work.

Bring' vor, was wahr ist;
Schreib' so, dass es klar ist
Und verficht's, bis es mit dir gar ist!

LUDWIG BOLTZMANN
Vorlesungen über die Principe der Mechanik

Contents

1. Introduction

To obtain physical properties of a material it often is not sufficient to know its chemical composition, but also its exact structure is significant. As a simple example, diamond and graphite have the same chemical composition – they both consist solely of carbon – but they have vastly different properties. Such different forms of one material having different crystal structures are called polymorphs.

In a bulk crystal – like diamond and graphite – the basis of the crystal is infinitely repeated in all three dimensions. In adsorbate layers, however, that periodicity is broken in the direction perpendicular to the surface. While polymorph formation in the second- and higher adsorbate layers is mostly governed by intermolecular interactions, the adsorbates in the first layer are strongly influenced by the molecule-substrate interactions. Consequently, each adsorbate layer can form a specific crystal lattice and hence show distinct physical properties. An example of such behavior is given by pentacene adsorbate layers on boron nitride (see Figure 1.1). Right at the substrate surface pentacene adsorbs as a non-conducting wetting layer of flat-lying molecules. Additional deposition then leads to the formation of a layer of upright-standing, tilted molecules on top of the wetting layer. This first layer (1L) displays features of hopping transport and a mobility of $\mu = 1.6\ cm^2/Vs$. On top of that forms a second layer (2L) consisting of molecules which are less tilted then the ones in the first layer. This low tilting allows for a consecutive overlap of π-orbital leading to a high charge carrier mobility ($\mu \approx 3\ cm^2/Vs$) and a bandlike transport mechanism [15].

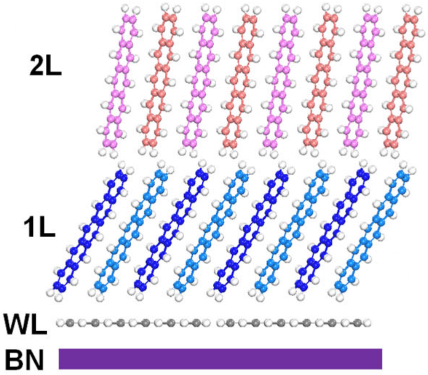

Figure 1.1.: Adsorbate layers of pentacene on boron nitride. Image taken from Ref. [15].

Predicting such complicated adsorbate layers directly from first principles would be a great step in the search for optimized materials. For now, however, making predictions for the system presented above is out of reach as the computational cost of the necessary simulations increases rapidly with the size of the studied system. Consequently, we choose a slightly smaller system: TCNE (tetracyanoethylene) on Cu(111). On Cu(111) a single, isolated TCNE molecule adsorbs in a flat-lying geometry with the four cyano groups slightly bend towards the surface. With further deposition the interaction between the adsorbed molecules becomes more important and at some point, Pauli repulsion will inhibit any closer packing of flat-lying TCNE molecules on the surface. At this point two fundamentally different scenarios are possible: A second adsorption layer could form on top of the first layer with no major reconstruction of the first layer. In an alternative scenario, the molecules of the first layer flip up and assume an upright-standing orientation (see Figure 1.2). The upright-standing TCNE molecules display a much weaker adsorption energy than the flat-lying ones. However, as their footprint is only approximately half as big, the packing density can be doubled without reducing the intermolecular distances.

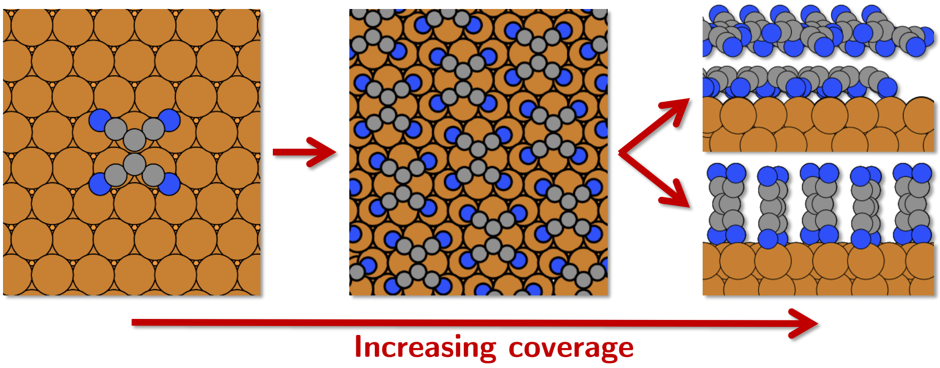

Figure 1.2.: Illustration of two possible scenarios upon deposition beyond the full monolayer coverage of flat-lying molecules.

To answer that question, we first have to find all polymorphs that the first layer can assume for a wide range of coverages. In the thermodynamic equilibrium, the polymorph that forms at a specific coverage is that one with the lowest Gibbs free energy (see chapter 1.2.1). The main contribution to the Gibbs free energy lies in the adsorption energy of each possible polymorph. As the energetic differences of competing polymorphs can be in the range of a few meV, we need sufficiently accurate methods like dispersion-corrected density functional theory (DFT) (see chapter 1.3.1). Unfortunately, such simulations are computationally far too costly to simulate every possible polymorph – even more so as there is in principle an infinite number of possible polymorphs. Therefore, we make use of sophisticated structure search algorithms. A short overview of structure search algorithms is given in chapter 1.1 and a more detailed look into the chosen method (SAMPLE) is then presented in chapter 2. The SAMPLE approach allows us to create an exhaustive set of possible polymorphs and predict their adsorption energies with an accuracy level comparable to DFT calculations, but with far less computational costs.

With these basics, we approach the core of this study: predicting the structure of an adsorbate layer of TCNE on Cu(111). We start from low adsorbate coverage, go via adsorbate layers of full monolayer coverage, and then apply the SAMPLE approach on upright-standing molecules.

It turns out that especially for those densely-packed systems the choice of the training data, which is used for the SAMPLE approach, has a major influence on the prediction accuracy. Selecting these training points randomly is not a sufficient strategy to achieve the necessary prediction quality at reasonably training set sizes. To tackle this problem, we implement different training set selections strategies. These strategies are then benchmark on a simplified version of the studied system in order to determine which one performs best (see chapter 3).

Having that, we switch back to the full surface system and fit SAMPLE's energy model using a D-optimally selected training set. Using this energy model, we then predict the adsorption energies of all discretized monolayer polymorphs (chapter 4.2). Based on energetically favorable monolayers we build bilayer polymorphs as described in chapter 4.3. In the following chapter 4.4 we combine all obtained polymorphs to create a surface phase diagram which allows us to determine how the surface structure changes with temperature and partial pressure of TCNE in the gas phase. With that knowledge, a comparison of the Gibbs free energy of a bilayer phase with that of a monolayer phase of equivalent coverage is conducted (chapter 4.4). This comparison is used to determine at which external conditions (pressure, temperature) the transition from a monolayer phase to a bilayer takes place.

To obtain additional information and enable a comparison to experimental data [13] we calculate the vibrational modes and frequencies of selected polymorphs. (see chapter 5) We find that the experimentally observed vibrational frequencies match those which are expected from our predictions. However, we also find that the formerly postulated existence of charge transfer to the second adsorbate layer conflicts with the vibrational frequencies that we calculated for adsorbates in the second layer. To confirm this finding, we additionally calculate the molecular orbital density of states (MODOS) for representative adsorption geometries. Again, we see that the experimentally observed singly charged adsorbate species cannot correspond to molecules in the second layer, but to upright-standing TCNE molecules in the first adsorbate layer.

1.1. Structure Search

As we want to find the polymorph that forms on a surface under specific conditions, we are in the field of (computational) structure search. As the term ’structure search’ is used for two different problems, let us clarify the differences of those:

Local Structure Search

The mathematical problem of finding a local extremum of a function is denoted as local optimization. A local extremum is defined as a point in parameter space where all partial derivatives are zero. Local structure search (also known as local geometry optimization) is a special case of that problem. In this case the parameter space is spanned by the coordinates of all atoms whose position shall be optimized and the target function is replaced by the potential energy surface (PES). A local minimum then equates to a geometry where each infinitely small deviation of one atom leads to an increase in energy.

Two characteristics of local optimization are important:
Firstly, potential energy surfaces usually have several local minima. Finding one of those does not give any information about the ’quality’ of that local minimum – i.e. the energy in that local minimum might be far from the global minimum. Secondly, the stability and efficiency of local optimization algorithms typically depend strongly on the starting point. Consequently, it is essential to start from a geometry that is already as close to the minimum as possible.

The most straightforward approach for local optimization is the gradient descent algorithm where the direction of each optimization step is given by the gradient of the potential energy surface. Even this simple algorithm is guaranteed to find the minimum, but there is no upper limit to the number of necessary steps. An improved method is the Broyden-Fletcher-Goldfarb-Shanno algorithm (BFGS) [16] belonging to the class of quasi-Newton methods. The basic idea is to locally approximate the potential energy surface with an harmonic function. This harmonic function is defined by a Hessian matrix that is estimated at the first step and then continuously improved using the local gradient. Within that approximation a one-dimensional line search is performed to find the next optimization point.

Global Structure Search

When we want to find the energetically most favorable polymorph a local structure search would only lead us to the next locally optimal polymorph. – i.e. a polymorph geometry that cannot be improved by infinitesimally small deviations. In the general case, however, the search space corresponds to a *non-convex* potential energy surface – i.e. there are multiple local minima which are separated by energy barriers. Finding the lowest local minimum – the *global minimum* – is known as global structure search. This task is fundamentally more challenging and less straight-forward than local structure search. For instance, there is in general no way to verify that a found minimum is indeed represents the global minimum of the search space.

Commonly used methods like Genetic Algorithms [17] or Neuronal Networks [18] which typically perform best with large training sets. Obtaining such large training sets of DFT calculations for surface systems is however hindered by the high computational costs. Therefore, we will use a more efficient method based on coarse-graining and Bayesian linear regression which will be explained more in detail later.

In this study both types of structure search are used. The main task of finding the polymorph that forms in thermodynamic equilibrium is a global structure search problem. In order to find this minimum efficiently we build up the guess polymorphs out of adsorbate molecules which are locally optimized. For that step the BFGS-algorithm is used as implemented in FHI-*aims* [19]. Furthermore, local structure search is used to find optimized geometries for the bilayer polymorphs constructed in chapter 4.3.

1.2. Physics of Surface Systems

When we want to find the best polymorph what we need first is a good definition for what is meant by 'best' polymorph. In the previous chapter we defined that rather mathematically as the global minimum of some non-convex function. It might seem apparent to use the adsorption energy of each polymorph for this no-convex function and, therefore, define the 'best' polymorph as that one which has the lowest adsorption energy. However, as the adsorbate layer is in contact with the surrounding gas, the situation becomes slightly more complex. In this case, TCNE molecules can adsorb to the surface or desorb to the gas phase. As the surrounding gas phase is several magnitudes larger than the studied surface, we can treat it as a reservoir of potential adsorbate molecules. Furthermore, in most experiments pressure and temperature are approximately constant (but different from zero!). Under these conditions the relevant thermodynamic potential is Gibbs free energy G. The polymorph that corresponds to the global minimum of Gibbs free energy constitutes the ground state of the studied system. In this study, we assume thermodynamic equilibrium, meaning that the adsorbate layer will always assume its ground state – even if the corresponding global minimum is surrounded by high energy barriers. In experiments this condition is challenging to achieve, but it can be aspired using e.g. hot-wall epitaxy. Now that we know we are searching for the polymorph of lowest

Gibbs free energy, we just need to know how to calculate the Gibbs free energy for the studied surface system.

1.2.1. Calculating Gibb's Free Energy for an Adsorbate Layer

The physics in the following section is mainly based on Rogal/Reuter [6] as well as Beret/Ghiringhelli [7].

To calculate the Gibbs free energy of each possible polymorph, we model the initial state before deposition and the final state when the adsorbate layer has formed as illustrated in Figure 1.3. The initial state is the 'clean' system that consists of the gas phase and the substrate without any adsorbed molecules as sketched on the left side of Figure 1.3. This initial state is identical for all possible polymorphs. On the contrary, each possible polymorph corresponds to a different final state. These final states are modeled as illustrated on the right side of the aforementioned figure. The 'best' polymorph is then the one where the Gibbs free energy decreases the most from the initial 'clean' state to the final state that includes the adsorbate layer. To calculate this difference, we split up both the initial as well as the final state into three parts:

1. The gas phase above a certain height – i.e. that part of the gas phase, which is unaffected by the adsorbate layer.
2. The lower layers of the substrate. This part is not influenced by the interface, too, and can, therefore, be described in both states by the same Gibbs free energy G_{subst}.
3. Finally, we have the vicinity of the interface. This is the only part that is changed during the deposition (in the approximation of this model).
   This part consists of N_{ads} molecules which represent the adsorbate layer of the final state and are part of the gas phase in the initial state. Furthermore, it holds N_{subst} atoms of the uppermost layers of the substrate. In the initial state, we can describe the chemical potential of each of those molecules/atoms with μ_{gas} and μ_{subst}, respectively. In the final state, however, we cannot split the contribution of single particles that easily, but we have to describe the Gibbs free energy of this part of the unit cell at once ($G_f(N_{ads}, N_{subst})$). With that, the difference in Gibbs free energy is given by:

$$\Delta G = G_f(N_{ads}, N_{subst}) - N_{subst}\ \mu_{subst} - N_{ads}\ \mu_{gas} \tag{1.1}$$

In this study, we conduct a geometry optimization for each polymorph that is plotted in the phase diagram. In this optimization, we include the uppermost substrate layers while the lower layers are assumed to stay constant. Consequently, the *change* of Gibbs free energy of the substrate, $G_f(0, N_{subst})$, is included in the adsorption energy of the adsorbates and, hence, we can substitute as following:

$$G_f(N_{ads}) = G_f(N_{subst}, N_{ads}) - N_{subst}\ \mu_{subst} \tag{1.2}$$

With that, we can simplify equation 1.1

$$\Delta G = G_f(N_{ads}) - N_{ads}\ \mu_{gas} \tag{1.3}$$

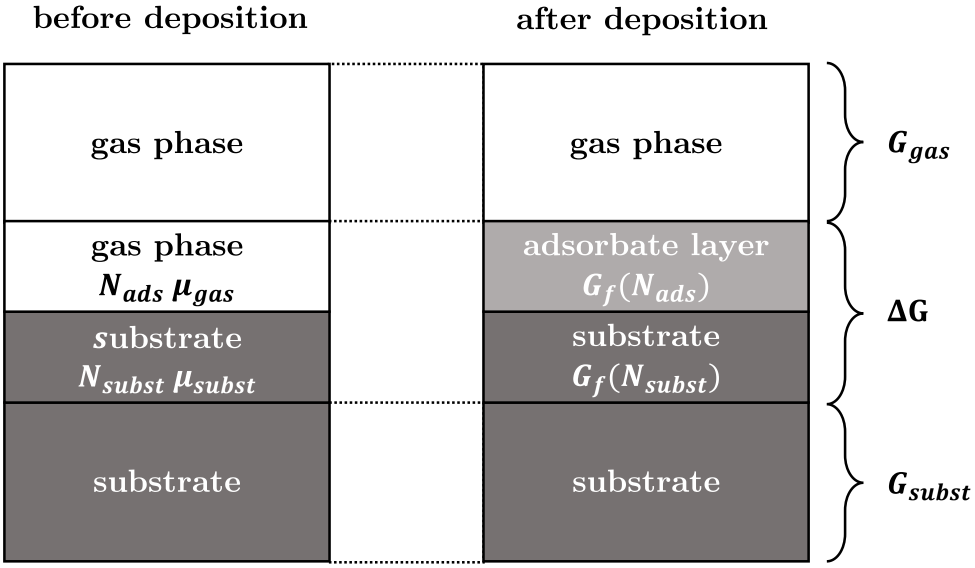

Figure 1.3.: Schematic view of the studied interface system

Furthermore, for a homogeneous surface, ΔG scales linearly with the lateral system size and so we are allowed to define the change in the free energy *per area*: $\Delta\gamma$.

$$\Delta\gamma = \Delta G/A \tag{1.4}$$

$$\Delta\gamma_{ads} = \frac{1}{A}\Big(G_f(N_{ads}) - N_{ads}\ \mu_{gas}\Big) \tag{1.5}$$

$\Delta\gamma_{ads}$ consists two contributions: First, $G_f(N_{ads})$, the Gibbs free energy of N_{ads} adsorbed molecules: This term depends on properties like the adsorption energy of the adsorbates, therefore it has a different value for each possible polymorph. The other term, $N_{ads}\ \mu_{gas}$, only depends on the external conditions pressure and temperature, but not on the specific polymorph that forms on the surface.

Let us first have a look at $G_f(N_{ads})$.
In general, Gibbs free energy G is given by:

$$G = \sum_i \mu_i \cdot N_i = F + p \cdot V = U - T \cdot S + p \cdot V \tag{1.6}$$

When we split off the configurational part F^{conf} and the vibrational part F^{vib} from the total Helmholtz free energy F we are left with E^{DFT} which denotes the ground state energy of N_{ads} TCNE molecules in the adsorbate layer. We will come back to that term later.

$$G_f(N_{ads}) = E_f^{DFT}(N_{ads}) + F_f^{conf}(N_{ads}) + F_f^{vib}(N_{ads}) + p \cdot V_f(N_{ads}) \tag{1.7}$$

Now we use equation 1.6 to split F^{conf} and F^{vib} further up into the inner energy (U) and the entropic contribution (TS). For F^{conf} the inner energy vanishes, and it therefore consists only of the entropic part TS^{conf}. The vibrational part F^{vib} consists of the zero-point energy E^{ZPE} (which is not covered in E^{DFT}) and the vibrational entropy TS^{vib} [20].

$$
\begin{aligned}
F^{conf} &= \qquad - TS^{conf} && (1.8)\\
F^{vib} &= E^{ZPE} - \ TS^{vib} && (1.9)
\end{aligned}
$$

With that we update equation 1.5:

$$\Delta\gamma_{ads} = \frac{1}{A}\left(E_f^{DFT}(N_{ads}) - TS_f^{conf}(N_{ads}) + E_f^{ZPE}(N_{ads}) - TS_f^{vib}(N_{ads}) + p\cdot V_f(N_{ads}) - N_{ads}\ \mu_{gas}\right) \tag{1.10}$$

Now let us turn to the chemical potential μ of the molecules in gas phase. Under the assumption of an ideal gas, the chemical potential μ at a given temperature T and pressure p is given by:

$$\mu_{gas}(T,p) = -\frac{1}{N}\left(k_BT\ ln\left(Q_{gas}^{tot}\right) + pV\right) \tag{1.11}$$

$$Q_{gas}^{tot} = \frac{1}{N!}\left(q^{trans}q^{rot}q^{vib}q^{electr}q^{nucl}\right)^N \tag{1.12}$$

Q_{gas}^{tot} denotes the partition function of N indistinguishable particles. Inserting Q_{gas}^{tot} into equation 1.11 and evaluating the logarithm yields the separated contributions to $\mu_{gas}(T,p)$ where the factor $\frac{1}{N!}$ accounting for the indistinguishable states has been moved into μ^{trans}.

$$\mu_{gas}(T,p) = \mu^{trans} + \mu^{rot} + \mu^{vib} + \mu^{electr} + \mu^{nucl} - \frac{pV}{N} \tag{1.13}$$

q^{electr} leads to μ_{gas}^{electr}, the electronic contribution of the chemical potential, which is described by E_{gas}^{DFT}, the ground state energy of one molecule in gas phase obtained via density functional theory. q^{nucl} can be neglected since the nuclear states will hardly change in the studied processes. The vibrational contribution q^{vib} results in μ_{gas}^{vib} and will be dealt with later.
The remaining translational and rotational contributions (μ^{trans} and μ^{rot}) can be calculated analytically under the assumption of an ideal gas. In the following text these two contributions will be denoted as $\Delta\mu_{gas}(T,p)$ to shorten the equations. Note that the formula for S_{rot} (equation 1.18) is only valid for non-linear molecules like TCNE.[a]

$$\mu_{gas}(T,p) = E_{gas}^{DFT} + \mu_{gas}^{vib} + \Delta\mu_{gas}(T,p) \tag{1.14}$$

$$\Delta\mu_{gas}(T,p) = E_{trans} + E_{rot} - T\ (S_{trans} + S_{rot}) \tag{1.15}$$

$$E_{trans} + E_{rot} = c_{v,trans} + c_{v,rot} + (c_v - c_p) = \frac{3}{2}k_BT + \frac{3}{2}k_BT + k_BT \tag{1.16}$$

$$S_{trans} = k_B\ ln\left(\left(\frac{2\pi m}{h^2}\right)^{3/2}\frac{(k_BT)^{5/2}}{p}\right) + \frac{5}{2}k_B \tag{1.17}$$

$$S_{rot} = k_B\ ln\left(\sqrt{\prod_{i=1}^{3}\frac{\pi}{\sigma}I_i}\ \left(\frac{8\pi^2k_BT}{h^2}\right)^{3/2}\right) + \frac{3}{2}k_B \tag{1.18}$$

I_i ... moments of inertia of the molecule in the gas phase

σ ... number of rotational symmetry equivalents of the molecule in the gas phase; TCNE: point group $= D_{2h}$; $\sigma_{TCNE} = 4$ (Ref. [22, table 10.1])

---

[a] A more detailed derivation of these terms can be found in Ref. [21, 22]. For the actual calculations the implementation of the python package ASE [23] has been used.

Inserting $\mu_{gas}(T,p)$ from equation 1.14 into Gibbs free energy per area $\Delta\gamma_{ads}$ (equation 1.5) leads to:

$$\begin{aligned}\Delta\gamma_{ads} = \frac{1}{A}\Big[&E_f^{DFT}(N_{ads}) - TS_f^{conf}(N_{ads}) + E_f^{ZPE}(N_{ads}) - TS_f^{vib}(N_{ads}) + p\cdot V_f(N_{ads})\\ &- N_{ads}\,\left(E_{gas}^{DFT} + \mu_{gas}^{vib} + \Delta\mu_{gas}(T,p)\right)\Big]\end{aligned} \tag{1.19}$$

In analogy to ΔG (see equation 1.1) we define the adsorption energy $E_{ads}^{DFT}(N_{ads})$ as the change in the inner energy of N_{ads} molecules upon adsorption from the gas phase.

$$E_{ads}^{DFT}(N_{ads}) = E_f^{DFT}(N_{ads}) - N_{ads}E_{gas}^{DFT} \tag{1.20}$$

$$\begin{aligned}\Delta\gamma_{ads} = \frac{E_{ads}^{DFT}(N_{ads})}{A} + \frac{1}{A}\Big[&- TS_f^{conf}(N_{ads}) + E_f^{ZPE}(N_{ads}) - TS_f^{vib}(N_{ads}) + p\cdot V_f(N_{ads})\\ &- N_{ads}\,\left(\mu_{gas}^{vib} + \Delta\mu_{gas}(T,p)\right)\Big]\end{aligned} \tag{1.21}$$

To estimate the maximal size of the work term $1/A\,\cdot\,p\cdot V_f(N_{ads})$ we assume a maximal pressure of $10^5\ Pa$ and a height of the adsorbate layer of at most $1\ nm$. Even in this upper limit the contribution of the work term is three orders of magnitude smaller than the typical adsorption energies which are in the range of 1 to 10 eV/nm^2. Consequently, we can safely neglect this term.

$$\begin{aligned}\frac{1}{A}\,p\cdot V_f(N_{ads}) &\leq 10^5\ Pa\cdot 1\ nm = 10^5\cdot 10^{-9}\ J/m^2\\ &\leq 10^{-4}\cdot\frac{6\cdot 10^{18}eV}{10^{18}\ nm^2}\\ &\leq 1\ meV/nm^2\end{aligned} \tag{1.22}$$

Next we search for an upper limit for the configurational entropy contribution $\frac{1}{A}\,TS_f^{conf}(N_{ads})$. To that end we model $S_f^{conf}(N_{ads})$ as an arrangement of k defect sites in N_{ads} total adsorption sites for each unit cell. Furthermore, we make use of the Stirling formula in equation 1.23. To find an upper limit we assume a rather high defect density of $k/N_{ads} = 0.1$ (equation 1.24).

$$\begin{aligned}\frac{1}{A}\,F_f^{conf}(N_{ads}) &\approx -\frac{1}{A}\,T\cdot S_f^{conf}\\ &\approx -\frac{1}{A}\,k_BT\cdot ln\left(\frac{N_{ads}!}{k!\,(N_{ads}-k)!}\right)\\ &\approx -\frac{1}{A}\,k_BT\cdot N_{ads}\,\left[\frac{k}{N_{ads}}\cdot ln\left(\frac{N_{ads}}{k}-1\right) - ln\left(1-\frac{k}{N_{ads}}\right)\right] &(1.23)\\ &\approx -\frac{N_{ads}}{A}\,0.3\,k_BT &(1.24)\end{aligned}$$

At a room temperature this approximation results in a configurational entropy in the order of 10 meV per adsorption site[b] or equivalently 10 to 40 meV/nm^2. This upper limit is still a good

---

[b] Assuming a coverage in the range of 1 to 5 adsorbate molecules per nm^2

deal smaller than the adsorption energies but already in the same order of magnitude as the numerical accuracy of DFT calculations[c] – which is a level of accuracy that we do not want to sacrifice.
However, as we are interested in differences between polymorphs, most of this contribution cancels out. For two phases with identical defect density and identical coverage $\Theta = N_{ads}/A$ there is no difference in configurational entropy at all. For different coverages Θ_1, Θ_2 the difference in configurational entropy can be approximated as $0.3\ k_BT \cdot (\Theta_1 - \Theta_2)$. Hence, we are allowed to neglect this contribution as long as we do comparisons between polymorphs of identical or similar coverage.

Finally let us have a look at the vibrational contributions $F_f^{vib}(N_{ads})$ and μ_{gas}^{vib} which are split up into a zero point energy and an entropic contribution.

$$\begin{aligned} F_f^{vib}(N_{ads}) - N_{ads}\ \mu_{gas}^{vib} &= E_f^{ZPE}(N_{ads}) - T \cdot S_f^{vib}(N_{ads}) - N_{ads}\ E_{gas}^{ZPE} + N_{ads}\ T \cdot S_{gas}^{vib} \\ &= \Delta E^{ZPE}(N_{ads}) - T \cdot \Delta S^{vib}(N_{ads}) \end{aligned} \tag{1.25}$$

Calculating vibrational properties is far costlier than calculating properties like the adsorption energy (see chapter 5.2). For the gas phase terms E_{gas}^{ZPE} and S_{gas}^{vib} this would still be possible as we assume all TCNE molecules in the gas phase to be identical. For the adsorbate terms $E_f^{ZPE}(N_{ads})$ and $S_f^{vib}(N_{ads})$, however, not only the calculations are far more costly due to the additional substrate atoms, but most importantly these terms have different values for each possible polymorph!

From equation 1.25 we see that only the change of the vibrational modes is relevant – therefore it might be possible to concentrate on a few most relevant modes. Additionally, for comparisons between polymorphs only the differences in their vibrational modes contribute and all changes that are equal for different configurations cancel out. At the same time this means that including only the vibrational terms of the gas phase molecules would worsen the results. All in all we are currently forced to neglect the vibrational terms.

Applying all those considerations to equation 1.21 gives us a final approximation for the Gibb's free energy of adsorption $\Delta\gamma_{ads}$. Furthermore, we plug equations 1.16, 1.18, and 1.17 into equation 1.15 to obtain an approximation for $\Delta\mu_{gas}(T,p)$.

$$\Delta\gamma_{ads} = \frac{E_{ads}^{DFT}(N_{ads})}{A}\ -\ \Theta \cdot \Delta\mu_{gas}(T,p) \tag{1.26}$$

$$\Delta\mu_{gas}(T,p) = -k_BT\ ln\left[\left(\frac{2\pi m}{h^2}\right)^{3/2}\frac{(k_BT)^{5/2}}{p}\ +\ \frac{\sqrt{\pi}}{\sigma}\sqrt{\prod_{i=1}^{3} I_i}\ \left(\frac{8\pi^2 k_BT}{h^2}\right)^{3/2}\right] \tag{1.27}$$

1.2.2. Calculating the Adsorption Energy

An important aspect of equation 1.26 is that when we rank polymorphs of identical coverage Θ by their Gibbs free energy only the adsorption energy E_{ads} is relevant, as the contribution of the chemical potential is identical.

In equation 1.20 E_{asd}^{DFT} is defined as the difference in the inner energy of N_{ads} molecules upon adsorption from the gas phase. For the practical use, however, we cannot calculate the energy

---

[c] DFT settings are converged to the order of 10 meV per unit cell.

of the adsorbed molecules directly. What we can calculate is the energy of a substrate covered by the adsorbate layer and the same substrate without the adsorbate layer. The difference yields the total energy of the adsorbed molecules in that unit cell[d].

$$E_f^{DFT}(N_{ads}) = E_{full\ geometry} - E_{substrate} \tag{1.28}$$

Inserting this definition in equation 1.20 gives:

$$E_{ads}^{DFT}(N_{ads}) = E_{full\ geometry} - E_{substrate} - N_{ads}E_{gas}^{DFT} \tag{1.29}$$

Furthermore, we define the mean adsorption energy per adsorbate molecule and the mean adsorption energy per surface area.

$$E_{ads,\ per\ molecule} = \frac{E_{ads}^{DFT}(N_{ads})}{N_{ads}} \tag{1.30}$$

$$E_{ads,\ per\ area} = \frac{E_{ads}^{DFT}(N_{ads})}{A} = E_{ads,\ per\ molecule} \cdot \Theta \tag{1.31}$$

1.3. Electronic Structure Theory

The following chapter is based on *Introduction to computational chemistry* by Frank Jensen [24]

All good things start with the Schrödinger equation [25], more precisely the time-independent Schrödinger equation:

$$\hat{H}\Psi = E\Psi \tag{1.32}$$

For a system of atom, nuclei, and electrons the Hamiltonian $\hat{H}$ is given by:

$$\hat{H} = \hat{T}_n + \hat{V}_{n\text{-}n} + \hat{T}_e + \hat{V}_{n\text{-}e} + \hat{V}_{e\text{-}e} \tag{1.33}$$

$$\hat{T}_n = \sum_A \frac{-\hbar}{2m_A}\nabla_A^2 \tag{1.34}$$

$$\hat{V}_{n\text{-}n} = \frac{1}{4\pi\epsilon_0}\sum_{A,B>A}\frac{Z_A Z_B e^2}{|\vec{R_B} - \vec{R_A}|} \tag{1.35}$$

$$\hat{T}_e = \sum_i \frac{-\hbar}{2m_i}\nabla_i^2 \tag{1.36}$$

$$\hat{V}_{n\text{-}e} = \frac{1}{4\pi\epsilon_0}\sum_{i,A}\frac{-Z_A e^2}{|\vec{r_i} - \vec{R_A}|} \tag{1.37}$$

$$\hat{V}_{e\text{-}e} = \frac{1}{4\pi\epsilon_0}\sum_{i,j>i}\frac{e^2}{|\vec{r_j} - \vec{r_i}|} \tag{1.38}$$

As the nuclei are far heavier than the electrons we assume that the electronic problem can be separated from that of the nuclei. I.e. we find a solution for the positions of the electrons assuming constant positions of the nuclei. This commonly used assumption is known as Born-Oppenheimer approximation [26]. Within that approximation the purely nuclear terms $\hat{T}_n$ and

---

[d] With a coverage of N_{ads} adsorbate molecules per unit cell.

$\hat{V}_{n\text{-}n}$ can be removed from the Hamiltonian of the electronic problem as they only add a constant contribution to the eigenvalues.

$$\hat{H}_{BO} = \hat{T}_e + \hat{V}_{n\text{-}e} + \hat{V}_{e\text{-}e} \tag{1.39}$$

But still, in this many-body problem each electron adds three spatial dimensions and one spin dimension to the problem. Therefore, for a system of N electrons we have to find a $4N$-dimensional wave function.

1.3.1. Density Functional Theory

Hohenberg and Kohn [27] showed that there exists a functional $E_{DFT}[\rho]$ that establishes a one-to-one correspondence between the electron density $\rho(r, \sigma)$ of the many-body wave function and its ground state properties. Hence, instead of searching for the high-dimensional ground state wave function, one directly calculates the electron density and from there on properties like the ground state energy E_{DFT}. The electron density is given by the square of the wave function summed over all electrons. Therefore, switching from the wave function to the electron density reduces the dimensionality of the problem from $4N$ to 4, where N is the number of electrons. This implies that the dimensionality does not grow with the number of electrons anymore. This proof provides the basic principle for density functional theory (DFT). Unfortunately, Hohenberg and Kohn only showed that there *exists* such a functional, but they provided no analytical form of that functional.

In analogy to equation 1.39 the (unknown) functional $E_{DFT}[\rho]$ is split up into the kinetic functional $T_e[\rho]$, the functional $E_{n\text{-}e}[\rho]$ describing the interaction between nuclei and electrons and the functional $E_{e\text{-}e}[\rho]$ for the interaction between electrons and electrons. $E_{e\text{-}e}[\rho]$ is then further split up into a Coulombic part $J[\rho]$ and the exchange functional $K[\rho]$.

$$E_{DFT}[\rho] = T_e[\rho] + E_{n\text{-}e}[\rho] + E_{e\text{-}e}[\rho] \tag{1.40}$$

$$= T_e[\rho] + E_{n\text{-}e}[\rho] + J[\rho] + K[\rho] \tag{1.41}$$

$E_{n\text{-}e}$ and $J[\rho]$ can be described classically:

$$\begin{aligned} E_{n\text{-}e}[\rho] &= -\sum_A \int \frac{Z_A(\vec{R}_A)\rho(\vec{r})}{|\vec{R}_A - \vec{r}|} \\ J[\rho] &= \frac{1}{2} \int \int \frac{\rho(\vec{r})\rho(\vec{r}')}{|\vec{r} - \vec{r}'|} d\vec{r} d\vec{r}' \end{aligned} \tag{1.42}$$

For the general case $T_e[\rho]$ and $K[\rho]$ are still unknown. For the special case of a uniform electron gas Thomas, Fermi and Dirac found functionals for $T_e[\rho]$ and $K[\rho]$ (Thomas-Fermi-Dirac model). This model works reasonably well for metallic systems but is a very poor model for every system with abrupt changes in the electron density – like atoms and molecules.

A major step towards actual application of the DFT approach was done by Walter Kohn and Lu Jeu Sham [28]. Their fundamental improvement was to split the kinetic functional $T_e[\rho]$ into two parts: A major part of $T_e[\rho]$ is already covered by the kinetic energy $\hat{T}_S$ of a auxiliary, *non-interacting* electron gas which can be solved exactly. We are then left with a comparably small correction E_{xc}. The price that comes with this Kohn-Sham approach is that auxiliary orbitals ϕ have to be re-introduced to describe the auxiliary electrons and hence the dimensionality is increased again.

It might seem as if this approach only shifted the problem to the new exchange-correlation functional E_{xc}. However, since E_{xc} makes up only a small part of the total energy, a bad approximation for E_{xc} has far less influence on the result than a bad approximation for the full kinetic functional $T_e[\rho]$.

To obtain the major contribution $\hat{T}_S$ a special Hamiltonian operator $\hat{H}_\lambda$ is introduced that allows to adjust the interaction strength of the electrons:

$$\hat{H}_\lambda = \hat{T}_e + \hat{V}_{ext}(\lambda) + \lambda \hat{V}_{e-e} \tag{1.43}$$

$\lambda = 1$ reproduces the real system: $\hat{V}_{ext}(\lambda{=}1) = \hat{V}_{n-e}$ whereas the other extremum $\lambda = 0$ represents a purely hypothetical system of non-interacting electrons where $\hat{V}_{ext}(\lambda)$ is adjusted such that the electron density stays constant. For this system of non-interacting electrons the *exact* kinetic energy functional can be written using an anti-symmetrized wave function composed of molecular orbitals (the Slater determinant):

$$\hat{T}_S = \sum_{i=1} \langle \psi_i | - \frac{\hbar}{2m_i} \nabla^2 | \phi_i \rangle \tag{1.44}$$

The full DFT functional can now be written as:

$$E_{DFT}[\rho] = T_S[\rho] + E_{n\text{-}e}[\rho] + J[\rho] + E_{xc}[\rho] \tag{1.45}$$

This expression for E_{DFT} must be equivalent to the expression in equation 1.40 and hence we can explicitly express the remaining exchange-correlation functional $E_{xc}[\rho]$.

$$E_{xc}[\rho] = (T_e[\rho] - T_S[\rho]) \ + \ (E_{e\text{-}e}[\rho] - J[\rho]) \tag{1.46}$$

The term in the first parenthesis of equation 1.46 corresponds to that share of the kinetic functional $T_e[\rho]$ which is not covered by $T_S[\rho]$. This missing kinetic correlation energy must, therefore, be included in this exchange-correlation functional. The term in the second parenthesis stands for the potential correlation and the exchange energy.

To sort different models for the exchange-correlation functionals, Perdew's ladder [29] was introduced. This hypothetical ladder starts from the crude approximation of missing the non-interaction electron gas to the heavenly realms of chemical accuracy. Each rung of this ladder corresponds to one class of functionals with (typically) increasing accuracy at the cost of increasing model complexity and, therefore, increasing computational cost.

Local Density Approximation (LDA)

The first rung, the Local Density Approximation, is the most basic approximation for the exchange-correlation functional $E_{xc}[\rho]$. It is based on the assumption that the electron density varies only slowly and thus $\rho(\vec{r})$ can be at least locally approximated by a uniform electron gas for which the exchange energy is given analytically (Thomas-Fermi-Dirac model). Consequently, this approximation works well for metals but not so well for atoms and molecules. For the case of different electronic densities of the two spin channels, this approach is generalized to the Local Spin Density Approximation (LSDA).

Generalized Gradient Approximation (GGA)

To model systems with non-vanishing electronic density gradients, gradient-corrected methods were introduced. Unfortunately, simply including the gradient $\Delta\rho$ actually leads to worse results than the Local Density Approximation. To overcome this problem functionals with several parameters were introduced. These parameters are either fitted using empirical data or such that they fulfill certain theoretical conditions.

Common examples of empirically fitted GGA functionals are the *B88* exchange functional proposed by A. D. Becke [30] and the *LYP* correlation functional introduced by Lee, Yang, and Parr [31]. Often those two functionals are used in combination, which is consequently denoted as *BLYP*.

Another common functional, that in contrast to *B88* and *LYP* is not based on empiric data but on theoretical conditions, is the *PBE*-functional introduced by Perdew, Burke and Ernzerhof [32].

Meta-GGA or higher order gradient methods

The next rung corresponds to the next higher is to add higher derivatives of the electronic density, namely, the second-order term (Laplacian, $\nabla^2\rho$) or, equivalently, the orbital kinetic energy density τ.

Examples of meta-GGA methods are the correlation functionals *B95* by Becke [33], the *SCAN* functional [34], and the *TPSS* functional [35].

Hyper-GGA or hybrid methods

One approach to improve the approximation is to mix the exchange-correlation functional with Hartree-Fock exact exchange. One of the simplest hybrid methods mixes one half of the *LDA* functional with one half of exact exchange. Consequently, this approach it known as *Half-and-Half* (*H+H*) method.

The next logical step is to replace the *LDA*-contribution by a *GGA*-functional. This is done, for instance, in the *Becke three parameter functional*(*B3*) which – as the name suggests – uses three mixing parameters. One often used version of this approach is the *B3LYP* approach which uses the *B88* exchange and the *LYP* correlation functional (equation 1.47, Ref. [24]. These functionals are then combined with exact exchange by the three parameters a, b, and c which are empirically fitted to experiments.

$$E_{xc}^{B3LYP} = (1-a)E_x^{LDA} + aE_x^{\text{exact}} + b\Delta E_x^{B88} + (1-c)E_c^{LDA} + cE_c^{LYP} \tag{1.47}$$

Another commonly used hybrid functional is *PBE0* [36] which mixes three parts of *PBE* exchange with one part of Hartree-Fock exchange while the correlation is evaluated at *PBE* level.

$$E_{xc}^{PBE0} = E_{xc}^{GGA} + \frac{1}{4}(E_x^{HF} - E_x^{GGA}) \tag{1.48}$$

Generalized random phase approximation (RPA)

For the accurate description of several effect – for instance van der Waals interactions — virtual (unoccupied) orbitals are necessary. Methods that include such virtual orbitals constitute the highest rung of Jacob's ladder.

Another classification that can be made for exchange-correlation functionals is if the used parameters are either derived from theoretical considerations or if they are fitted to experimental data. Those functionals that are based on experimental data (e.g. *BLYP* and *B3LYP*) typically yield best results for system which are similar to those on which they were parametrized. Typically, these parametrization sets include molecular systems and, hence, empirically fitted functionals are often used for such molecular systems. On the other hand, non-empirical functions (e.g. *PW91* and *PBE*) are more generally applicable and, therefore, more often used for periodic systems like metals.

1.3.2. Charge Partitioning – Mulliken Population Analysis

For the comparison with the experimental data we need to know how the spatial charge distribution on the TCNE molecules changes upon adsorption and more specifically how these charge rearrangements differ for different adsorption geometries.

However, all charge partitioning schemes have one big, common problem: There is no physical observable and consequently no unique definition of how to assign the charge density of the molecule to single atoms. Consequently, several different partitioning schemes exist, with each of them using different definitions and each of them yielding different results. One popular scheme is the Mulliken population analysis [37]. This scheme is a rather simple one and has several imperfections, but still it allows at least for a qualitative analysis of charge transfer processes and is therefore a useful tool for better understanding of the adsorption process.

In a molecule, electronic orbitals are typically not limited to just one atom, but they are spread out over the whole molecule. Consequently, they are called molecular orbitals (MOs). The total electron density ρ of a molecule can be split up into contributions ρ_i corresponding to each of those molecular orbitals.

$$\rho_i(\vec{r}) = \phi_i^2(\vec{r}) \tag{1.49}$$

When we expand those molecular orbitals in a set of non-orthogonal basis functions $\mathcal{X}_\alpha$ with coefficients c_α, the square of the MOs is given by:

$$\phi_i^2 = \sum_{\alpha\beta}^{M_{basis}} c_{\alpha i} c_{\beta i} \mathcal{X}_\alpha \mathcal{X}_\beta \tag{1.50}$$

The second index β had to be introduced as the square of the MOs using *non-orthogonal* basis functions contains non-vanishing off-diagonal elements.

The total number of electrons N_{elec} of a molecule is obtained by integrating the charge density of each MO over $d\vec{r}$ and then summing over all MOs with weights n_i accounting for the occupation of each MO.

$$N_{elec} = \sum_{i}^{N_{orb}} n_i \int \rho(\vec{r}) d\vec{r} \tag{1.51}$$

Applying equations 1.49 and 1.50 yields:

$$N_{elec} = \sum_{i}^{N_{orb}} n_i \int \sum_{\alpha\beta}^{M_{basis}} c_{\alpha i} c_{\beta i} \mathcal{X}_\alpha \mathcal{X}_\beta d\vec{r} \tag{1.52}$$

$$= \sum_{\alpha\beta}^{M_{basis}} \sum_{i}^{N_{orb}} n_i c_{\alpha i} c_{\beta i} \int \mathcal{X}_\alpha \mathcal{X}_\beta d\vec{r} \tag{1.53}$$

$$= \sum_{\alpha\beta}^{M_{basis}} D_{\alpha\beta}\ S_{\alpha\beta} \tag{1.54}$$

$$D_{\alpha\beta} = \sum_{i}^{N_{orb}} n_i c_{\alpha i} c_{\beta i} \tag{1.55}$$

$$S_{\alpha\beta} = \int \mathcal{X}_\alpha \mathcal{X}_\beta d\vec{r} \tag{1.56}$$

Here we introduced the density matrix $D_{\alpha\beta}$ accounting for the occupation of basis functions $\mathcal{X}_\alpha$ and $\mathcal{X}_\beta$ as well as the overlap matrix $S_{\alpha\beta}$ accounting for the spatial overlap of that pair of basis functions.

The elementwise(!) product $D_{\alpha\beta}\ S_{\alpha\beta}$ yields a matrix that is used by the Mulliken Population Analysis. At this point we face two challenges:
First, we have to assign the elements of that matrix to the basis functions $\mathcal{X}$. The diagonal elements $D_{\alpha\alpha}\ S_{\alpha\alpha}$ are easy, they solely contribute to the basis function $\mathcal{X}_\alpha$. The off-diagonal elements, however, indicate electrons that are shared by both basis functions $\mathcal{X}_\alpha$ and $\mathcal{X}_\beta$. Here, Mulliken population analysis uses the most straightforward way of assigning one half to each contributing basis function. This is actually a rather strong simplification as it does not consider differences in the electronegativity of the involved atoms. Additionally, this scheme does *not* guarantee that the matrix elements are bound by 0 and 2 as it would be expected from a physical standpoint.

The second decision lies in the assignment of basis function to atoms. Gladly, in the case of atom-centered orbitals (as they are used in FHI-*aims*) there is a natural correspondence between basis functions and atoms. But even in this seemingly clear case, problems can occur. Diffuse basis functions extend by design into regions far from the nucleus they are centered on. If there is another nucleus in that region the diffuse basis function might describe electron density that would more reasonably be assigned to the second nucleus and not that nucleus on which the basis function is centered. Furthermore, when we enlarge the basis set – i.e. we use more basis functions – this effect does not vanish, but it actually becomes even worse.

With these assumptions we can formulate the Mulliken electron population ρ_A for an atom A. As each off-diagonal element occurs twice in the $D_{\alpha\beta}S_{\alpha\beta}$-matrix, the assignment factor of 1/2 cancels out.

$$\rho_A = \sum_{\alpha\in A}^{M_{basis}} \sum_{\beta}^{M_{basis}} D_{\alpha\beta}\ S_{\alpha\beta} \tag{1.57}$$

The gross charge on that atom is then given by $Q_A = Z_A - \rho_A$ with Z_A being the atomic number. One final limitation should be noted on Mulliken population analysis: It does not reproduce multipole moments (dipole, quadrupole, etc.), i.e. calculating the multipole moments of the molecule from the Mulliken charges will not lead to the same result as a direct integral over the electron density $\rho(\vec{r})$.

There are several alternative charge partitioning schemes. A very brief overview shall be given in the following:

Löwdin analysis

This is merely a variation of Mulliken population analysis which generalizes the multiplication DS to S^nDS^{n-1} with $n = 1/2$. This orthogonalization step leads to a more reasonable assignment of the electrons to the atoms which, for instance, limits the occupation numbers to the range of 0 to 2. On the other hand, this scheme is not rotational invariant anymore.

Electrostatic potential (ESP) charges

The idea of this approach is to first calculate the electrostatic potential in a region around the studied atoms. Subsequently, point charges are placed on the positions of the nuclei with the amount of charge chosen such that they reproduce the external electrostatic potential. This leads directly to the strength of this approach: in contrast to other charge partition schemes, the evaluated charges can correctly reproduce the electrostatic behavior! A problem of this scheme can be the fitting procedure: For instance, in a close packing of atoms, there is simply too little space between the atoms to enable a good fitting accuracy.

Bader charges / Atoms in molecules

This analysis is based on the electron density. The boundary between two atoms is then defined as that line where the gradient of the electronic density is zero. The electron density inside this boundary is then integrated and assigned to the corresponding nucleus. In the general case, this nucleus is not in the exact center of this region of electron density and, therefore, these Bader charges do not reproduce electrostatics. A benefit of this approach is that the evaluated saddle point of the electron density between the atoms is also a measure for the strength of this bond.

Hirshfeld charges

In this partition scheme, electron density $\rho(\vec{r})$ is assigned to neighboring nuclei according to the electron density that virtual free atoms, placed on these nuclei, would have in $\vec{r}$. The difference between this assigned charge and the charge of the free atom is then used as partial charge of that atom. This scheme is, for instance, used for the Tkatchenko-Scheffler van der Waals correction [38].

2. SAMPLE – Surface Adsorbate Prediction With Little Effort

Surface **A**dsorbate Poly**m**orph **P**rediction with **L**ittle **E**ffort (SAMPLE) [8, 9, 10] is an approach for global structure search that is based on Bayesian linear regression. It enables the prediction of physical properties for an exhaustive set of candidate polymorphs with an accuracy of the same order of magnitude as the numerical accuracy of DFT calculations.

The following chapter 2.1 collects all fundamental assumptions and limitations of the SAMPLE approach. Chapter 2.2 then explains how the physical system is modeled, i.e., it will describe (I) how the search space is discretized into so called configurations and (II) the energy model that will be used for the predictions. Finally, chapter 2.3 presents how machine learning is employed to find the best set of fitting parameters for the energy model.

For the further way from the predicted adsorption energies, via ab initio thermodynamics, to the surface phase diagram please see chapter 1.2 as well as the results in chapter 4.2 ff.

2.1. Assumptions and Limitations

The SAMPLE approach is a powerful tool. However, to enable efficient training and prediction of immense numbers of candidate polymorphs, it is limited to a specific range of operations.

Commensurable Structures

Adsorbate layers are commensurable on a specific surface if the unit cell of the adsorbate layer can be defined using only integer multiples of the substrate unit vectors. This implicates that each periodic unit of the adsorbate layer 'sees' an identical patch of surface below it. This characteristic is typically fulfilled if the corrugation of the potential energy surface is high – i.e. when the energy barrier between two local adsorption geometries is higher than the interaction energy of two molecules sitting in those local minima. Commensurability of the adsorbate layer is a necessary assumption for the SAMPLE approach as it allows us to describe the surface system with periodic unit cells.

Thermodynamic equilibrium

For evaluations like the surface phase diagram, we assume that the adsorbate layer assumes the global minimum of Gibbs free energy for each combination of pressure and temperature. However, just like mathematical optimization algorithms do not always find the global optimum of a function, physical systems do not always assume the global energetic optimum. Energetic barriers can separate local minima from the global energetic minimum and hence lead to the formation of kinetically trapped states. In the SAMPLE approach such metastable states are part of the search space and hence they are included in the predictions. As of now, however, the heights and positions of energy barriers between such states are unknown and therefore we must assume that always that polymorph which corresponds to the global energetic minimum will form. Furthermore, we do not model the detailed formation kinetics of the second layer but we assume a layer-by-layer growth and then compare the final homogeneous bilayer polymorphs with corresponding monolayer polymorphs.

Geometry relaxation

The candidate polymorphs are built out of building blocks named local adsorption geometries (see chapter 2.2.1). The geometry of these local adsorption geometries are optimized *before* they are combined to candidate polymorphs, but not *after* that step. Especially when two neighboring molecules get very close (limited by d_{min}) strong pair interactions and consequently large energy gradients occur. In that case it might be beneficial for molecules to deviate from their local adsorption geometry to greatly lower the pair-interaction energy at the cost of only a little increase in potential energy.

On the one hand, preliminary tests showed that this case indeed occurs and that geometry optimizations of those candidate polymorphs can result in rather strong changes in the geometry in which adsorbate molecules even move into different local adsorption geometries. On the other hand, configurations with high interaction energies leading to such problems are obviously energetically not particularly favorable anyway. Furthermore, if the molecules move into different local adsorption geometries, then these polymorphs can be neglected since that target polymorph is already included in the search space.

During the geometry optimization of local adsorption geometries, the substrate is kept fixed as otherwise the final polymorphs could not be built out of multiple local adsorption geometries.

For the comparison of monolayer and bilayer polymorphs a more extensive geometry optimization is used which is described in chapter 4.3.

Minimal distance d_{min}

A minimal intermolecular distance d_{min} is necessary since the interaction energies diverge for small distances due to Pauli repulsion. This threshold introduces an artificial limit to the search space and could, in principle, lead to missing relevant polymorphs. The value for this parameter is estimated from the interaction energies of dimers in the gas phase. In hindsight, the obtained results (see discussion of Figure 4.3) reassure the choice of the d_{min} as they indicate that there are no energetically favorable polymorphs with intermolecular distances below that threshold.

2.2. Modeling the System

2.2.1. Coarse Graining of the Search Space

If we assume rigid molecules and a fixed adsorption height, each single molecule in the adsorbate layer would have five continuous degrees of freedom (two spatial and three rotational ones). When we further assume that the adsorbate molecules form commensurate structures consisting of unit cells with up to N molecules in the basis we would end up with 5 N continuous degrees of freedom in total. If we would now attempt a straight forward discretization using only 10 steps for each degree of freedom, we would get 10^{5N} distinct combinations. Allowing for a maximum number of three different molecules per unit cell, this would result in 10^{15} combinations. For the presented system one DFT calculation takes around 100 core hours. Consequently, even with a very low rate of 0.01 $€/h_{core}$, conducting a brute force calculation of all combinations would cost roughly 1 quadrillion € – definitely unfeasible!

Fortunately, in the studied system the potential energy surface for single molecules is strongly corrugated. This means that there are adsorption positions and -orientations which are energetically far more favorable than others and those favorable geometries are separated by energy barriers. Each single molecule aims to assume one of these local adsorption geometries which can be found using local geometry optimization from different starting points.

For the system of TCNE on Cu(111) there are three local adsorption geometries that are oriented flat-lying (face-on oriented) and eight local adsorption geometries in which the TCNE molecule is upright-standing (edge-on oriented) as can be seen in Figure 2.1

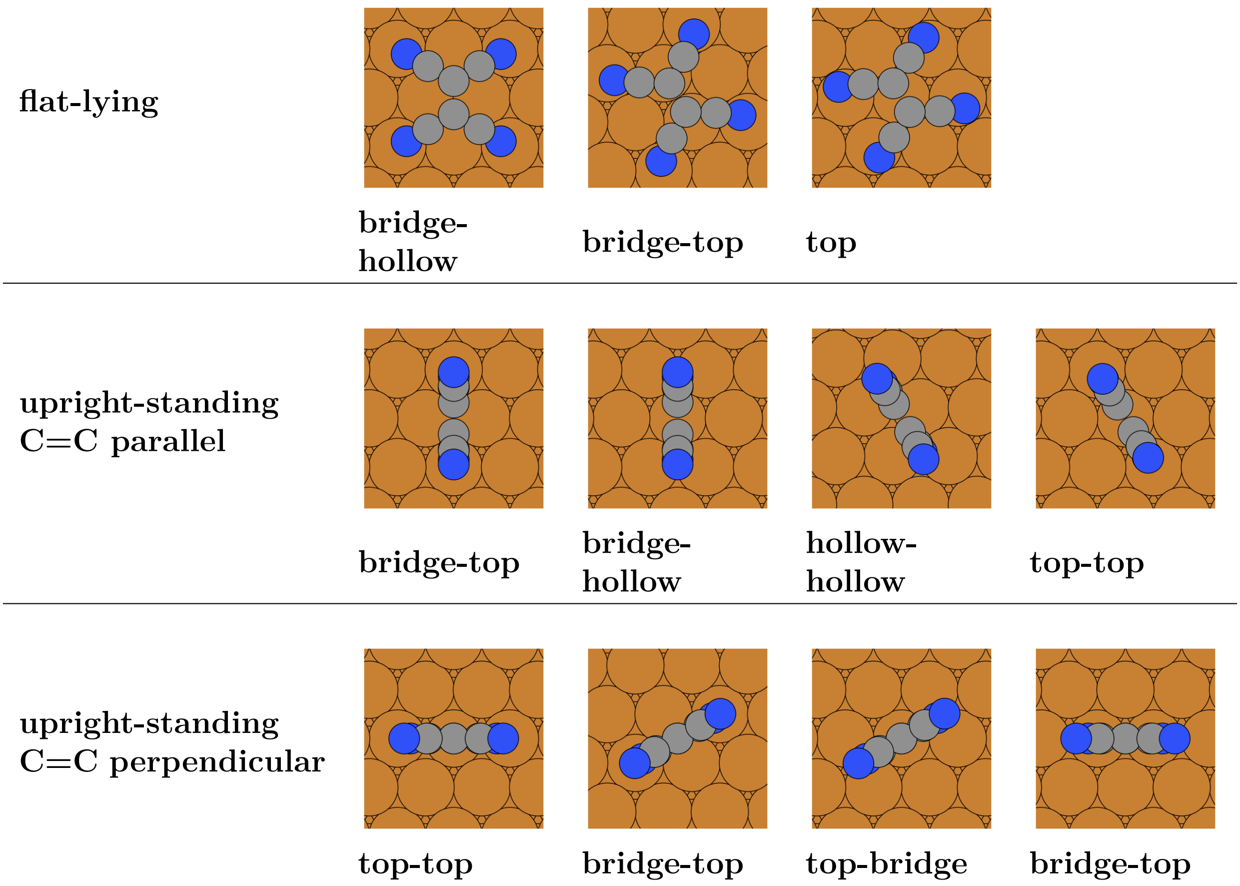

Figure 2.1.: Local adsorption geometries for TCNE on Cu(111) sorted by adsorption energy.
First row: flat-lying
Second row: upright-standing – central C=C bond parallel to substrate surface
Third row: upright-standing – central C=C bond perpendicular to substrate surface. This type of local adsorption geometries is not used in this study (see text).

These local adsorption geometries were originally obtained by Michael Scherbela [12]. I used his geometries as starting points, conducted a local geometry optimization and removed numerical deviations from the symmetry axes by averaging the molecule coordinates with respect to each symmetry axis.

For TCNE on Cu(111), the adsorption geometries with the central C=C bond orientated perpendicular to the surface are generally energetically less favorable than those adsorption geometries where this bond is oriented parallel to the surface. This energetic difference also persists when the pair interactions are included (for details, see Appendix A). Furthermore, the TCNE molecule is approximately square. Hence both orientations of the upright-standing molecules have nearly identical surface footprints and, therefore, the maximal packing density of both orientations is nearly identical, too. Considering all those properties, it seems highly unlikely that the local adsorption geometries with the C=C bond perpendicular to the surface are part of relevant polymorphs. Consequently, only the flat-lying local adsorption geometries as well as those upright-standing geometries where the central C=C bond is parallel to the substrate surface are used in the present study.

To construct a specific polymorph, we have to define its unit cell as well as the arrangement of the molecules inside that unit cell (see Figure 2.2). To that end, we first create all unit cells that can be built on the chosen substrate up to a maximum area A_{max}. A large share of those unit cells is equivalent due to symmetries of the substrate. Therefore, we apply all symmetry operations of the substrate to each unit cell and compare the transformed geometries. Of all sets of symmetry equivalent unit cells only one unique instance is kept. Further details on this step can be found in [10, 39].

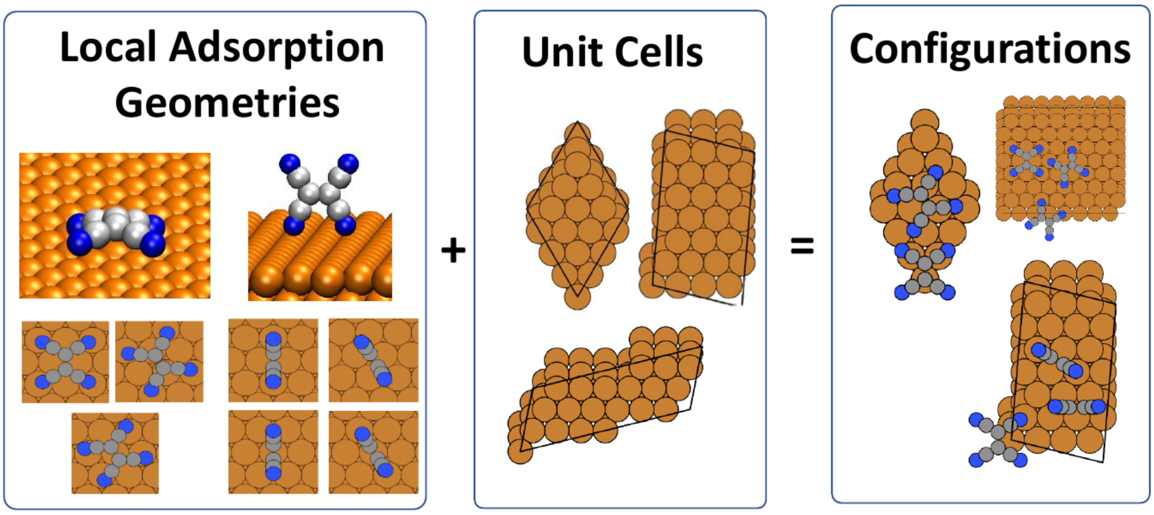

Figure 2.2.: Configurations are discretized, commensurate adsorbate layers consisting of locally optimized single-molecule adsorption geometries in a periodic repetition unit.

Each of those unit cells then serves as one starting point to build configurations. For that endeavor we make use of the fact that each unit cell consist of several *substrate cells* which are spanned by the substrate lattice as can be seen in the first image of Figure 2.3. Due to the translational symmetry of the substrate, all substrate cells are equivalent, and each local adsorption geometry can be placed in each substrate cell. Therefore, the exact adsorption geometry of one molecule can be defined using two indices: one for the local adsorption geometry that it assumes and one for the substrate cell in which it is placed.

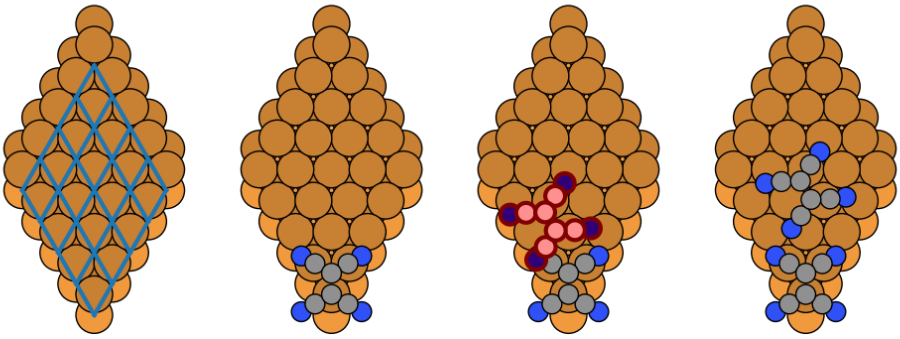

Figure 2.3.: Each local adsorption geometry is proposed to be placed in each substrate cell. If the resulting geometry contains no inter-atomic distance below d_{min} the proposed configuration is added to the pool of possible configurations.

We now start with the first unit cell and put one molecule in the first substrate cell using the first local adsorption geometry. For this first molecule all substrate cells are translational invariant. Hence, we are allowed to place it in the first substrate cell (the one nearest to the origin) without loss of generality. Then we take a second molecule and propose to place it in the next substrate cell using one of the local adsorption geometries. Collisions between this new molecule and all already placed ones are calculated using the Euclidean distances of each atom of the first molecule with each atom of the second one. If one of these distances falls below a predefined threshold the new geometry is rejected. This procedure yields a tree of possible configurations. Each branch of that tree is extended until no further molecule can be placed or until the aspired number of molecules per unit cell (n_{mol}) is reached (whatever happens first). All configurations that fit n_{mol} molecules without collisions are added to the set of possible configurations. All those configurations will later be used for prediction and evaluation. In this study the maximal unit cell size is limited by $A_{max} = 36\ A_{uc}$ with $A_{uc} = 0.0564\ nm^2$ being the surface area of a substrate unit cell. n_{mol} is set to 1, 2, and 3. This way, nearly 11 million configurations consisting of 1, 2 *or* 3 TCNE molecules are created in 215 different unit cells having a surface area of up to 2.03 nm^2.

2.2.2. Energy Model

Now that we have created an exhaustive set of configurations we have to introduce an energy model that will later enable us to predict the adsorption energy of each configuration.

The total adsorption energy of a configuration can be separated into one-body terms E_i, two-body terms $E_{i,j}$ and higher order terms (equation 2.1). In the case of adsorbate layers, the one-body terms correspond to the interaction energies between each single adsorbed molecule and the substrate. The two-body terms account for the energy of each pairwise interaction. Higher order terms account for all interactions which include at least three molecules, and which *cannot* be described by a sum of two-body interactions.

$$E = \sum_i E_i + \sum_i \sum_{j,i<j} E_{i,j} + \dots \tag{2.1}$$

In the present study this series expansion is truncated after one- and two-body terms. It should be noted that we do not limit the two-body terms to nearest neighbors but instead include all pairs up to a certain maximal distance d_{max} which typically also includes next nearest neighbors. This cutoff parameter is chosen high enough such that pair energies above that distance can be safely neglected.[a]

$$E_{pred} = \sum_{geoms} N_g \cdot U_g + \sum_{pairs} N_p \cdot V_p \tag{2.2}$$

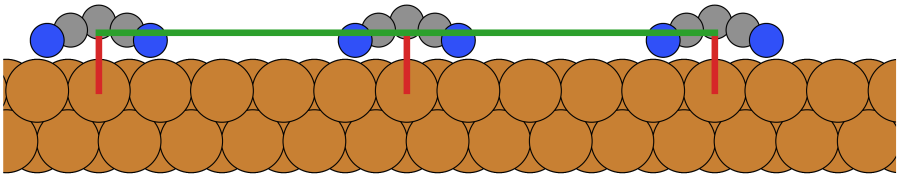

Figure 2.4.: Energy model for the adsorption energy of a configuration
Red lines represent interactions between molecules and the substrate (U_g)
Green lines represent interactions interactions between molecules (V_p)

The main benefit of introducing this energy model is that it limits the number of free parameters. Without this energy model each new configuration would correspond to a new free parameter that needs to be fitted. With this model we have one parameter for each local adsorption geometry as well as one parameter for each distinct pair of molecules. The number of local adsorption geometries is constant since they are predefined and as we use a maximum cutoff distance the number of distinct pairs is limited, too. Each new configuration is build up out of the fixed set of local adsorption geometries and molecule pairs. Therefore, the complexity of our model – i.e. the number of free parameters – does not increase with the number of configurations anymore!

For each configuration the fitting parameters U_g and V_p are combined into a line vector $\vec{\omega}$ and the corresponding coefficients N_g and N_p are combined into one column vector $\vec{N}$. For a whole set of configurations these vectors $\vec{N}$ are then combined to the *model matrix* $\hat{X}$ so that we can write the energy model very compactly:

$$\hat{X} = (\vec{N}_1, ..., \vec{N}_{n_{set}})^T \tag{2.3}$$

$$\vec{\omega} = (U_1, ..., U_{n_{geoms}}, V_1, ..., V_{n_{pairs}}) \tag{2.4}$$

$$\vec{E}_{pred} = \hat{X}\,\vec{\omega}^T \tag{2.5}$$

2.3. Machine Learning

2.3.1. Bayesian Linear Regression

Now we have defined an energy model for the studied system. But with that we also introduced a set of parameters[b] $\vec{\omega}$ for which we have to find the best values.

---

[a] In this study $d_{max} = 1.6\ nm$ is chosen. The fitted interaction energies for pairs of such a distance are around or below 1 meV.

[b] In this study, seven fitting parameters are used for the one-body interactions (one for each local adsorption geometry) and 502 fitting parameters for the two-body interactions.

So what are 'best values'?

The property that we want to predict is the adsorption energy obtained via density functional theory E_{DFT}. For the training of the model we calculate this energy for a set of configurations (the training set) which has the model matrix $\hat{X}$. If we assume any specific set of fitting parameters $\vec{\omega}$, we could use the energy model and predict the adsorption energies of that training set as $\vec{E}_{pred} = \hat{X}\vec{\omega}^T$. The difference between these predictions and the actual, calculated DFT-energies E_{train} of the training set could then be minimized.

However, DFT calculation always include some amount of numerical uncertainties. Trying to fit the DFT energies exactly could lead to overfitting, namely fitting numerical noise instead of physical interactions. To avoid this, we simply account for the numerical uncertainties ϵ when comparing the predicted energies $\vec{E}_{pred}$ and the calculated energies E_{train}.

$$\vec{E}_{train} = \vec{E}_{pred} + \vec{\epsilon} = \hat{X}\vec{\omega}^T + \vec{\epsilon} \tag{2.6}$$

We model the numerical uncertainties as a Gaussian distribution $\mathcal{N}$ with mean 0 and variance σ^2_{model}. The probability for a specific set of energies $\vec{E}_{train}$ is therefore given by a Gaussian distribution $\mathcal{N}$ with mean $\hat{X}\vec{\omega}^T$ and variance σ^2_{model}. A reasonable value for σ_{model} is given by the level of accuracy to which the DFT calculations are converged. In this study, $\sigma_{model} = 0.01\ eV$ is chosen.

$$\begin{aligned}
\vec{E}_{DFT} &= \vec{E}_{pred} + \vec{\epsilon} = \hat{X}\vec{\omega}^T + \vec{\epsilon} && (2.7)\\
p(\vec{E}_{train} \mid \vec{\omega}, \pi) &= p(\vec{\epsilon} = \vec{E}_{train} - \hat{X}\vec{\omega}^T \mid \vec{\omega}) = \mathcal{N}(\hat{X}\vec{\omega}^T, \sigma^2_{model}) &&\\
p(\vec{E}_{train} \mid \vec{\omega}, \pi) &\propto \exp\left(-\frac{(\vec{E}_{train} - \hat{X}\vec{\omega}^T)^2}{2\sigma^2_{model}}\right) && (2.8)
\end{aligned}$$

In principle, one could now take the derivative of $p(\vec{E}_{train} \mid \vec{\omega}, \pi)$ with respect to $\vec{\omega}$ and equal it to zero to find the set of parameters which fit best with the 'measured' values $\vec{E}_{train}$ – this would be a linear regression fitted with a maximum likelihood estimation. There is just one big problem: Since density functional calculations are computationally so costly, we typically have fewer training points then fitting parameters and consequently the system is under-determined.

This is where *Bayesian* comes into linear regression. Actually, even before we consider the training data we have some prior knowledge about the fitting parameters. As an example, we know that interactions between molecules are typically stronger for close pairs than for pairs which are farther apart. When we account for this, and several other, physically motivated, assumptions, some sets of fitting parameters $\vec{\omega}$ are more likely than others. Mathematically speaking, we have prior knowledge π that is encoded as a prior probability distribution $p(\vec{\omega} \mid \pi)$.

Bayesian linear regression now makes use of the Bayesian theorem (2.9) to combine this prior knowledge $p(\vec{\omega} \mid \pi)$ with the likelihood stemming from the training data $p(\vec{E}_{train} \mid \vec{\omega}, \pi)$.

$$\overbrace{p(\vec{\omega} \mid \vec{E}_{train}, \pi)}^{\text{posterior}} = \frac{\overbrace{p(\vec{\omega} \mid \pi)}^{\text{prior}}\ \overbrace{p(\vec{E}_{train} \mid \vec{\omega}, \pi)}^{\text{likelihood}}}{\underbrace{p(\vec{E}_{train} \mid \pi)}_{\text{model evidence}}} \tag{2.9}$$

The combined *posterior* probability distribution $p(\vec{\omega} \mid \vec{E}_{train}, \pi)$ tells us which set of parameters is most reasonable from a physical perspective (enforced by the prior) *AND* at the same time fit the training data well (enforced by the likelihood). Figure 2.5 illustrates this characteristic: First, the prior knowledge is used to localize the range of possible values within a comparably broad range. This step is crucial since it allows us find reasonable fitting parameters even for interactions that never show up in the training data! For all interactions that show up in the training data – and those interaction that are similar to them – we then further narrow down the range of interaction energies using the likelihood of the training data.

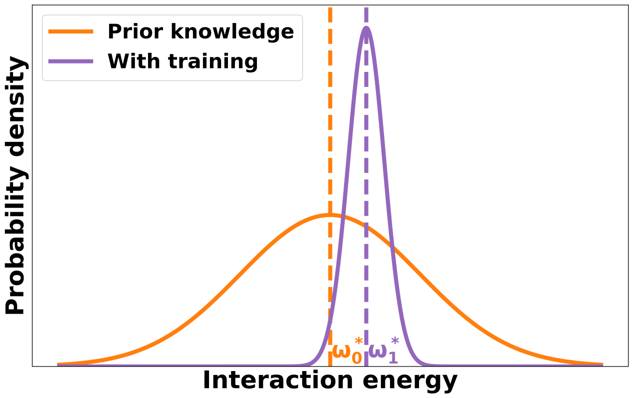

Figure 2.5.: Schematic, one-dimensional illustration of the learning process ω_0^* and ω_1^* denote the most likely interaction energies of the prior / posterior probability distribution.

Let us have another look at the Bayes' theorem (equation 2.9). It includes two new terms which we have not yet defined. First, there is the model evidence $p(\vec{E}_{train} \mid \pi)$ in the denominator. This factor becomes relevant if we want to compare different models. In the current use case, however, it cancels out as it yields the same value for all $\vec{\omega}$.

The second new term, $p(\vec{\omega} \mid \pi)$, is important as it accounts for the prior knowledge. In general, we model the prior probability distribution of each parameter as a Gaussian with mean value ω_p and a corresponding variance σ^2. Combining these Gaussian probability distributions of each $p(\omega_i \mid \pi)$ over the full vector $\vec{\omega}$ of interaction energies leads to a multivariate Gaussian (equation 2.10).

$$p(\vec{\omega} \mid \pi) \propto \exp\left(-\frac{1}{2}(\vec{\omega} - \vec{\omega}_p)^T \, \hat{C}_p^{-1} \, (\vec{\omega} - \vec{\omega}_p)\right) \tag{2.10}$$

The covariance matrix $\hat{C}_p$ is the multi-dimensional equivalent to the single-dimensional variances σ_i^2 and it additionally accounts for the correlation between different interaction energies ω_i. The covariance matrix will be explained more in detail in section 2.3.2 and 2.3.3.

Plugging this general formula for the prior into Bayes' theorem allows us to state the full *posterior* probability distribution of the fitting parameters:

$$p(\vec{\omega} \mid \vec{E}_{train}, \pi) \propto \underbrace{\exp\left(-\frac{1}{2}(\vec{\omega} - \vec{\omega}_p)^T \, C_p^{-1} \, (\vec{\omega} - \vec{\omega}_p)\right)}_{\text{prior}} \cdot \underbrace{\exp\left(-\frac{(\vec{E}_{train} - \hat{X}\vec{\omega}^T)^2}{2\sigma_{model}^2}\right)}_{\text{likelihood}} \tag{2.11}$$

Now, with the prior knowledge, the system is not under-determined anymore, and we can apply the maximum likelihood method to find that set of interaction parameters $\vec{\omega}^*$ that corresponds to the maximum of the posterior probability distribution. Consequently, the derivative of $p(\vec{\omega} \mid \vec{E}_{train}, \pi)$ with respect to $\vec{\omega}$ must be equal to zero for $\vec{\omega} = \vec{\omega}^*$.

$$\frac{\partial}{\partial \vec{\omega}} \; p(\vec{\omega} \mid \vec{E}_{train}, \pi)\Big|_{\vec{\omega}=\vec{\omega}^*} \overset{!}{=} 0 \tag{2.12}$$

As the posterior probability distribution consists of exponential functions, its derivative is given by the functions itself times the derivatives of the exponents.

$$p(\vec{\omega}^* \mid \vec{E}_{train}, \pi) \;\cdot\; \frac{\partial}{\partial \vec{\omega}} \left(-\frac{1}{2}(\vec{\omega} - \vec{\omega}_p)^T \, C_p^{-1} \, (\vec{\omega} - \vec{\omega}_p) - \frac{(\vec{E}_{train} - \hat{X}\vec{\omega})^2}{2\sigma^2_{model}} \right) \Bigg|_{\vec{\omega}=\vec{\omega}^*} \overset{!}{=} 0 \tag{2.13}$$

The first term of equation 2.13 (the posterior probability distribution) does not become zero for any finite argument[c]. Consequently, the second term (the derivative) must become equal to zero to fulfill the equation.

$$-\hat{C}_p^{-1}(\vec{\omega}^* - \vec{\omega}_p) + \frac{\hat{X}^T\left(\vec{E}_{train} - \hat{X}\vec{\omega^*}\right)}{\sigma^2_{model}} \overset{!}{=} 0 \tag{2.14}$$

$$\left(\hat{C}_p^{-1} + \frac{\hat{X}^T\hat{X}}{\sigma^2_{model}}\right) \vec{\omega}^* \overset{!}{=} \hat{C}_p^{-1}\vec{\omega}_p + \frac{\hat{X}^T\vec{E}_{train}}{\sigma^2_{model}} \tag{2.15}$$

$$\vec{\omega}^* = \underbrace{\left(\hat{C}_p^{-1} + \frac{\hat{X}^T\hat{X}}{\sigma^2_{model}}\right)^{-1}}_{\hat{C}_{post}^{-1}} \left(\underbrace{\hat{C}_p^{-1}\vec{\omega}_p}_{\text{prior}} + \underbrace{\frac{\hat{X}^T\vec{E}_{train}}{\sigma^2_{model}}}_{\text{likelihood}} \right) \tag{2.16}$$

Finally, we have found an explicit expression for $\vec{\omega}^*$, the best estimate for the set of interaction energies! From the structure of that expression we can clearly see how the prior estimates $\vec{\omega}_p$ are refined using the likelihood of the training data. Furthermore, we can identify the first term as the inverse of the new posterior covariance matrix $\hat{C}_{post}$ which again consists of a prior contribution and a contribution of the training data.

$$\hat{C}_{post}^{-1} := \hat{C}_p^{-1} + \frac{\hat{X}^T X}{\sigma^2_{model}} \tag{2.17}$$

Since both the prior and the likelihood are modeled by Gaussian functions, the posterior probability distribution can written as a multivariate Gaussian using the estimated interaction energies and the posterior covariance.

$$p(\vec{\omega} \mid \vec{E}_{train}, \pi) \propto \exp\left(-\frac{1}{2}(\vec{\omega} - \vec{\omega}^*)^T \, \hat{C}_{post}^{-1} \, (\vec{\omega} - \vec{\omega}^*)\right) \tag{2.18}$$

Due to this iterative structure of the learning process, we have two options if we later want to add more training data: Either we start again from the prior knowledge and learn on all training points at once. Alternatively, we can use the *posterior* probability distribution which already includes the old training data as new prior and then update it using the likelihood of

---

[c] Assuming $\sigma^2_{model} \neq 0$.

the new training data. Both approaches are equivalent as can be derived from the structure of the Bayes' theorem (equation 2.9). The second approach in principle saves computational time, however, as we are limited to small training sets the computational cost for the fitting process is generally negligible.

2.3.2. Defining the Priors

In the last chapter we derived a formula for the best estimate of the interaction energies $\vec{\omega}^*$ as a function of the priors and the training data (equation 2.16). The next step now is to exactly define the priors. (The other contribution – the training data – will be discussed more in detail in chapter 3.2.)

Since the prior is modeled as a multivariate Gaussian distribution, it is defined by two properties: the mean values $\vec{\omega}_p$ and the corresponding prior covariance matrix $\hat{C}_p$.

Prior Interaction Energies $\vec{\omega}_p$

Let us first have a look at the prior interaction energies $\vec{\omega}_p$. This vector actually consists of two parts: the prior one-body interactions $\vec{\omega}_{p,1body}$ and the prior two-body interactions $\vec{\omega}_{p,2body}$ (in analogy with equation 2.4).As prior of the one-body interactions $\vec{\omega}_{p,1body}$ we use the adsorption energy of a single, isolated adsorbate molecule in a large unit cell[d]. For the two-body interactions $\vec{\omega}_{p,2body}$ we have no solid *prior* knowledge. In general, the molecules might attract or repulse each other and we do not want to introduce any bias at this point. Consequently we set the prior mean of the two-body interactions $\vec{\omega}_{p,2body}$ to zero and allow for large variances σ^2_{2body} around this mean value.

Prior Covariance Matrix $\hat{C}_p$

Next, we need to define the prior covariance matrix $\hat{C}_p$. The elements on the diagonal of this matrix account for the variance of each interaction energy while the off-diagonal elements account for the correlation between each set of two interaction energies.
Similar to the mean values, the covariance matrix consists of two parts: one block for the one-body interactions and another one for the two-body interactions. The adsorption energies of local geometries (the one-body terms) are assumed not to be correlated with the pair interaction energies (the two-body terms). Therefore, $\hat{C}_p$ decomposes into a block matrix consisting of those two blocks.

$$C_p = \begin{bmatrix} \hat{C}_{1body} & 0 \\ 0 & \hat{C}_{2body} \end{bmatrix} \tag{2.19}$$

Prior Covariance of One-Body Interactions $\hat{C}_{1body}$

The diagonal elements of $\hat{C}_{1body}$ stand for the prior variance of the adsorption energy of each local adsorption geometry. Or in other words, these elements encode how sure we are about our prior guesses for the adsorption energies of single molecules.

The main uncertainty for these adsorption energies comes from an effect called charge depolarization. Upon adsorption charge can be transferred between the substrate and

---

[d] Due to the periodic nature of the unit cells, even a single molecule is affected by the interaction energy with its periodic replica. Therefore, 'large' means that the lattice vectors of the unit cell must be in the order of 1 nm such that these interactions have decayed nearly to zero.

the adsorbate molecule and the amount of this charge transfer influences the strength of the adsorption energy. This effect is stronger for a single, isolated molecule than for a molecule that is situated in the middle of a full monolayer as there are in the latter case simply less substrate atoms from which each adsorbate molecule can draw charge from.

Consequently, we can use the adsorption energies of single, isolated molecules as prior, but we have to allow for deviations between these values and the adsorption energies of molecules in full monolayers. To this end we introduce the hyperparameter σ_{1body} that controls the *prior one-body interaction uncertainty.* We set $\sigma_{1body} = 0.2\ eV$ with $0.2\ eV$ being approximately 10 % of the total adsorption energy. The adsorption energies of all local adsorption geometries have approximately the same variance and there is no correlation between different local adsorption geometries. Therefore, the prior covariance matrix $\hat{C}_{1body}$ of local adsorption geometries is given by a unitary matrix times σ^2_{1body}.

$$\hat{C}_{1body} = \sigma^2_{1body} \cdot \mathbb{I} \tag{2.20}$$

Prior Covariance of Two-Body Interactions $\hat{C}_{2body}$

The diagonal elements of the prior covariance matrix again denote the variance of each interaction. However, in contrast to the one-body interactions, assuming one identical variance for all two-body interactions would not be optimal. As noted before, even without training data we know that interaction energies that correspond to close pairs typically vary much stronger than interaction energies of far separated pairs which will be always close to zero. Consequently, we model the prior two-body variances as an exponentially decaying function:

$$(\hat{C}_{2body})_{i,i} = \sigma^*_{2body,i} \cdot \sigma^*_{2body,i} \tag{2.21}$$

$$\sigma^*_{2body,i} = \sigma^{max}_{2body} \cdot \exp\left(-\frac{d_i - d_{min}}{\tau}\right) \tag{2.22}$$

σ^{max}_{2body} here denotes the maximal variance which is reached at the minimal intermolecular distance $d_i = d_{min}$. From there on $\sigma^*_{2body,i}$ decays with a rate that is determined by the *decay length* τ. For $d_i = d_{min} + \tau$ the variance decays to $\frac{1}{e}$ of the maximum. This way, τ defines what is meant by close pairs and far separated pairs.
The choice of the hyperparameter σ^{max}_{2body} can be based on physical knowledge and experience. Pair interactions for this system are typically in the range of at most a few hundred meV. Interestingly, bigger values for that parameter seem to lead to slightly better predictions (see Appendix: B) and so we choose $\sigma^{max}_{2body} = 1\ eV$ as a conservative value allowing for a broad spread.

The decay length τ has to be chosen such that it roughly follows the actual decay of the interaction and σ^*_{2body} should decay to the range of meV for distances near the cut-off distance d_{max}. In this study, $\tau = 0.3\ nm$ is chosen as the molecule-molecule interactions in this system decay rather quickly.

Defining the prior *covariance* of two-body interactions will be a bit more complicated. At the same time this part of the prior covariance is crucial for the efficiency of Bayesian linear regression and thus the efficiency of the SAMPLE approach! Therefore, we devote the full next section to that topic.

2.3.3. Distinguishing Pairs – Feature Vectors and Correlation

When we discussed the covariance matrix of the *one-body* interactions we assumed that the adsorption energies of different local adsorption geometries are uncorrelated. This means that the off-diagonal elements of the covariance matrix $\hat{C}_{2body}$ are all zero and so each of the adsorption energies has to be optimized independently. This is of course not optimal with respect to computational time, but for local adsorption geometries there is simply no physically meaningful correlation between their corresponding adsorption energies.

For the two-body interaction the situation looks completely different. First of all, there are far more fitting parameters for two-body interactions than for the one-body interactions. Consequently, fitting them efficiently has a much bigger influence on the total efficiency compared to the fitting of the one-body interactions. But most importantly, for two-body interactions there clearly is a correlation between different interaction energies. Figure 2.6 displays two pairs of molecules which are definitely not exactly equal. At the same time these two pairs are not completely different either and consequently their interaction energies will not be not completely different, too. This is exactly what correlation means and what allows us to train our energy model with as little training data as possible and still achieve good prediction accuracy.

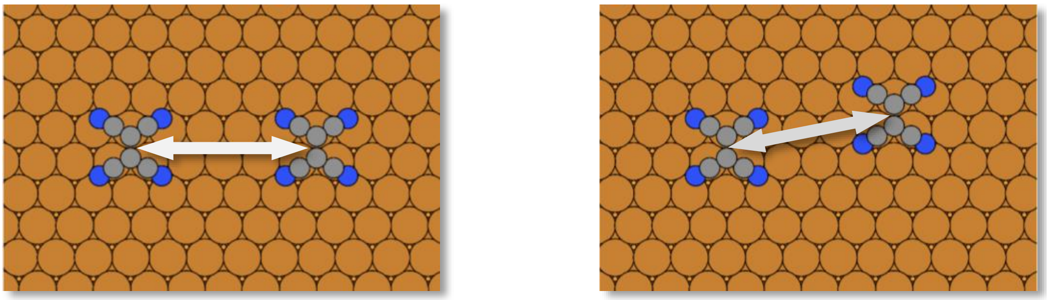

Figure 2.6.: Similar pairs of adsorbate molecules typically also have similar interaction energies. With the help of feature vectors, we can use this characteristic to improve the training efficiency.

However, to use this kind of knowledge, we have to define 'similarity of pairs' in a mathematical way. To make this possible we first need to find some kind of fingerprint to identify each distinct pair – this 'fingerprint' is commonly called a feature vector. With that we can then define a measure for the similarity between pairs.

Such a feature vector should obey the following criteria:

Uniqueness

Different pairs must lead to different feature vectors. This criterion is compulsory since otherwise we are not able to reliably distinguish different pairs which might have different interaction energies.

Conservation of symmetries

Pairs which are physically identical should always lead to the same feature vector. E.g. if one pair can be transformed into another one by a symmetry transformation of the substrate, these pairs are equivalent and should be described by one common fitting parameter. This condition is not compulsory, but it helps minimizing the number of fitting parameter and therefore improves the efficiency of the SAMPLE-algorithm.

Similarity

Similar pairs should lead to similar feature vectors. This is necessary to correctly account for correlation of interaction energies.

In this study an approach similar to Coulomb matrices[40] is used. On each molecule we choose a set of 'marker' atoms. Choosing the outermost ones proved to be effective (for TCNE: the four nitrogens). Then the distances d_α between each of those atoms on one molecule to each of the corresponding atoms on the second molecule are calculated as illustrated in Figure 2.7.

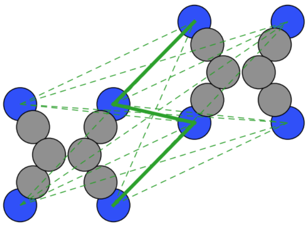

Figure 2.7.: Illustration of the distances used for construction of the feature vector. Shortest distances are highlighted for visibility but for this system currently all distances are taken into account.

A problem of Coulomb matrices is that they are not invariant with respect to rotation. If we rotate one of the interacting TCNE molecules by 180 degrees, their interaction energy must not change as this is a symmetry operation of TCNE. However, to construct a Coulomb matrix the atoms of the involved molecules have to be enumerated and, consequently, the list of distances changes if the order of the atoms is changed. We avoid this problem by sorting the obtained distances by their length in ascending order.

At this point, it is possible to use only the first n_{crop} distances to save computational time. However, for the present system, that step is not necessary and therefore all 16 distances are used.

The pair interactions typically change far more rapidly in the range of small intermolecular distances. Consequently, we need to have a finer resolution for such close-range features than for features corresponding to far separated pairs. To achieve that, we model the feature values f_α as an inverse power function of the pair distance (see equation 2.23). Furthermore, the distances are scaled by the minimal intermolecular distance d_{min}. This way the feature values f_α can range from one for $d_\alpha = d_{min}$ to approximately zero for $d_\alpha \gg d_{min}$. The exponent (*decay power* n) must be chosen greater than one to yield the desired effect on the feature resolution. Higher values for n result in more aggressive scaling. Using a decay power n in the range of two to three empirically proofed to work best (see Appendix B).

$$f_\alpha = \left(\frac{d_{min}}{d_\alpha}\right)^n \tag{2.23}$$

Using equation 2.23 each pair of molecules is now represented by vector $\vec{f}$ in the feature space. The dimensionality of this feature space is determined by the number of interatomic distances which is 16 in this study. Defining a similarity between these pairs is now rather straightforward. We use the L1-norm, also known as Manhattan norm, to measure the distance between two feature vectors in the feature space and then divide that distance by the *correlation length* ξ.

The correlation length is a hyperparameter that scales the feature space. Small values for ξ lead to less correlation between features whereas high values for ξ introduce correlations even for features which are not that similar. For the diagonal elements $(\hat{C}_{pair})_{i,i}$ (the variances) both features are identical, and this term becomes one.

With that the covariance of two pairs i and j can be written as the product of their standard deviations times their similarity which is given by an exponential function of their scaled distance in feature space.

$$(\widetilde{\hat{C}_{2body}})_{i,j} = \underbrace{\sigma^*_{2body,i} \cdot \sigma^*_{2body,j}}_{variances} \cdot \underbrace{\exp\left(-\frac{||\vec{f}_i - \vec{f}_j||_{L1}}{\xi}\right)}_{similarity} \tag{2.24}$$

If there were only flat-lying molecules, we could immediately use equation 2.24 to calculate the covariance. However, in order to study densely-packed polymorphs, we need to include upright-standing molecules as well. This leads to a problem which is illustrated in Figure 2.8. Since the feature vectors depend only on the relative position of the molecules, but not on their absolute position, they cannot distinguish a pair of flat-lying molecules from a pair of two upright-standing molecules. This fact would violate the first criterion of the feature vectors – the uniqueness. Flat-lying adsorbate molecules and upright-standing ones have different electronic structures and, consequently, we cannot assume that they result in identical two-body interactions.

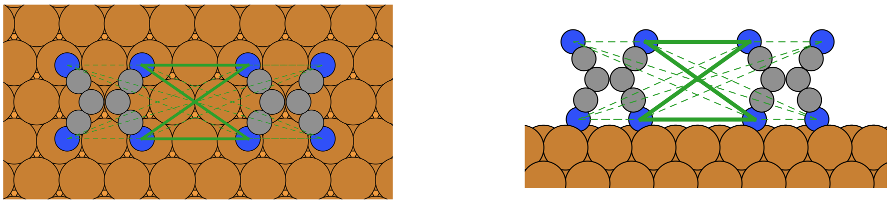

Figure 2.8.: Feature vector of a pair of flat-lying molecules and a pair of standing molecules. Note the nearly identical geometry of the distance vectors.

To overcome this problem, three different types of pairs are defined: *flat*, *mixed* and *standing*. The *flat* and *standing* groups contain pairs consisting of two flat-lying molecules or two upright-standing molecules, respectively, while the *mixed*-type is used for interactions of upright-standing molecules with flat-lying molecules. We then introduce a factor $B_{i,j}$ and set it to zero if the two pairs i and j have different types and one for identical types.

$$B_{i,j} = \delta\left(type_i,\ type_j\right) \tag{2.25}$$

This procedure removes the correlation of those pairs which cannot be distinguished by the feature vector, or in other words: It tells the SAMPLE algorithm to learn those pairs independently. With this modification, the prior covariance matrix splits up into three uncorrelated block matrices corresponding to the three types of pairs (see Figure 2.9).

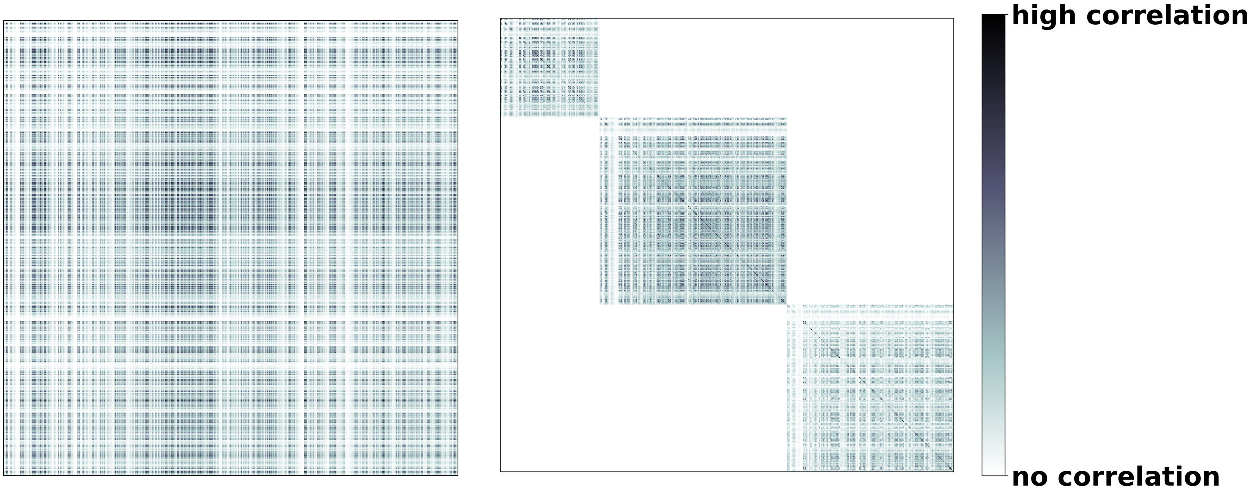

(a) *correlated* covariance matrix without factor $B_{i,j}$

(b) *blockwise* covariance matrix including factor $B_{i,j}$

Figure 2.9.: To ensure the uniqueness of feature vectors the covariance between different types of interactions is set to zero.

With this new factor, the covariance $(\hat{C}_{2body})_{i,j}$ between each pair of two-body interaction energies reads as following:

$$(\hat{C}_{2body})_{i,j} = \underbrace{\sigma^*_{2body,i} \cdot \sigma^*_{2body,j}}_{variances} \cdot \underbrace{\exp\left(-\frac{||\vec{f}_i - \vec{f}_j||_{L1}}{\xi}\right)}_{similarity} \cdot \underbrace{B_{i,j}}_{\text{split types}} \tag{2.26}$$

One technical detail is yet missing: To obtain all features, we create all possible pairs of two adsorbate molecules in the discretization of the adsorption sites. However, there is no point in using different fitting parameters if two pairs are extremely similar. Therefore, if two pairs lead to feature vectors which are separated in the feature space by less than a minimal feature distance Δf, both are fitted with the same two-body interaction energy. In this study a minimal feature distance of $\Delta f = 0.02$ is chosen. A summary of all chosen hyperparameters is given in Table 2.1 and additional data describing the sensitivity of these parameters can be found in Appendix B.

parameter	**symbol**	**value**
prior one-body interaction uncertainty	σ_{1body}	0.2 eV
prior two-body interaction uncertainty	σ^{max}_{2body}	1.0 eV
real space decay length	τ	0.3 nm
model uncertainty	σ_{model}	0.01 eV
feature correlation length	ξ	1.0
decay power	n	-2
feature threshold	Δf	0.02
minimal intermolecular distance	d_{min}	0.26 nm
maximal intermolecular distance	d_{max}	1.6 nm

Table 2.1.: Chosen set of hyperparameters

3. Increasing the Packing Density

A central aim of this thesis is to study how the adsorbate layer changes when the coverage is increased. To that end, we first find the adsorbate layers that form at sub-monolayer coverage. At such low coverage, the energetically most favorable configurations consist of flat-lying molecules, since this is the preferred orientation of single, isolated adsorbate molecules. When we increase the coverage, the molecules are forced into configurations with smaller intermolecular distances. At some point molecules start to flip from the flat-lying orientation to an upright-standing orientation as these standing molecules have a smaller footprint on the surface than the lying ones and consequently the standing molecules can pack more densely at the same intermolecular distances.

Unfortunately, these standing molecules are more delicate than the flat-lying ones. As a simple, geometric example, when we place two flat-lying TCNE molecules next to other, not more than two cyano groups – which account for the strongest interactions – can come into close range. On the other hand, in exactly upright-standing TCNE molecules, there are always two cyano groups exactly above each other. Consequently, when we place two standing TCNE molecules next to other, instead of one or two cyano groups, two or four of them come into close contact. This implies, that also the interactions of those molecules increase far more rapidly with decreasing intermolecular distance.

To investigate and overcome this problem, a test system is built (see chapter 3.1) such that we can cheaply create lots of training data and, based on that, study different training set selection strategies (see chapter 3.2).

3.1. Test System 'Gas Phase'

Preliminary tests with the existing algorithm showed that the prediction quality was often not sufficiently stable. In particular, the prediction quality fluctuated heavily find different training sets of identical size. This is a problem, as it complicates all evaluations that rely on a stable prediction quality. For instance, when we search for an optimal set of hyperparameters, we need to resolve rather small changes of the prediction accuracy. When the prediction quality, however, fluctuates heavily with different training sets, we cannot resolve those small differences anymore. On the other hand, this problem is also a chance, as it shows that a well-chosen training set can further improve our prediction quality.

To investigate this problem more systematically, we need to evaluate the adsorption energy via density functional theory (DFT) for a large set of configurations. From that pool of calculated configurations, we could then choose several different training sets and compare their resulting prediction accuracy. However, conducting a large set of simulations on the full surface system would result in immense computational costs. Therefore, a simplified test system is created for these investigations: The configurations of this test system are at first build as described in section 2.2.1 but then the substrate is removed from the final geometry before starting the

DFT calculations. Hence this system can be seen as a two-dimensional monolayer of molecules in vacuum. Consequently, all interactions with the surface vanish, but all molecule-molecule interactions persist. Note that the discretization of the local adsorption geometries is still used – i.e. each single molecule must still assume one of the local adsorption geometries that it would assume on the substrate surface.

For all unit cells with 3 to 36 surface atoms such 'gas phase'-configurations are built using a minimal intermolecular distance $d_{min} = 0.26\ nm$. For up to two molecules, this results in approximately $3 \cdot 10^5$ configurations. From those configurations, two sets of 1000 configurations each are randomly chosen. The first set (*flat*) contains only flat-lying molecules (with respect to the removed surface) and the second one (*standing*) contains only standing molecules.[a]

Each of these sets is then split up into a training pool of $n_{pool} = 500$ configurations from which the training sets will be chosen as well as a test set of $n_{test} = 500$ configurations which will be solely used to evaluate the prediction accuracy. These two sets are chosen randomly but kept fixed throughout all further evaluations.

To compare the nature of the *flat* set and the *standing* set, we now choose ten randomly drawn training sets from each of their respective training pools.[b] With each of these training sets the energy model is fitted and with that the adsorption energy $E_{pred,i}$ of each configuration in the test set is predicted. Each of those predicted energies is then compared to the adsorption energy of that configuration evaluated directly via first principles calculation $E_{DFT,i}$. From that difference the root mean square error (RMSE) is calculated.

$$RMSE = \frac{1}{n_{test}} \sqrt{\sum_i \left(E_{pred,i} - E_{DFT,i}\right)^2} \tag{3.1}$$

When we now compare the prediction accuracy that we obtain from training sets of those two different types of configurations (Figure 3.1), we clearly see that the flat-lying molecules lead to a rather small spread and stable, accurate predictions for all training sets containing more than about 100 configurations. On the other hand, the training on the upright-standing molecules has severe problems. Sometimes we end up with a poor prediction accuracy even with training sets of 200 configurations! But at the same time, we see that good training sets *do exist* for this type of configurations. So we need to make sure that we always choose such a good training set.

But before that, let us quickly have a look on the interaction energies that occur in those two different sets. A histogram of those interaction energies is shown in Figure 3.2. Most configurations have interaction energies of less than 100 meV (the y-axis is truncated at 30 counts). However, we also see some configurations with exceptionally high interaction energies $E \geq 200\ meV$ and nearly all of those configurations consist of standing molecules. These configurations contain pairs of standing molecules with minimal intermolecular distances just slightly above the threshold d_{min}. As we have discussed before, the repulsion between such pairs of standing molecules increases particularly rapidly due to the higher number of interacting cyano groups. Additionally, due to the discretized local adsorption geometries, only specific intermolecular distances can exist. This leads to the large energetic gaps between configurations

---

[a] A third set of 1000 configurations with each of them consisting of exactly one standing and one lying molecule per unit cell has been chosen, too. The characteristics of this set are always somewhere between the other two, which is why it is not further used for any of the evaluations.

[b] These training sets are chosen subsequently. I.e., a training set of n configurations consists of the training set with $(n-1)$ configuration plus one new configuration.

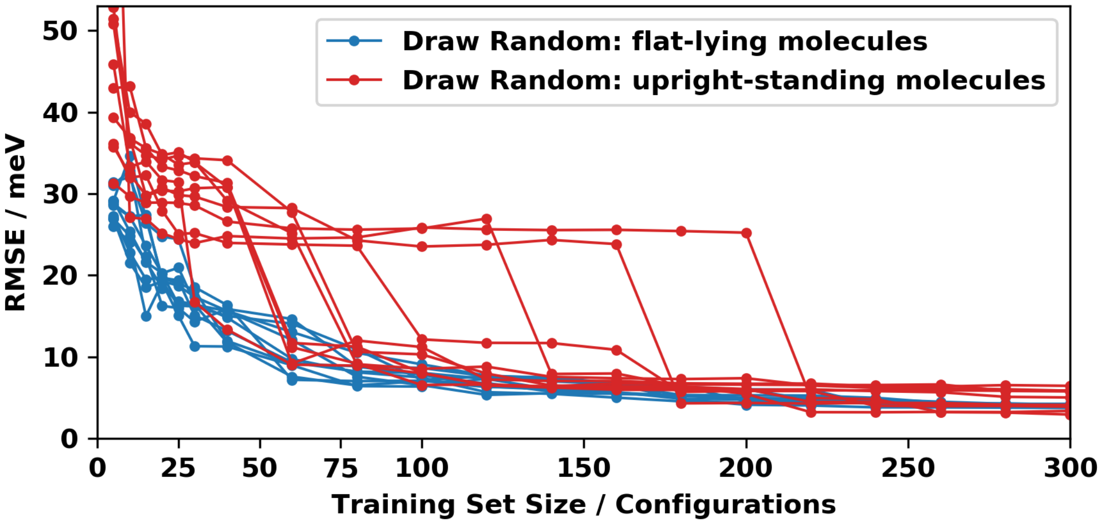

Figure 3.1.: Comparison of prediction accuracy as a function of the training set size for configurations consisting of flat-lying and upright-standing molecules, respectively.

as we see it for instance in the range of 250 meV to 300 meV. A way to deal with these problems is to improve the training set selection such that the configurations of these training sets include as many and as diverse interactions as possible.

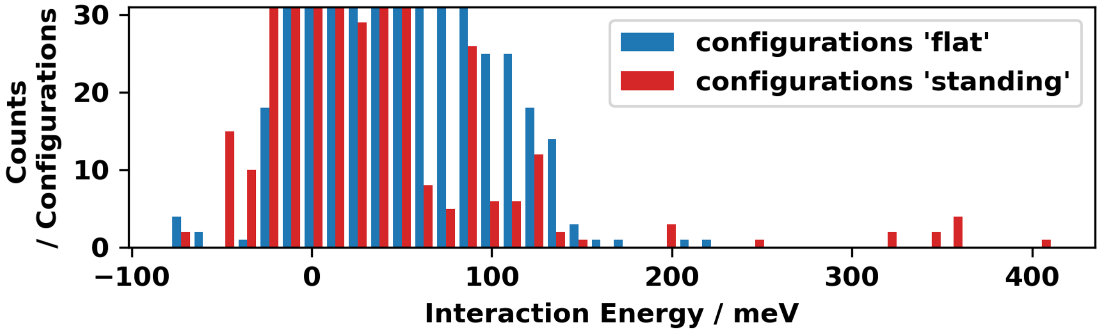

Figure 3.2.: Distribution of interaction energies in the *flat* and *standing* sets of the 'gas phase' system. Count is truncated at 30.

3.2. Training Set Selection

The following section deals with the challenge of choosing the most efficient training set, i.e. a training set which allows us to improve the prediction quality as much as possible while needing as few DFT calculations as possible.

To that end, we need some indicator of the quality of a training set. Based on that indicator we can then can rank the training sets, select the best one and *then* calculate their adsorption energies using DFT.

The accuracy of the predicted energies would be the most straightforward indicator, but of course the exact prediction accuracy is only known as soon as we have conducted the training and checked the predictions on a test set. So, we would have to first invest a lot of cpu time and could only afterwards see if this training set was indeed a good choice.

Fortunately, from the fitting process of Bayesian linear regression we obtain not only the set of most likely fitting parameters (the interaction energies ω^*), but we additionally receive the full posterior probability distribution $p(\vec{\omega} \mid \vec{E}_{train}, \pi)$ for those parameters! On top of that, the posterior covariance depends only on the interaction which occur in the configurations of the training set but not on the adsorption energies of these configurations - i.e. we do not need the actual DFT calculations to know the shape of $p(\omega \mid E_{train}, \pi)$! With that in mind, we can propose multiple training sets and calculate the posterior covariance matrix that we would gain from each of them. Based on that posterior covariance matrix we can then choose the optimal training set.

A schematic illustration of the posterior probability distribution for the strongly simplified case of a single fitting parameter ω, is given in Figure 2.5. In this one-dimensional illustration the shape of $p(\vec{\omega} \mid \vec{E}_{train}, \pi)$ is given by its variance and it is obvious that we have to minimize this variance to maximize the accuracy of the corresponding fitting parameter. Since the variance is a real valued scalar, there exists a well-defined order relation. Consequently, if there was only a single fitting parameter, we would simply choose that training set which decreases the posterior variance of that fitting parameter the most.

However, when we leave this minimal example and increase the number of parameters from one to two, the situation immediately becomes more complicated. For two parameters[c], the posterior probability distribution becomes a bivariate Gaussian as depicted in Figure 3.3. Furthermore, the scalar variance σ^2 of the one-dimensional example becomes a matrix: the posterior covariance matrix $\hat{C}_{post}$ which we have already seen in section 2.3.1 (equation 2.17).

The diagonal elements of $\hat{C}_{post}$ correspond to the squared variance of the posterior probability distribution $p(\omega_i \mid \vec{E}_{train}, \pi)$ of each fitting parameter. The off-diagonal elements indicate the correlation of the fitting parameters. I.e., a positive (negative) off-diagonal element $(C)_{1,2}$ implies that a large value for ω_1 increases (decreases) the probability for a large value for ω_2. The reverse implication (from ω_2 to ω_1) is true as well since $\hat{C}_{post}$ is symmetric. Please bear in mind that the elements of the covariance matrix are – in contrast to the *correlation* matrix – in general not normed to the range of -1 to $+1$.

One way to understand the covariance matrix better, is to find its eigenvalues λ_i and eigenvectors $\vec{\nu}_i$ by solving the eigenvalue problem $\hat{C}_{post}\, \vec{\nu} = \lambda \vec{\nu}$. The eigenvectors $\vec{\nu}$ can be understood

---

[c] The general case of n fitting parameters does not differ substantially from the presented special case of $n = 2$ anymore - it just becomes harder to visualize, which is why here the two-dimensional case is presented.

in the following way: If we would create a new parameter from a linear combination of fitting parameters with weights defined by an eigenvector $\vec{\nu}$, the resulting, combined parameter would have a variance equal to that eigenvalue λ which corresponds to $\vec{\nu}$. This characteristic is visualized in Figure 3.3 for two different covariance matrices. The ellipses in this plot depict lines of constant probability of the bivariate probability distribution defined by $\hat{C}_{post}$.

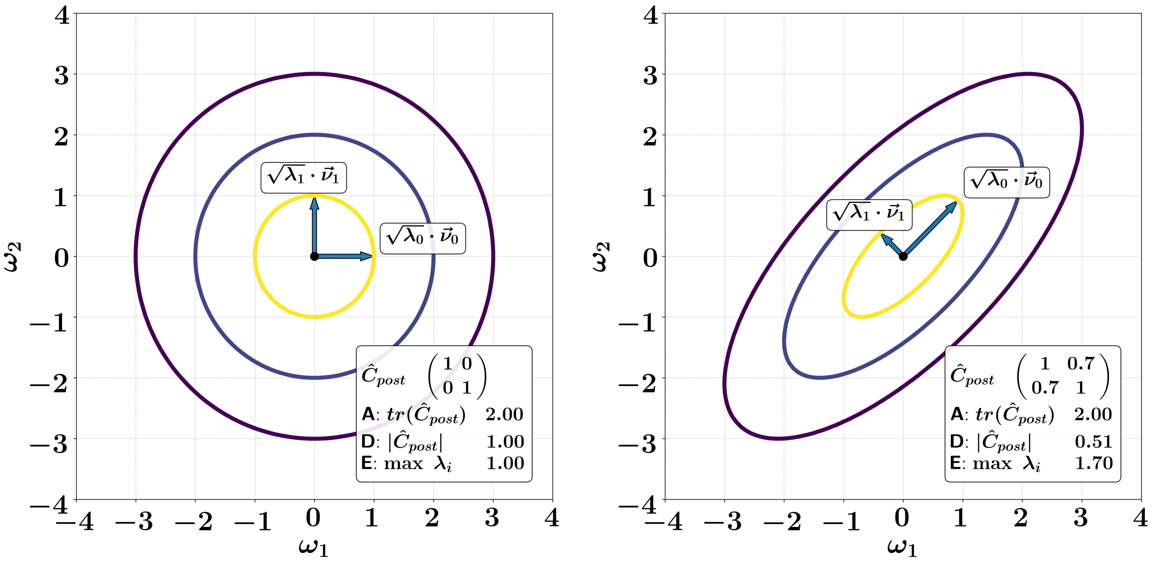

Figure 3.3.: Iso-probability contours at 1σ, 2σ, and 3σ of two bivariate Gaussian distributions defined by different covariance matrices $\hat{C}_{post}$.

With that, let us go back to the original question: how can we define which posterior covariance matrix – and therefore which training set – is the best one? As there exists no unique order relation for matrices in general, there exists no such definition for the covariance matrix, either. Consequently, in the multi-dimensional case, we cannot decide which is the best training set as straightforwardly as in the one-dimensional case. However, there are some reasonable choices (*optimality criteria*) to define such an order relationship which will be presented in the following.

A minor detail is yet missing before we can dive into the zoo of optimality criteria: Up to now, we implicitly assumed that the posterior covariance matrix $\hat{C}_{post}$ is obtained from some fixed training set with model matrix $\hat{X}$. As we now deal with different training sets, we have to use a more explicit notation.

$$\hat{C}^{-1}_{post}(\xi_t) = \hat{C}^{-1}_{p} + \frac{\hat{X}^T(\xi_t)\,\hat{X}(\xi_t)}{\sigma^2_{DFT}} \tag{3.2}$$

Furthermore, we will use $\Psi_\alpha(\xi_t)$ to denote the score of a proposed training set ξ_t evaluated with a particular optimality criterion α. In all presented optimality criteria, the training set that achieves the lowest score is considered as the best one.

A-Optimality

The A-optimality criterion aims to minimize the **A**verage variance of the fitting parameters [41]. This is achieved by minimizing the trace of the covariance matrix which is equivalent to the sum of the eigenvalues λ.

$$\Psi_A(\xi_t) = tr(\hat{C}_{post}(\xi_t)) = \sum_i^d \lambda_i \tag{3.3}$$

This criterion does not take into account the off-diagonal elements of the covariance matrix. Therefore, it yields the same score for training sets where the occurring interactions are highly correlated as for a training set with uncorrelated interactions. This is demonstrated by by the identical A-optimality-scores of the different covariance matrices in Figure 3.3.

D-Optimality

The D-optimality criterion aims to minimize the *generalized* variance of the parameter estimates. This is achieved by minimizing the product of the eigenvalues λ of the covariance matrix or equivalently minimizing its **D**eterminant [42, 43].

$$\Psi_D(\xi_t) = \prod_i^d \lambda_i = |\hat{C}_{post}(\xi_t)| \tag{3.4}$$

It is therefore proportional of the area of the iso-probability curve and so this criterion minimizes the region of possible values for the fitting parameters.

An equivalent, but maybe more comprehensible interpretation of the D-optimality can be derived using the information matrix $M(\xi_t)$ which is a direct measure of the knowledge that we obtain from the prior and the training set combined, compared to a completely flat prior (= a constant probability distribution). Obviously, we want to maximize this information gain.
The information matrix $M(\xi_t)$ is defined as the inverse of the covariance matrix and, hence, the *inverse* D-optimality is given by the determinant of $M(\xi_t)$ (equation 3.5). Each configuration that is added to the training set increases $|\hat{X}^T(\xi_t)\hat{X}(\xi_t)|$ and, therefore, the determinant of the information matrix. In that sense, D-optimality corresponds to that training set (of constant size) that increases the information gain, measured by $M(\xi_t)$, the most. Apart from its descriptive benefit, this formulation of D-optimality saves a costly computational step as we directly *maximize* the determinant of the inverted posterior covariance matrix $\hat{C}^{-1}_{post}(\xi_t)$ (that we have to calculate anyway) instead of applying another inversion just to *minimize* the determinant of the resulting $\hat{C}_{post}$. This equivalence of maximization and minimization of the inverse is a specific property of the determinant[d], which is why this trick cannot be applied to the other optimality criteria.

$$\Psi_D^{-1}(\xi_t) = \left|\hat{C}^{-1}_{post}(\xi_t)\right| = |M| = \left|\hat{C}_p^{-1} + \frac{\hat{X}^T(\xi_t)\hat{X}(\xi_t)}{\sigma^2_{DFT}}\right| \tag{3.5}$$

E-Optimality

The E-optimality criterion minimizes the maximal or **E**xtremal variance of parameter estimates.

$$\Psi_E(\xi_t) = \max_i \lambda_i \tag{3.6}$$

---

[d] Compare to equation 3.4: The eigenvalues λ'_i of the inverse covariance $\hat{C}^{-1}_{post}(\xi_t)$ are given by the $1/\lambda_i$. The determinant of $\hat{C}^{-1}_{post}(\xi_t)$ is then given by the product of those eigenvalues $\lambda'_i = 1/\lambda_i$. Furthermore, inversion and product operator can be swapped, i.e. the product of the inverse eigenvalues is equivalent to the inverse of the product of eigenvalues. For a sum (A-optimality) or maximum function (E-optimality) instead of the product, this correspondence does not hold.

> Consequently, this criterion should first reduce the uncertainty of the interaction energies of close pairs as they have the highest prior uncertainty. However, this criterion might then underestimate the combined influence of far separated pairs, which have low prior uncertainties but occur more often.

All the optimality presented above have in common that they aim to minimize the uncertainty of the fitting parameters, i.e. the one-body and two-body interaction energies. However, these parameters are just an intermediate step towards the predicted adsorption energies of full configurations. So, we could instead directly aim to minimize the prediction uncertainty of configurations.

The predicted adsorption energy of one configuration x is given by a sum of the interaction energies $E_{pred}(x) = \vec{N}(x)\ \vec{\omega}^T$. In that sense, a configuration is similar to an eigenvector of the covariance matrix. As we have seen at the beginning of this section, the variance of an eigenvector corresponds to the variance of the multivariate Gaussian in the direction of that eigenvector. In analogy to that, we project the multivariate Gaussian onto the model vector $\vec{N}(x)$ to obtain the marginal probability distribution for a configuration x. The variance of that probability distribution then tells us what prediction uncertainty we can expect on that configuration x if we train on a specific training set ξ_t.

$$\sigma^2\left(\xi_t, x\right) = \vec{N}^T(x)\ \hat{C}_{post}(\xi_t)\ \vec{N}(x) \tag{3.7}$$

Of course, we do not want to base the training set selection on the prediction of a single configuration. Instead, we define an *evaluation set* ξ_e on which the expected prediction accuracy is evaluated. So we apply equation 3.7 and obtain a vector that consists of the expected variances of each configuration of the evaluation set ξ_e assuming that the energy model has been fitted on the proposed training set ξ_t.

$$\overrightarrow{\sigma^2}\left(\xi_t, \xi_e\right) = \left[\sigma^2\left(\xi_t, x_1\right) \quad \dots \quad \sigma^2\left(\xi_t, x_i\right) \quad \dots \quad \sigma^2\left(\xi_t, x_{N_e}\right)\right] \tag{3.8}$$

In order to use these variances for the training set selection, we have to make two more decisions:

1. For the full evaluation set, we obtain a vector of variances. Similar to matrices, there are multiple, different ways to define an order relation for vectors. Possible properties to define for such an order relation are, for instance, the sum (or equivalently, the average) or the maximum of the variances.

2. The choice of the evaluation set ξ_e. The basic choice would be to use all created configurations, which would consequently lead to a good prediction accuracy in the whole search space. However, we could also choose ξ_e such, that the training is tailored onto a specific range of configurations. For instance, when we are most interested in configurations of high coverage, we would use such configurations for the evaluation set and hence end up with a training set that leads to most accurate predictions specifically on that region of the search space.

G-Optimality

For the G-optimality criterion, we sort the training set by the maximal prediction variance of all configurations of the evaluation set. Hence, we aim to minimize the deviation of the worst prediction of that evaluation set and in that sense this criterion is the counterpart to the E-optimality criterion.

$$\Psi_G\left(\xi_t, \xi_e\right) = \max_{x_i \in \xi_e} \sigma^2\left(\xi_t, x_i\right) \tag{3.9}$$

V-Optimality

V-Optimality (also known as IV-Optimality) is similar to the A-optimality criterion as it aims to minimize the **I**ntegrated **V**ariance of the predictions. Originally, this criterion is defined on a *continuous* design space R, in contrast to our discretized design space ξ [43].

$$\Psi_V\left(\xi_t, R_e\right) = \int_{x \in R_e} \sigma^2(\xi_t, x)\, \mu(x)\, dx \tag{3.10}$$

Where $\mu(x)$ stands for the density at the point $x \in R_e$. In our discrete design space these points x_i become configurations. Consequently, the continuous region of interest R_e is replaced by the evaluation set ξ_e and the integral becomes a sum. The V-optimality score for a training set ξ_t and a *discrete* evaluation set ξ_e therefore reads:

$$\Psi_V\left(\xi_t, \xi_e\right) = \sum_{x_i \in \xi_e} w(x_i)\, \sigma^2(\xi_t, x_i) \tag{3.11}$$

The weights $w(x_i)$ are the counterpart to the densities $\mu(x_i)$ and can be used to put more weight on specific configurations. Furthermore, when we divide $\Psi_V\left(\xi_t, \xi_e\right)$ by the size of the evaluation set, we obtain the more intuitive average variance for that specific combination of training- and evaluation set [41]. Of course, this does not change the outcome of this criterion.

In combination with the choice of the evaluation set ξ_e, these optimality criteria could be used together in an interesting way: One could first concentrate on exploration – i.e. to limit the worst-case uncertainty, a first training set is chosen using the G-optimality criterion in conjunction with a large evaluation set that is representative for all possible configurations. Based on that, a subsequent exploitation step could be done: A second training set would be chosen based on the V-optimality criterion in combination with a smaller evaluation set of promising, low energy configurations. This way, one would end up with a rough estimate of the full potential energy surface and a precise estimate in the most important range of low energy configurations.

With all that, one should not forget that this broad exploration is actually a strength of the SAMPLE approach. In contrast to other machine learning methods that quickly concentrate their evaluation area around the lowest minima, the SAMPLE approach is able to predict the whole potential energy surface. From that, one can not only find the global minimum, but also other local minima which correspond to metastable states and can be important as well (e.g. in the form of defects).

3.2.1. Fedorov Exchange algorithm

In principle, one could now do an exhaustive search by creating all possible training sets, calculating the optimality of each of them using one of the presented criteria and finally use that proposed training set which achieves the best score. However, the number of possible training sets having k configurations, chosen from a pool of n configurations, is given by the binomial coefficient $\frac{n!}{(n-k)!\,k!}$. If we want to choose a set of 100 configurations from a pool of 300 thousand possible configurations there are approximately $5 \cdot 10^{368}$ possible choices.[e] It is, therefore, absolutely impossible to directly evaluate each possible training set.

An alternative, iterative approach by Fedorov [44] is used to overcome this problem. A sketch of our implementation of Fedorov's algorithm is given in Figure 3.4.

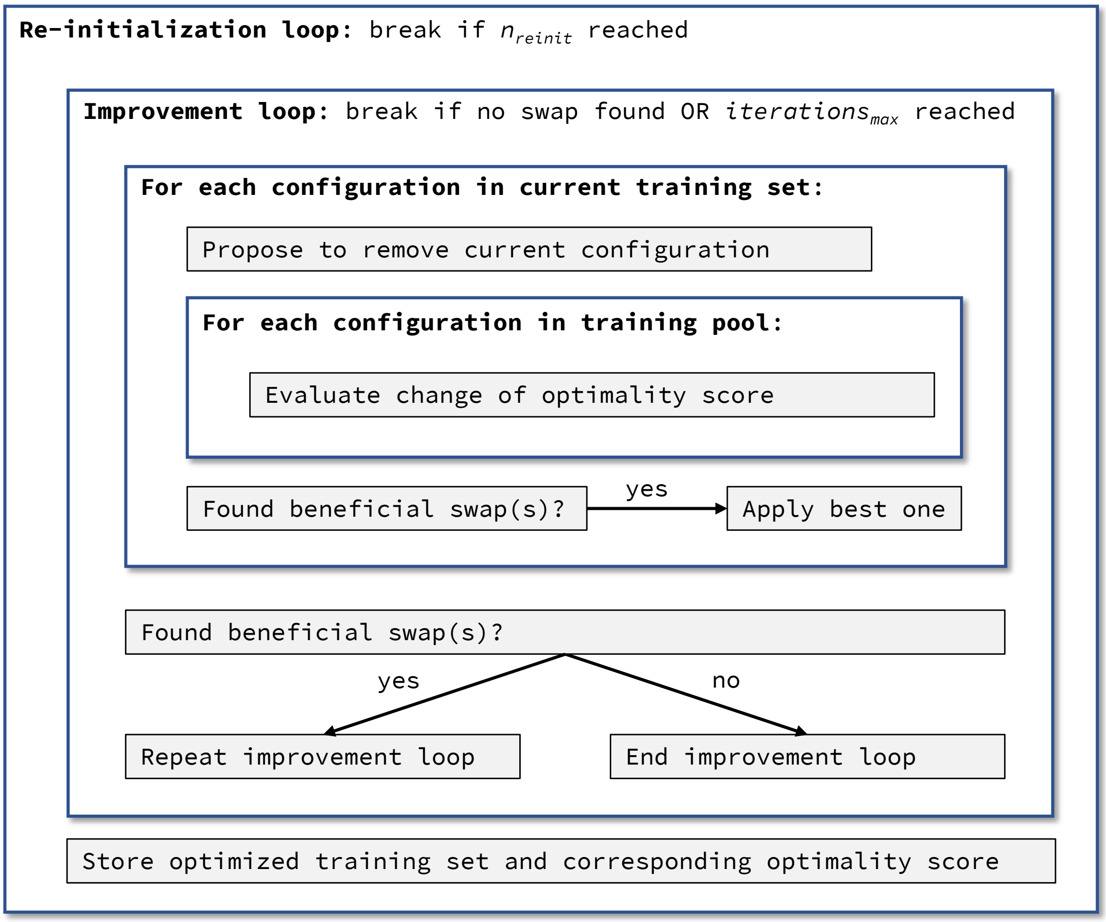

Figure 3.4.: Implementation of the Fedorov exchange algorithm

To start the algorithm, we propose a randomly chosen, initial training set (**re-initialization loop**). This set is now denoted as current (proposed) training set which we aim to optimize. Furthermore, we define the training pool as the set of all configurations that can be added to the current training set.

Now, we subsequently select **each configuration from the current training set**. For each of these configurations, we then we propose **each configurations of the training pool** as replacement for the selected configuration. Each of those *proposed* swaps is rated using the chosen optimality criterion. As soon as we have proposed each configuration of the training pool once, we check whether at least one of those proposed swaps would be beneficial – i.e. if it would lead to a better score according to the chosen optimality criterion. If one or multiple

---

[e] For the studied system there are 300 048 configurations with a maximum of two molecules in the unit cell. Since the energy model is truncated after the two-body interactions, all interactions can be found in the subset of those configurations which have two molecules in their unit cell.

beneficial swaps are found, we immediately apply the best one of those This search for a beneficial swap is repeated for each configuration of the current training set.

When we have finished the last configuration of the current training set, we check whether beneficial swaps were found during the loop over the current training set. If at least one beneficial swap was found – and, consequently, the current training set was improved during that loop – the whole **improvement loop** is repeated. If no beneficial swap was found, the current training set has reached a local minimum with respect to one-configuration swaps and so we can end this loop. Additionally, we stop improving the current training set if a maximal number of allowed iterations has been reached. In this study, the maximal number of allowed iterations is set to 10. However, most swaps happen during the first iteration and usually the search converged after three or four iterations, therefore this threshold is hardly ever reached.

Now, we have a training set that constitutes a local minimum as it cannot be improved any further by swapping any single configuration with one from the training pool. It might, however, still be possible to jump from the current locally optimized training set to a better one by swapping two or more configurations simultaneously. Though, including these multi-configuration swaps into Fedorov's algorithm would increase the computational complexity immensely as we would have to consider all combination of configurations: In the extreme case of swapping the whole training set, we would be back to the original $5 \cdot 10^{368}$ possible choices! Therefore, we instead repeat the whole procedure from several randomly-chosen initial training sets and use the best optimized training set found in those runs. The validity of this approach can be checked immediately by comparing the outcomes of these runs. It turned out that most of the time most runs lead to the same minimum indicating that this one is indeed the global minimum. Moreover, even training sets that correspond to other local minima yield optimality scores far better than almost all random initial sets – i.e. these locally optimal training set are nearly as good as the set that constitutes the global minimum.

Apart from this re-initialization step, the algorithm is fully deterministic. This property is highly valuable – e.g. for the comparison of different hyper parameters we have to avoid any noise coming from randomly-chosen training sets.

3.3. Comparison of Optimality Criteria

The optimality criteria presented in chapter 3.2 were implemented[f] and will now be benchmarked based on the training sets they select from the training pool. As we have seen, the *standing* set as it is introduced in chapter 3.1 is the most interesting one, hence this one is used for the following evaluation.

For each optimality criterion, sets of n_{train} configurations are chosen independently with n_{train} increasing stepwise from 5 to 100. To ensure that the training set selection start from the same initial training sets for each optimality criterion, the pseudo-random number generator that selects these initial sets is re-seeded to the same value (42) whenever we switch to a new optimality criterion.
As evaluation set for the G- and V-optimality criterion, we use the training pool from which the training sets are chosen. The training pool is completely separated from the test set – on

---

[f] The D-optimality criterion has been implemented before by Michael Scherbela [45] but, due to a bug, not all possible swaps were considered and the resulting training sets were, therefore, not fully optimized.

which the RMSE is evaluated – and so this choice of the evaluation set represents the general case and does not introduce any kind of overfitting.

The evolution of the prediction accuracy that results from these training sets is displayed in Figure 3.5. Five series of randomly drawn training sets have been added for comparison.

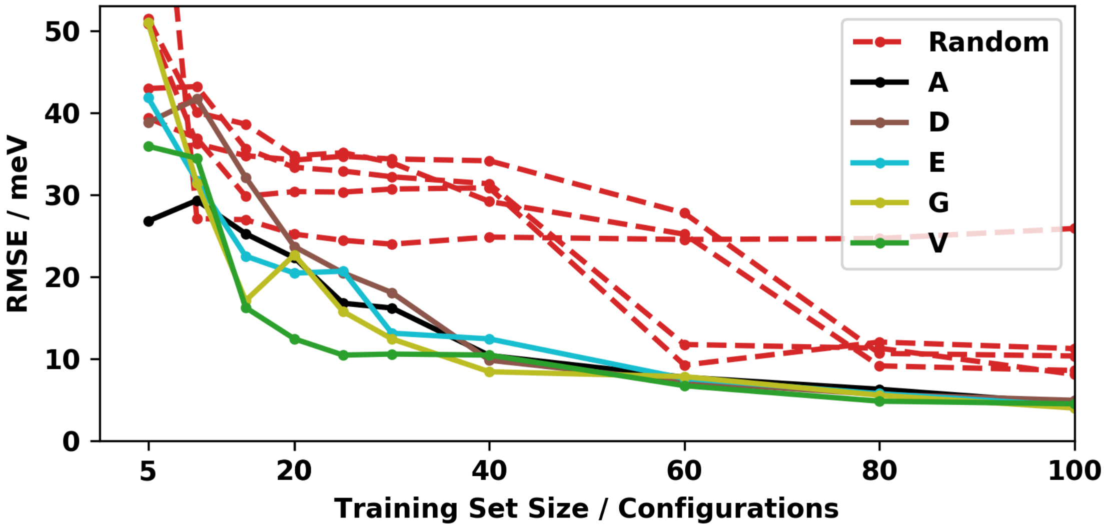

Figure 3.5.: Evolution of prediction accuracy for different training set selection strategies.

The main insight that we gain from this comparison, is that all optimality criteria perform significantly better than random selection for all training set sizes $n_{train} \geq 20$. Furthermore, for large training sets $n_{train} \geq 60$ the prediction accuracies are virtually identical for all implemented optimality criteria.

For small training sets, the differences between the optimality criteria are larger. However, for such small training sets, each new data point can shift several interaction energies at the same time. This leads to fluctuations like the one we see in the G-optimality curve at $n_{train} = 20$. Consequently, we need to use larger training sets such that the prediction accuracy is nearly converged. With that, we are back in the range of $n_{train} \geq 60$ where all optimality criteria are significantly better than a random selection, but there are hardly any differences between them. (Each of the optimality criteria leads to an RMSE of 4 to 5 meV at $n_{train} = 100$.)

As we have to decide on one of those optimality criteria and all of them perform equally well in the relevant range, we consider other, lesser factors. One factor that does not matter in theory, but can matter in real life, is the computational run time of the algorithm. For all those presented optimality criteria, the run time scales with the size of the training set to be chosen and the size of the training pool. The effort to find a G- or V-optimal training set additionally scales with the size of the evaluation set. Hence, for a broad search these two algorithms might become computationally costly. From the remaining optimality criteria, the E-optimality criterion is the slowest as it involves the maximum function which inhibits some mathematical simplifications. The D-optimality criterion, on the other hand, can be equivalently used in its inverse form (equation 3.5) which saves us one costly matrix inversion. Furthermore, this criterion in principle includes A-optimality as it considers the diagonal- *and* off-diagonal elements of the covariance matrix. All in all, we decide to use the D-optimality criterion for all further evaluations, but it should be noted, that any of the other presented optimality criteria would probably perform equally as well.

4. Results: Predicted Polymorphs

Now that we have a machine learning approach to predict the adsorption energy of discretized polymorphs and we know how to select a good training set, we switch back from the test system to the actual surface system.
The following chapters 4.1 and 4.2 will describe all polymorphs of the first adsorbate layer together with their predicted adsorption energy. These predicted energies enable us to find out which polymorph forms in thermodynamic equilibrium at low coverages.
Having that we turn towards bilayers to resolve the question whether a monolayer of densely-packed, upright-standing molecules or a bilayer polymorph is energetically more favorable at an increased coverage. To that end, we calculate the Gibbs free energy as a function of the chemical potential for all monolayer polymorphs and selected bilayers.
Finally, we compare our findings with an experiment. More precisely, this comparison will be based on the vibrational frequencies and the density of states of the molecular orbitals which we calculate for specific representative polymorphs.

4.1. Training on Full Surface System

From the exhaustive set of $3 \cdot 10^5$ configurations, 300 configurations are chosen using the D-optimality criterion. This step results in a set consisting of 63 *flat*, 136 *mixed* and 101 *standing* configurations. The apparent imbalance between the size of these sets comes from the different number of local adsorption geometries in each of these sets. More local adsorption geometries lead to more different interactions which necessitates larger training sets to achieve a similar prediction accuracy. At the same time, some configurations of each type have to be spared for the test set. Therefore, we want to make sure that the full training set contains at least 100 configurations of each of the aforementioned types. To that end, we choose another 50 *standing* configurations – again using the D-optimality criterion whereby we account for the already chosen configurations.

To ensure that the one-body interactions can be learned properly, for each of the seven local adsorption geometries one additional configuration is build consisting of only one molecule in that local adsorption geometry in a large unit cell having 36 substrate surface atoms. Furthermore, as soon as we have found the energetically most favorable *flat* and *standing* configurations we check their adsorption energy directly via DFT calculations and then added those configurations to the available DFT data. With all that we end up with a set of 359 configurations together with the corresponding dispersion-corrected DFT calculations (settings are listed in Appendix D). The distribution of these configurations is as following: 117 *flat*, 136 *mixed* and 106 *standing*.

From this source set a D-optimal subset of 300 configurations is chosen as training pool and the remaining calculated configurations are used as test set.

In order to monitor the evolution of the prediction accuracy, training sets of increasing size are drawn D-optimally from that training pool. The energy model is fitted independently on each of these training sets and for each fitted model the RMSE of the energy prediction is evaluated on the test test (Figure 4.1). The dashed, black line indicates the thermal energy at room temperature ($k_B T$), which marks the level of accuracy that the predictions should not exceed.

The difference between the blue curve (*mixed*) and the orange curve (*blockwise*) corresponds to the problem of distinguishing the feature vectors of flat and standing pairs of molecules as discussed in chapter 2.3.3. For small training sets of $n \leq 30$ training points the basic learner yields in some cases better prediction accuracies. This behavior is expected, since this learner assumes more correlations and therefore need fewer fitting parameters, leading to a higher learning rate. However, in the most relevant range of $n \geq 50$ training points the modified learner has finally fitted most of its parameter and consistently yields a better prediction accuracy than the other learner which is not able to completely separate the different types of interactions.

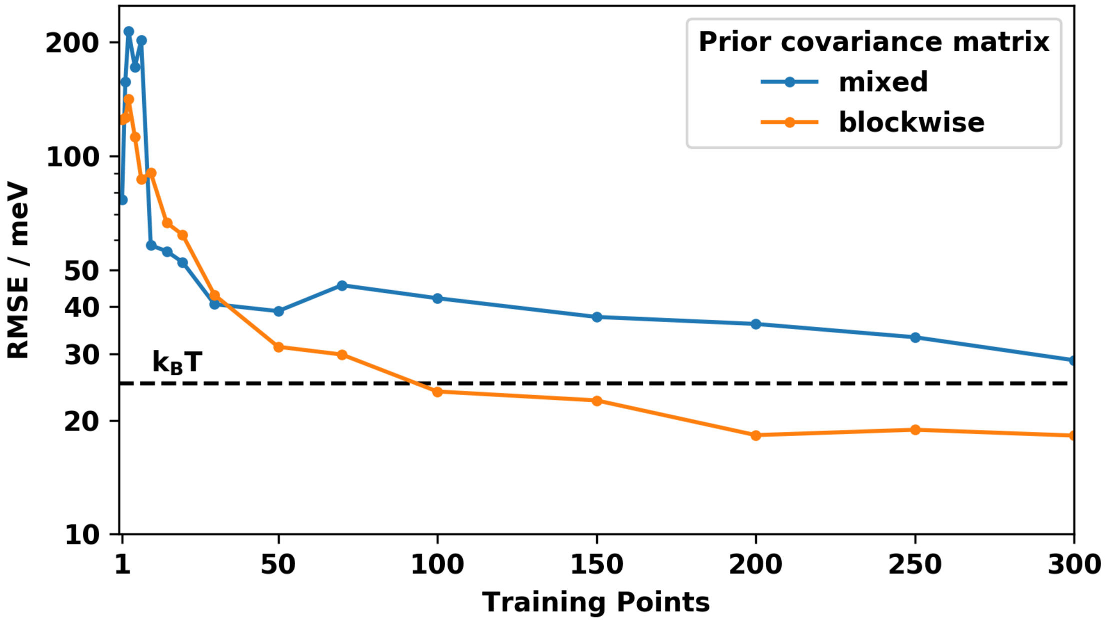

Figure 4.1.: Prediction accuracy on the full surface system.
blockwise indicates the use of the improved prior covariance matrix in which the correlation between different types of interactions is set to zero. Note that the *RMSE* is shown on a logarithmic scale.

4.2. Adsorption Energies of Monolayers

For the actual predictions, Bayesian linear regression is trained on the whole D-optimally chosen training set. For each of the 10 972 902 configurations with up to three molecules per unit cell, the adsorption energy is predicted and then divided by the number of molecules. Figure 4.3 shows those energies plotted as a function of the coverage Θ. Each predicted energy is depicted by a small horizontal dash. The colors indicate the three different types of configurations. Green dashes (*standing*) depict the adsorption energy of those configurations which consist purely of upright-standing molecules, blue dashes (*flat*) those of flat-lying molecules, and orange dashes (*mixed*) correspond to those configurations that contain both types of adsorption geometries. The different width of the dashes has no physical meaning; it only ensures that all markers can be seen.

In the range of low coverage (i.e. sparse configurations), *flat* configurations reach the strongest adsorption energy per molecule. This is as expected since each single, isolated molecule has a stronger adsorption energy when it is adsorbed in a flat-lying geometry than in an upright-standing orientation *and* at the same time pair interaction energies are comparably weak at low coverage.

Going towards higher coverage (e.g. from $\Theta = 0.5\ N_{adsorbate}/nm^2$ to $\Theta = 1.5\ N_{adsorbate}/nm^2$), the predicted adsorption energy of *flat* configurations decreases. Back in the energy model we defined the adsorption energy as one-body interactions between molecules and the surface plus two-body interactions of all pairs of molecules. Furthermore, the one-body interactions are approximately equal for different configurations that consist of the same local adsorption geometries. Consequently, the fact that the predicted adsorption energy decreases indicates that the two-body interaction energies decrease – i.e. pairs of flat-lying molecules with attractive interactions must exist.

Slightly below a coverage of $\Theta = 2\ N_{adsorbate}/nm^2$ the adsorption energy of those flat configurations reaches a minimum and starts to rise from here on. This corresponds to the point where the molecules are forced to lower their pair-distances below the equilibrium distance, which leads to the increase in energy. The minimal intermolecular distance $d_{min} = 2.6$ Å limits this trend as it is not possible to form any configuration with a coverage $\Theta > 2.3\ N_{adsorbate}/nm^2$ using only flat-lying molecules without undercutting this minimal intermolecular distance. The fact that the energetic minimum occurs before this point indicates that the choice of d_{min} is adequate. To create reasonable configurations with coverages higher than $\Theta = 2.3\ N_{adsorbate}/nm^2$ at least one upright-standing molecule must be used in those configurations. In an intermediate range up to $3\ N_{adsorbate}/nm^2$ the *mixed* configurations are the energetically most favorable ones until even those cannot be packed any denser and only configurations consisting purely of upright-standing molecule remain. However, in terms of the adsorption energy *per molecule* the standing molecules never outperform the flat-lying ones.

As we have just seen, we can learn a lot from the energy *per area* of the polymorphs. However, as described in chapter 1.2, at full monolayer coverage and under the assumption that a reservoir of adsorbates exists, the property that is actually minimized is the Gibbs free energy of adsorption (equation 1.26). The Gibbs free energy consists of the chemical potential μ in gas phase and the adsorption energy *per area*. However, μ only plays a role when we compare configurations with (largely) different coverages. Consequently, to find out which configuration is the energetically most favorable one at a specific coverage, we only need the adsorption energy per area which is shown in Figure 4.4. This transformation changes the situation completely. Due to their higher packing density, the *mixed* and *standing* configurations become more relevant and the energetic minimum shifts to a configuration consisting of upright-standing molecules.

As the chemical potential μ is equal to zero at a temperature of $T = 0\ K$, the adsorption energy per area – as depicted in Figure 4.4 – corresponds to the Gibbs free energy of those configurations at $T = 0\ K$.

Using the predicted adsorption energies, we can identify three important configurations that are described in the following as well as listed in 4.2.

Monolayer flat A This is the *flat* configuration with lowest energy **per molecule**. Both molecules in the unit cell of this configuration assume the energetically most favorable local adsorption geometry.

Monolayer flat B This is the *flat* configuration with lowest energy **per area**. It consists of the same local adsorption geometries as *flat A* but has a smaller unit cell consisting of just one adsorbate molecule and it is slightly denser packed. This configuration will be later used for the *flat-on-flat* bilayer.

Monolayer standing The configuration with lowest energy **per area** of all *standing* configurations. It will be later used for the *flat-on-standing* bilayer.

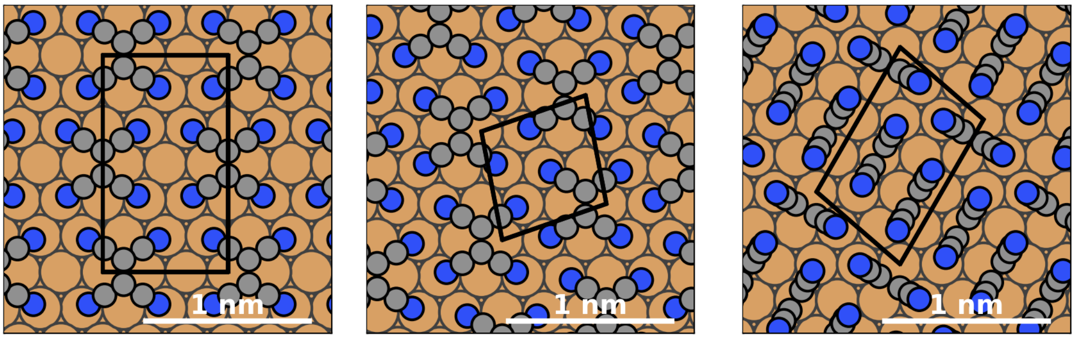

name	**Flat A**	**Flat B**	**Standing**
Identifier	(3,3,-3,3 | 1,1 | 0,5)	(3,-2,1,2 | 1 | 0)	(4,0,-1,3 | 12,17,20 | 0,4,5)
Coverage	1.97 N_{ads}/nm^2	2.22 N_{ads}/nm^2	4.43 N_{ads}/nm^2
$\boldsymbol{E}_{\textbf{pred, per mol}}$	$\mathbf{-2.05}\ \boldsymbol{eV}$	$-1.98\ eV$	$-1.72\ eV$
$\boldsymbol{E}_{\textbf{pred, per area}}$	$-4.04\ eV/nm^2$	$\mathbf{-4.39}\ \boldsymbol{eV/nm^2}$	$\mathbf{-7.62}\ \boldsymbol{eV/nm^2}$

Figure 4.2.: Selected monolayer configurations
$\boldsymbol{E}_{\textbf{pred}}$ ... predicted adsorption energy (*per area* and *per molecule*)
Identifier ... internal code to describe the unit cell, the local adsorption geometry, and the position of those (corresponding to the three parts of the identifier which are separated by vertical bars).

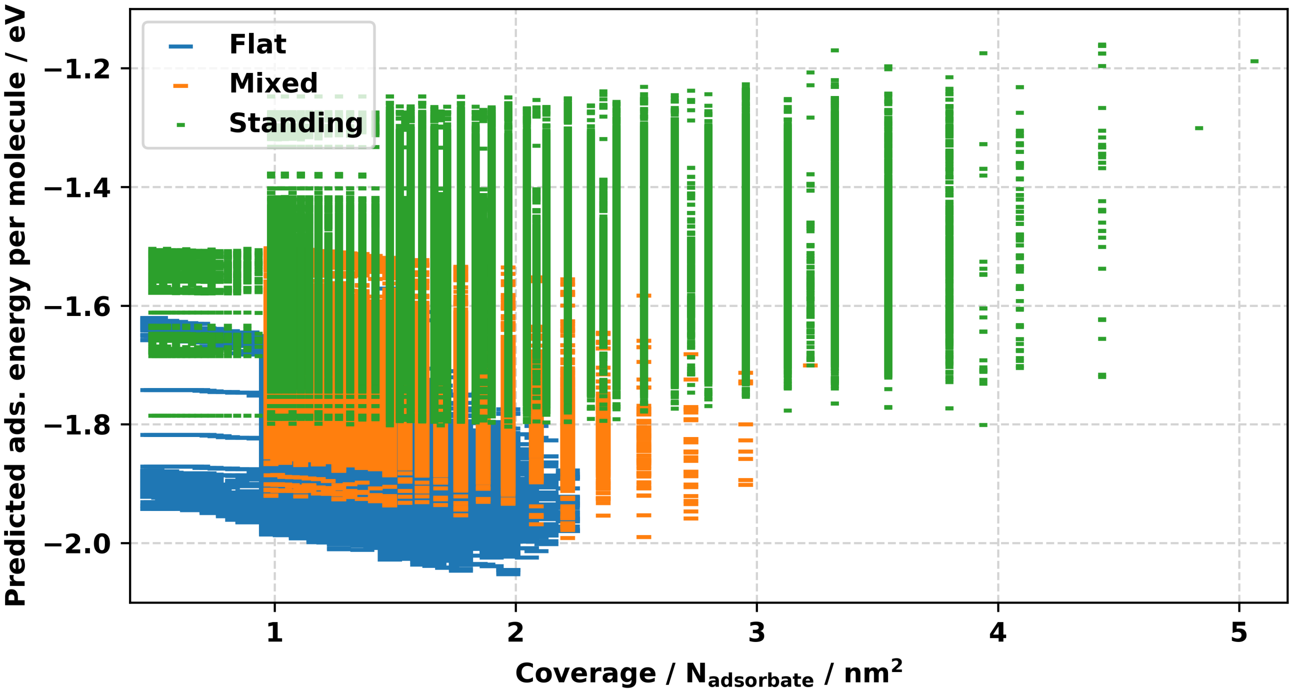

Figure 4.3.: Predicted adsorption energy *per molecule* of all monolayer-configurations of TCNE on Cu(111). The apparent step at $\Theta \approx 1\ N_{adsorbate}/nm^2$ is an effect of the finite size of the unit cells and does not affect energetically favorable configurations.

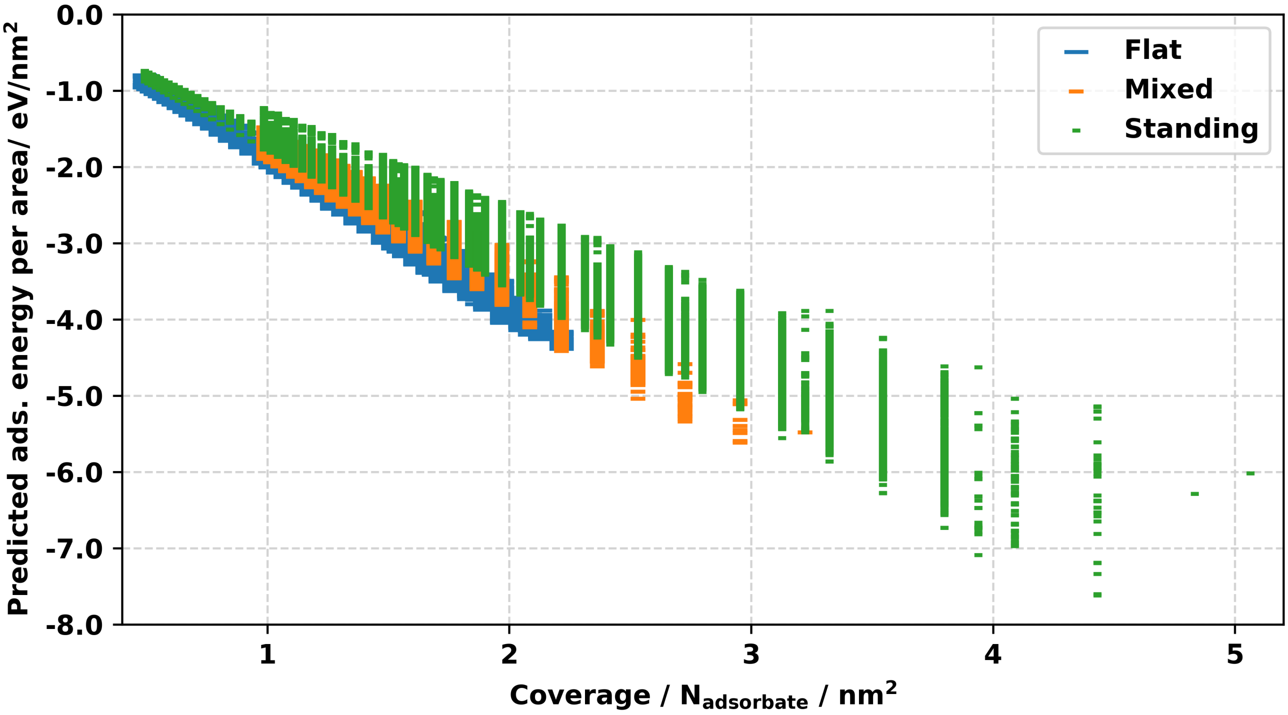

Figure 4.4.: Predicted adsorption energy *per area* of all monolayer-configurations.

4.3. Building Bilayer Configurations

A central question of this study is: How does the structure of the adsorbate layers change with increasing coverage. A major part of this question, namely how the structure of the *first* adsorbate layer changes as a function of the coverage, can be solved directly from the predicted adsorption energies that are shown in the previous section. However, to answer the question if – and at which point – the transition to bilayer polymorphs happens we obviously have to create such bilayer polymorphs and then compare them with densely-packed monolayer polymorphs. This step will be described in the following section.

In general, the interaction with the substrate is much stronger in the first layer than in the second layer and so molecules 'prefer' to adsorb in the first layer. However, when the packing density is increased the intermolecular distances have to shrink and with that the Pauli repulsion between molecules starts to increase rapidly. Consequently, at some point it becomes energetically more favorable to form a second adsorbate layer instead of a densely packed monolayer. In fact, the phase transition in the first layer (from flat-lying to upright-standing molecules) might never be reached because the molecules already moved to the second layer before the transition point is reached.

To resolve this question, we have to compare the Gibbs free energy of densely-packed monolayers with that of bilayers that have the same number of molecules per substrate surface area. But of course, this comparison only makes sense if we compare the monolayers to the energetically most favorable bilayer polymorphs. So, we have to find them. Unfortunately, there are several reasons why we cannot apply the SAMPLE approach straightforward to bilayer polymorphs.

The main difference between this task and that of finding the best polymorphs of the first layer is that for the first layer we had a well-defined supporting layer: the surface of the substrate which is kept fixed during the structure search. For the second layer the supporting layer is given by the first layer for which lots of different polymorphs exist and which is not as rigid as the substrate. Therefore, we face two major challenges:

1. For the first layer, it can be justified to neglect the relaxation of the substrate atoms below it, since these atoms are limited in their movement and a relaxation of the substrate surface leads to similar energy offsets for all local adsorption geometries. Contrary to that, the first adsorbate layer might show major reconstructions when a second layer is built on top of it and these reconstructions will be different for different polymorphs of the second layer. Consequently, the first and the second adsorbate layer cannot be evaluated independently but we have to evaluate the reconstructions of all combinations. This means that the number of possible polymorphs is approximately squared!

2. For the first layer, the periodic unit is given by the substrate lattice, meaning that the surface of each substrate unit cell is equivalent. For the second adsorbate layer, however, the periodic unit is defined by the unit cell of the first layer which is typically at least an order of magnitude bigger than the substrate unit cell. Therefore, the potential energy surface of the unit cell is much bigger and hence finding local adsorption geometries of the second layer needs much more computational time. On top of that, each first layer constitutes a different surface, so we would need to find a new set of second-layer local adsorption geometries for each first layer.

Considering those obstacles two approaches seem feasible:

- **Second-layer local adsorption geometries**
  Similar to the first layer, one could employ methods like Gaussian process regression to find local minima of the potential energy surface for a single molecule in the second layer. Once these building blocks are found, they could be added to the energy model and in that way one could overcome the configuration explosion concerning combinations of the first and second layer. This approach seems most promising for rather stable first layers, like phases consisting of flat-lying molecules. It is, however, unclear how building blocks for the second layer can be found such that they can be used with (most of) all first layer building blocks. Furthermore, this approach does not account for reconstructions of the first layer.
- **Build second layer on top of promising first layers and conduct geometry optimization**
  In this approach, one would pick a set of most relevant first layer polymorphs – e.g. energetically favorable, stable, close-packed configurations. An additional adsorbate molecule is then placed on top of this first layer in a reasonable (guessed) second layer adsorption geometry. From that starting point, a local geometry optimization of both adsorbate layers has to be done. On the one hand, this approach is less systematic due to the necessary choice of the guess geometries. But on the other hand, it allows to check for reconstructions of the first layer.

As we need to check if reconstructions of the first layer occur in the studied system, we have to choose the second approach.
For the first layer, we pick two very complementary configurations: The first configuration is that one which has the lowest adsorption energy per area of all *flat* configurations. We choose this configuration as it represents the most close-packed monolayer that can be built using only flat-lying local adsorption geometries and without going below the intermolecular distance threshold d_{min}. This configuration is depicted in Figure 4.2 (*flat B*) and it corresponds to the lowest orange dash in Figure 4.4. It consists of one flat-lying molecule in a unit cell that contains 8 surface atoms. Therefore, it has a coverage of $\Theta_{1a} = \frac{1}{8}\ N_{ads}/N_{substrate} = 2.22\ N_{ads}/nm^2$. Two guess-bilayers are built from this first layer: One is denoted as *flat-on-flat/vertical* as the second-layer molecule is put exactly above the molecule of the first layer. The other is named *flat-on-flat/shifted* as the second-layer molecule is additionally shifted horizontally such that the center of its central C=C-bond is situated above a nitrogen of the TCNE that is adsorbed in the first layer. Both of these *flat-on-flat* bilayer geometries consist of two molecules per adsorbate unit cell and therefore have a coverage of $\Theta_{1b} = \frac{2}{8}\ N_{ads}/N_{substrate} = 4.43\ N_{ads}/nm^2$. We will later compare this bilayer with a densely packed monolayer of the same total coverage[a] to find out if TCNE rather forms such a densely packed monolayer or if it transforms directly to a bilayer.

To study the behavior at even higher coverages, we additionally start from a densely packed monolayer-configuration. Namely, we pick that configuration which reaches the lowest energy *per area* of all configurations – the lowest green dash (configuration *standing*) in Figure 4.4. It consists of three standing molecules in a unit cell which has twelve surface copper atoms: $\Theta_{2a} = \frac{3}{12}\ N_{ads}/N_{substrate} = 4.43\ N_{ads}/nm^2$. To build a bilayer from that, we place an additional TCNE molecule right above the two molecules that stand parallel to each other (see right

---

[a] The term *total coverage* here denotes the *total* number of molecules of *all* adsorbate layers divided by the substrate surface area.

hand side of Figure 4.5). The resulting bilayer – which will be named *flat-on-standing* – has a coverage of $\Theta_{2b} = \frac{4}{12}\ N_{ads}/N_{substrate} = 5.91\ N_{ads}/nm^2$.

As these bilayers are only guessed geometries, we have to conduct local geometry optimizations with these configurations as starting points. To that end, we first optimize the first adsorbate layer (before we add the second layer to the geometry). In that step we also include the two uppermost layers of the substrate into the optimization procedure. I.e. the lower lying layers of the substrate are constrained, but the two uppermost layers of the substrate can adapt to the adsorbate layer above it. We then constrain the complete substrate, add the second layer as described above and optimize the geometry of both adsorbate layers. To improve the convergence speed for optimization of the *flat-on-standing* bilayer, we conduct additional optimization runs of the second layer during which we keep the first layer fixed. Furthermore, after twelve optimization steps it becomes obvious that the adsorption energy of the *flat-on-flat/vertical* configurations converges to a level far above the *flat-on-flat/shifted* configuration. Therefore, the *flat-on-flat/vertical* configuration is dropped and only the *flat-on-flat/shifted* configuration is used in the further evaluations using the shortened name *flat-on-flat*.

A detail of the nomenclature should be clarified at this point: Up to now, we always dealt with *configurations* – i.e. *discretized* polymorphs that consist of molecules in their undisturbed local adsorption geometry. But as soon as we optimize the geometry of these *configurations*, the constituting molecules are not anymore exactly in these local adsorption geometries. Consequently, we will not call these optimized geometries *configurations*, but switch back to the more general term *polymorph*.

We end up with two different, optimized bilayers polymorphs: (see Figure 4.5)

Bilayer: flat-on-flat The first layer of this polymorph is similar to the monolayer configuration *flat B*, on which it is based (see Figure 4.2). During the geometry optimization, the molecule in the second layer tilts from its flat orientation to an orientation that has an angle of $\Theta \approx 25°$ with respect to the surface plane. Furthermore, this molecule loses its bending and assume a planar geometry.

Bilayer: flat-on-standing The first layer of this polymorph is based on the monolayer configuration *standing*. The molecule that is adsorbed in the second layer is less tilted than that in the *flat-on-flat* polymorph ($\Theta \approx 18°$). This difference probably stems from the weaker packing density of this second layer compared to the second layer of the *flat-on-flat* polymorph.

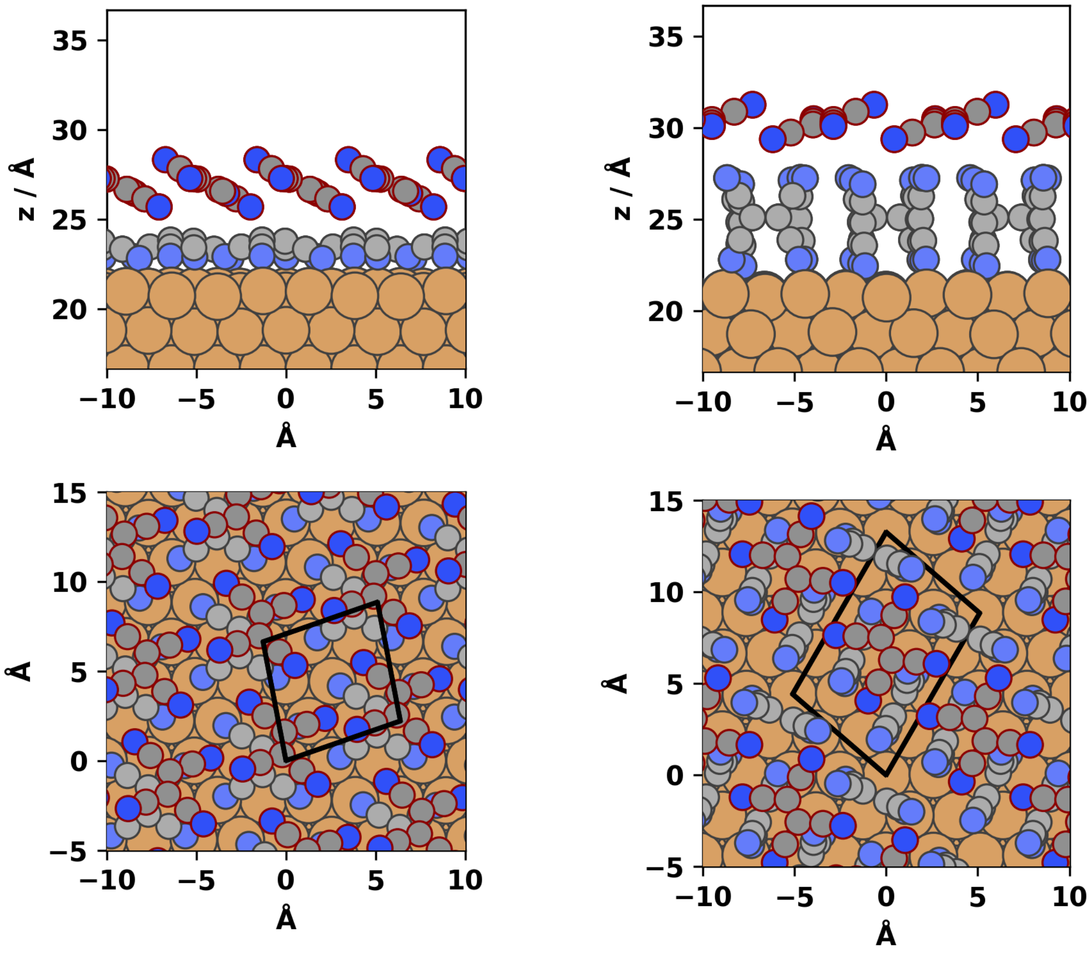

(a) **Bilayer: flat-on-flat** (b) **Bilayer: flat-on-standing**

Figure 4.5.: Front- and top view of the two bilayer polymorphs that are used for comparison with densely-packed monolayer polymorphs. The red edges mark molecules of the second layer. The unit cell of the adsorbate layer is indicated in black.

4.4. Comparing Monolayers and Bilayers – Surface Phase Diagram

Now that we have an exhaustive set of monolayers and selected bilayers, we can calculate Gibbs free energy for each of them as derived in chapter 1.2.1 (eq. 1.26). However, we must be careful when comparing the monolayer configurations with the bilayer polymorphs that we just built. The bilayers polymorphs are locally optimized. In contrast to that, the monolayer configurations are built from locally optimized adsorption geometries but have not been geometry optimized *after* these local adsorption geometries have been combined to configurations.

To resolve this discrepancy, we first create a separate, preliminary plot of Gibbs free energy versus chemical potential only for the monolayer configurations. From that data, we find all those configurations, which – at any chemical potential – yield the lowest Gibbs free energy of all configurations. For each configuration of that subset we then conduct a local geometry optimization that includes the adsorbates as well as the two uppermost substrate layers. These optimized monolayers can now be safely compared to the optimized bilayers.

Figure 4.6 shows the Gibbs free energy of the optimized mono- and bilayers as a function of the chemical potential μ, which itself is a function of pressure and temperature. $\mu = 0\ eV/molecule$ corresponds to either infinitely high partial pressure of TCNE in the gas phase and/or a temperature of $0\ K$. On the other hand, a more negative chemical potential corresponds to lower partial gas pressure and/or higher temperature.

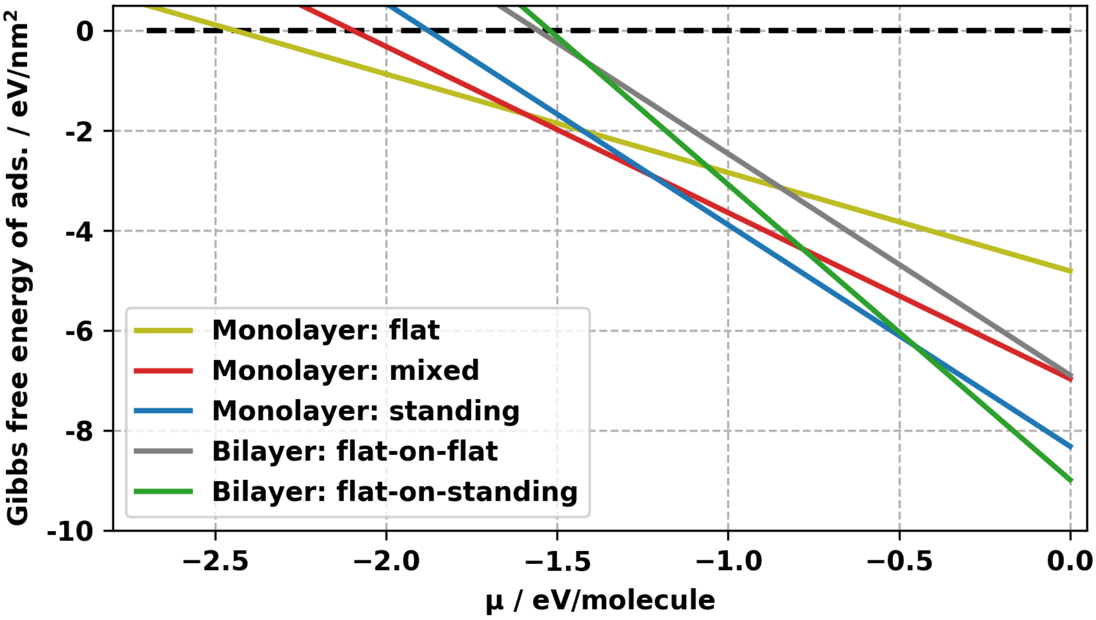

Figure 4.6.: Gibbs free energy of selected configurations
At $\mu = 0\ eV/molecule$, the Gibbs free energy is identical to the energy per area of that configuration. The derivative with respect to μ corresponds to the coverage.

Let us start at the very left-hand side of the plot – i.e at low partial gas pressure and/or high temperature. Below $\mu = -2.5\ eV/molecule$ the chemical potential of the molecules in gas phase is so strong that no stable adsorbate layer forms on the surface – this threshold is indicated by the dashed horizontal line. When we then continue towards higher pressure and/or lower temperature, the first phase that reaches a negative Gibbs free energy is one that consists purely of flat-lying molecules (yellow, *Monolayer: flat*). At a chemical potential $\mu \approx -1.5 eV/molecule$ and above, a mixed configuration (red, *Monolayer: mixed*) leads to a lower Gibbs free energy than the *flat* configuration. Shortly after that point, a configuration consisting purely of densely-packed upright-standing molecules becomes the energetically most favorable one (blue, *Monolayer: standing*).

Turning to the two bilayer polymorphs we see that the one with the lower coverage (gray, *Bilayer: flat-on-flat*) at no point becomes the energetically most favorably polymorph. In fact, the energetic difference to the most favorable polymorph is at all times larger than $1\ eV/molecule$! Furthermore, this gray curve is exactly parallel to the blue one (*Monolayer: standing*). This is no coincidence but due to the fact that we created this bilayer such that is has the same total coverage as the densely-packed monolayer and in the chosen model the coverage Θ corresponds to the derivative of Gibbs free energy with respect to μ. In contrast to that, the bilayer of higher coverage (green, *Bilayer: flat-on-standing*) finally crosses the Gibbs free energy of the densely-packed monolayer at weak low temperature and/or high pressure.

From these findings we can conclude that indeed a phase transition from flat-lying to upright-standing molecules occurs at around $\mu \geq -1.5 eV/molecule$. Furthermore, we can eliminate the possibility that this transition is masked by a bilayer polymorph as the studied bilayer polymorphs are far higher in Gibbs free energy at this point. Only at a lower temperature and/or higher pressure a transition towards bilayer polymorphs is predicted. The difference between both evaluated bilayers indicates that possible bilayers (and multilayers) would form on top of a first layer consisting of upright-standing TCNE molecules.

As the chemical potential is a function of partial gas pressure and temperature (eq. 1.27) we can transform the previous plot into a more descriptive two-dimensional surface phase diagram (Figure 4.7). Low chemical potential here corresponds to high temperature and/or low partial pressure – i.e. the lower right corner. Going from here towards high pressure and low temperature we can read the same progression as in the previous plot.

It should be noted here that this phase diagram is mostly of qualitative nature. It is a useful tool to visualize the succession of predicted polymorphs. However, due to glancing intersection in the Gibbs free energy curves, small prediction errors can shift the boundaries between surface phases by large amounts. Furthermore, one should keep in mind that the bilayer *flat-on-standing* represents the edge of the studied search area. Consequently, at low temperatures (i.e. on the very left hand side of the phase diagram) other polymorphs which are currently not part of the search space – for instance such with more than two layers – could potentially replace the *flat-on-standing* bilayer. However, this limitation does not affect the outcome of this study, as we do not rely on this specific range of the phase diagram.

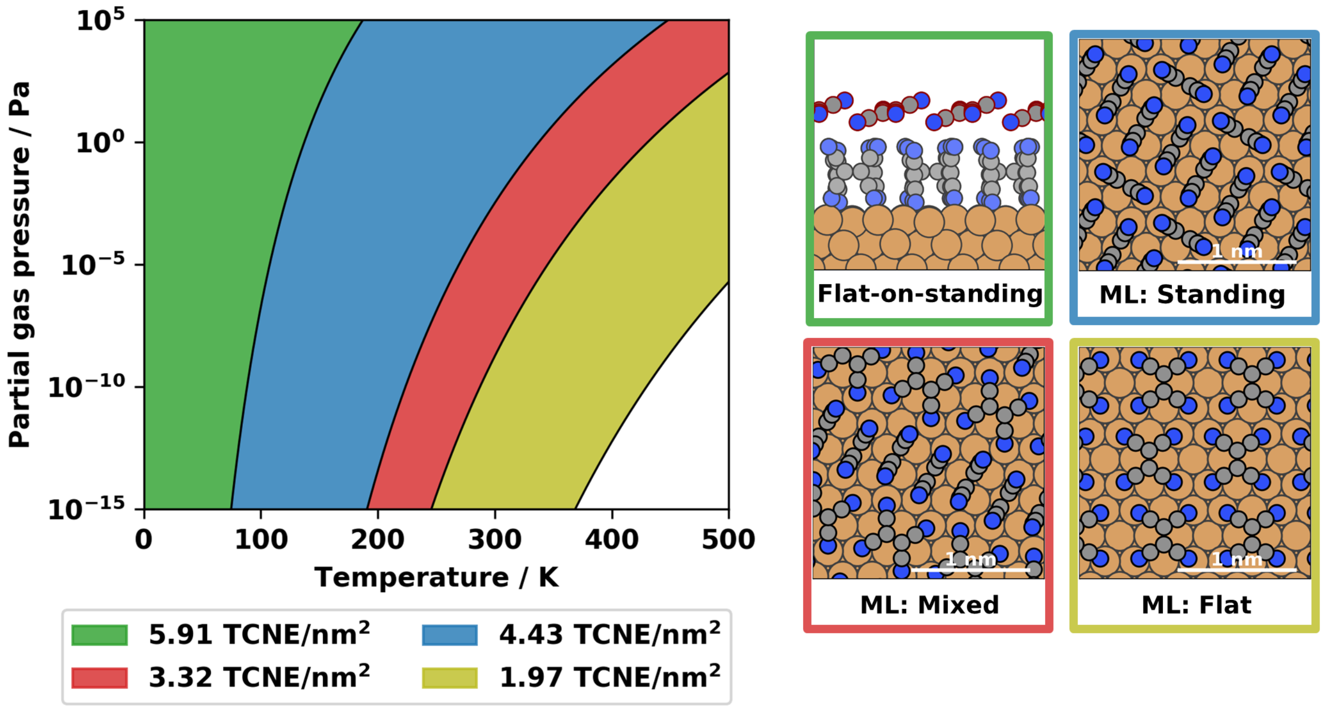

Figure 4.7.: Surface phase diagram.
Each combination of pressure p and temperature T leads to a chemical potential $\mu(p, T)$. The color at each point (p, T) corresponds to the configuration, which has the lowest Gibbs free energy for $\mu(p, T)$.

5. Comparison to Experiment

To check the predicted phases, we will now compare them with an experiment. Unfortunately, for the studied combination of substrate and adsorbate material, not a lot of experiments exist. However, we found one by Wulf Erley and Harald Ibach [13] in which they measured the formation of the adsorbate layers in an indirect way.
In the following we will quickly summarize the experimental setup they used. Then we discuss the involved techniques and which simulations we can use to compare our polymorphs with their results. Finally, we conduct those simulations and discuss the findings.

5.1. Summary of the Experiment

In the experiment [13], TCNE is deposited from the gas phase onto a cleaned Cu(111) surface in ultrahigh vacuum (UHV) at a temperature of $T = 100\ K$. At several points during the deposition process, the coverage is measured using auger electron spectroscopy (AES). The coverage at each of these deposition steps is estimated by comparing the ratio of the C_{272} and the Cu_{920} peak of the AES measurements with reference values obtained from a monolayer of CO on Cu(111). Additionally, electron energy loss spectra (EELS) were measured with an electron energy of $5.0\ eV$ and a resolution of $\Delta\tilde{\nu} = 30\ cm^{-1}$ at each of those steps. A more detailed explanation of the experimental setup is given in the original publication [46].

The central method of this experiment is electron energy loss spectroscopy (EELS). In this technique a beam of well-defined kinetic energy is directed onto the material. When inelastic scattering of the electron beam occurs, energy is transferred from that beam to the specimen. This energy loss of the beam can then be measured using an energy dispersive electron detector.

An important inelastic scattering process are excitations of vibrational modes of the adsorbate molecules. Hence, the positions of the peaks in the EELS curve indicate which vibrational frequencies and corresponding vibrational modes can be stimulated in the specimen. At the same time, the vibrational frequencies of a molecule depend on its geometry, its local environment and its electronic structure. Furthermore, when additional charge is accumulated on the molecule its electronic structure changes. With all that the vibrational spectrum becomes a function of the charge distribution of the molecule.

In the interpretation of the experiment the inverse of this relationship is used: At each deposition step, the observed frequencies are compared to the literature values for neutral, singly-, doubly- and triply-charged TCNE. This way, the experimenters estimate the electronic gross charge of all adsorbed TCNE molecules. As the gross charge changes with the position and orientation of an adsorbate molecule, the experimenters are able to see at which points of the deposition process new 'types' of adsorbate molecules appear.

Back to the experiment: At low coverage Erley and Ibach see, among other EELS-peaks, one at a frequency of $\tilde{\nu}_{cc} \approx 1275 cm^{-1}$ (see Table 5.1). This vibration is interpreted as symmetric stretching mode of the central C=C double bond of a doubly to triply charged TCNE molecule (see Figure 5.1). They conclude that this multiply ionized species corresponds to flat-lying TCNE molecules in the first adsorption layer.

Upon increasing the coverage, they observe a new peak at $\tilde{\nu}_{cc} \approx 1375\ cm^{-1}$. This vibration belongs to the same C=C stretching mode, but the frequency shift indicates that the adsorbate molecules that lead to this vibration are roughly singly charged. From that, they conclude that at this point – at an Auger value of $C_{272}/Cu_{920} \approx 0.11$ – a full monolayer coverage must have been reached and now molecules start to form a second layer on top of the first one.

At about twice that coverage – $C_{272}/Cu_{920} \approx 0.24$ – another new peak with a frequency of approximately $\tilde{\nu}_{cc} \approx 1565\ cm^{-1}$ becomes visible. This peak corresponds again to the same vibrational mode but this time under the assumption that the corresponding TCNE molecule is neutral. This leads to the conclusion that now the second layer is filled up, too, and so adsorbate molecules are forced into higher layers in which they interact only weakly with the substrate and hence maintain their neutral charge state.

molecule	$\tilde{\nu}_{cc}$ / cm^{-1}
low coverage / first layer	1275
medium coverage / second layer	1375
high coverage / higher layers	1565

Table 5.1.: Experimental vibration frequencies of the C=C stretching mode measured via EELS in different adsorbate geometries.

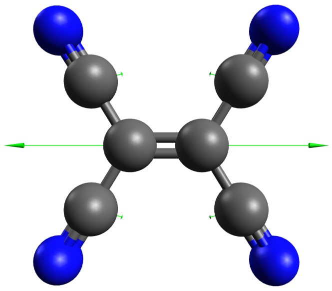

Figure 5.1.: Symmetric stretching mode of the central C=C double bond of TCNE

In a follow-up paper, the main findings of the experiment were summarized:

> ‘Evidence has been found that TCNE adsorbs as multiply charged species (TCNE^{x-}, x = 2 to 3) in the first layer, as singly charged anions (TCNE^{-}) in the second layer, and as neutral molecules (TCNE^{0}) in all subsequent layers.’ (Wulf Erley [47])

5.2. Vibrational Calculations

Now that we know the experiment, we need a way to compare its results with the properties of the polymorphs that we predict. One property that can be used for this comparison are the vibrational frequencies. To calculate these, we must first make sure that the geometry of the molecules in consideration represent a local minimum of the potential energy surface. This is achieved by a local geometry optimization. If a geometry is not in such a local minimum, but in a saddle point or local maxima, a distortion of the geometry would lead to a lower energy and hence to an imaginary vibrational frequency in that direction.

To actually calculate the vibrational modes the harmonic approximation is used. Within this approximation the second derivatives of the energy of the molecule, or equivalently the first derivatives of the forces, correspond to the molecule's force constant matrix (Hessian matrix).

$$H_{i,j} = \frac{\partial^2 E}{\partial x_i \, \partial x_j} = \frac{\partial F_{x_j}}{\partial x_i} \tag{5.1}$$

FHI-*aims* supports analytical calculation of forces, but currently not the analytical calculation of the second derivatives (i.e. the Hessian matrix). Therefore, forces are calculated directly in FHI-*aims* and the second derivatives are subsequently calculated using finite differences of the forces.

To that end, we set up a series of DFT calculations. In each calculation one atom of the molecule in consideration is displaced from its equilibrium position while all other atoms are kept in their original position. For the displacement a small, constant shift of $\delta = \pm 0.0025$ Å in x-, y- *or* z-direction is used.

For all these distorted geometries the forces on all atoms are calculated, divided by the displacement δ and then added to the Hessian matrix. The eigenmodes of that Hessian then provide the vibrational modes of the atoms and the eigenvalues yield their vibrational frequencies.

5.2.1. Calculated Vibrational Frequencies

For the vibrational analysis we choose the following representative adsorption geometries:

First layer, flat-lying This is the adsorption geometry of the TCNE in the configuration *Monolayer Flat B* which occurs at low coverage.

First layer, standing Here we use the local adsorption geometry of a single upright-standing TCNE molecule with the central C=C bond parallel to the surface. (see geometry *'bridge-top'* in Figure 2.1). This molecule is evaluated in a unit cell of four times four surface atoms.

Second layer, flat-on-flat This geometry corresponds to the molecule in the second layer of bilayer *Flat-on-flat* (see Figure 4.5 (a)).

Second layer, flat-on-standing This is the adsorbate molecule that forms the second layer of bilayer *Flat-on-standing* (see Figure 4.5 (b)).

Gas phase As a reference, we additionally calculate the vibrational modes of neutral, singly, and doubly charged TCNE in their relaxed geometry in the gas phase.

All vibration calculations are done in FHI-*aims* [19] using the PBE functional [32] and the vdW^{surf} correction [38]. To compare the experiment with our predictions, we calculate the vibrational frequencies of the predicted configurations. We then visually inspected the calculated vibrations in Avogadro [48] and selected that vibrational mode that corresponds to a symmetric stretching of the bond between the two central carbon atoms. Table 5.2 shows the frequencies of these modes.

Adsorbate geometry	FHI-*aims*, PBE $\tilde{\nu}_{cc}$ / cm^{-1}
First layer, flat-lying	1263
First layer, upright-standing	1341
Second layer, flat-on-flat	1482
Second layer, flat-on-standing	1512
Gas phase $q = -2\ q_e$	1273
Gas phase $q = -1\ q_e$	1375
Gas phase $q = 0\ q_e$	1491

Table 5.2.: Calculated vibration frequencies of the C=C stretching mode for the adsorbate molecules as well as in the gas phase for different charge states.

A first glance at the calculated values shows that the frequency of this vibrational mode increases when we go from the flat-lying, via an upright-standing to the second-layer adsorption geometries (the first four calculations). Furthermore, the difference between the two second-layer geometries (*flat-on-flat* and *flat-on-standing*) $\Delta\tilde{\nu}_{cc} = 30\ cm^{-1}$ is significantly smaller then the two steps before.

Figure 5.2 combines the calculated vibration frequencies with the experimental data. The frequencies of the flat-lying geometry (blue, dashed) fit nicely with the experimental data of the first-layer geometry (blue, continuous). Both frequencies are very close to the vibrational frequency of a doubly charged molecule in vacuum (black dashed line, 1273 cm^{-1}). This finding supports the experimental interpretation that at low coverage TCNE adsorbs as an approximately doubly charged species in a flat-lying orientation.

The next peak of the EELS spectrum is situated at 1375 cm^{-1}, which is exactly the vibrational frequency that we calculate for the singly charged molecule in the gas phase. Hence, we confirm that this EELS-peak corresponds to a singly charged TCNE molecule.

However, there are also discrepancies: First of all, the calculated frequency of the *standing* molecule (orange, dashed) is lower than the expected, experimental value (orange, continuous) as well as the gas phase value (black, dashed). However, the difference ($\Delta\tilde{\nu}_{cc} = 34\ cm^{-1}$) lies still within the range of error for this type of calculation.

More importantly, the experiment suggests that the singly charged species of TCNE corresponds to adsorbates in the second adsorbate layer. But when we have a look at the vibrational frequencies of the second-layer adsorbates (both green dashed lines) we see that they are far from the frequency of the singly-charged species! Both, the second layer that is situated on top of a flat layer (*flat-on-flat*, dark green, dashed) as well as the second layer on top of a standing layer (*flat-on-standing*, light green, dashed) lead to vibrational frequencies very close to that of a neutral molecule in the gas phase ($\Delta\tilde{\nu}_{cc} = 9\ cm^{-1}$ and $\Delta\tilde{\nu}_{cc} = 21\ cm^{-1}$, respectively). Consequently, a TCNE molecule that is adsorbed in the second layer appears to be nearly neutral regardless if the first layer consists of flat-lying or upright-standing TCNE molecules.

From the vibrational data we can conclude that in contrast to the interpretation of the experiment, the singly charged species cannot be explained by TCNE molecules adsorbed in the second layer. Instead, the experimentally observed singly charged species fits the vibrational properties of upright-standing TCNE molecules in the first adsorbate layer which are predicted using the SAMPLE approach.

At this point, the attentive reader certainly observed the discrepancy between the calculated second-layer geometries (green, dashed) and the experimental value that is attributed to the neutral species (green, continuous). Although this difference does not alter the results of the evaluation additional calculations concerning this issue are conducted and can be found Appendix C).

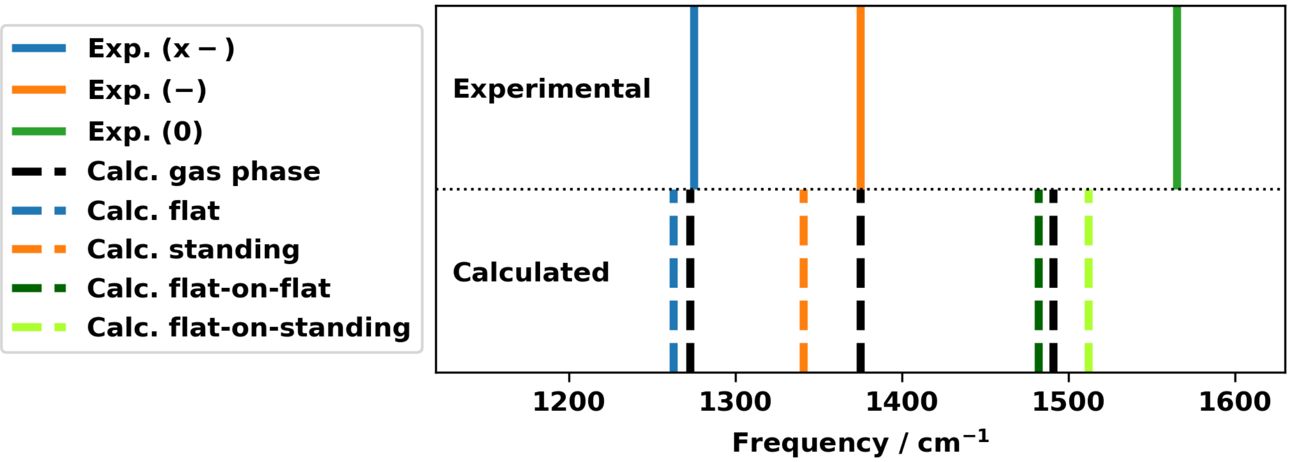

Figure 5.2.: Vibration frequencies of the central C=C double bond.
Experimental data is listed in Table 5.1
Calculated data (including gas phase) is listed in Table 5.2

5.3. MODOS

Another, alternative way to analyze the predicted polymorphs is to evaluate the molecular orbitals. From their occupation, we are not only able to calculate the gross charge of the adsorbed molecule, but we can also understand why molecular bonds are strengthened or weakened upon adsorption.

When a molecule is adsorbed onto a surface, it is possible that electrons are transferred from the surface to the adsorbate (or vice versa). The additional electrons will then occupy molecular orbitals of the adsorbate molecule. Two important molecular orbitals of neutral TCNE in the gas phase are depicted in Figure 5.3: The highest occupied molecular orbital (HOMO) contributes to the double-bond of the two central carbon atoms while the lowest unoccupied molecular orbital (LUMO) has a node between those two carbon atoms and therefore is anti-bonding with respect to that bond. Consequently, when the LUMO gets filled, the bond between the central carbon atoms gets weaker and the corresponding vibrational frequency is decreased as we have seen in the vibrational calculations.

To quantify the change of the occupation of the LUMO, we project the density of states (DOS) of the molecule in its adsorbate state onto the molecular orbitals of the gas phase state – i.e. the state that the molecule would assume in the gas phase. This projection is called molecular orbital density of states (MODOS) [49] and we conduct it using a Fortran script written by

Yong Xu (FHI, Berlin, Germany) based on the PhD thesis of Lorenz Romaner [50]. Having done that, we integrate the DOS of each of molecular orbital from minus infinity up to the Fermi energy E_F to obtain the occupation of that molecular orbital (see Figure 5.4 and 5.5).


(a) HOMO **(b)** LUMO

Figure 5.3.: Highest occupied and lowest unoccupied molecular orbital of neutral TCNE in the gas phase

5.3.1. Occupation of the LUMO

For this evaluation we took the same adsorption geometries as in the vibrational calculations with one change: To simplify the plots only the *flat-on-standing* geometry is used as example of a bilayer but not the *flat-on-flat* bilayer which shows similar results.

Figure 5.4 displays the molecular orbital density of states as a function of the energy. The shading indicates the occupied states below the Fermi energy E_F. The uppermost curve corresponds to a TCNE molecule that is adsorbed in the second layer on top of a first layer of upright-standing molecules. The two obvious peaks in the DOS correspond to the HOMO and LUMO of TCNE in the gas phase. A quick glance at this shaded area already reveals that the HOMO of TCNE stays filled upon adsorption, and the LUMO stays completely empty – i.e. in contrast to the experimental findings, we observe *no* substantial charge transfer to the second layer. To quantify this observation, the occupation of each molecular orbital is calculated and displayed in Figure 5.5. In this plot the HOMO and LUMO correspond to molecular orbitals 32 and 33, respectively. The occupation of each molecular orbital is given by the integral over the density of states of that orbital up to the Fermi energy E_F. Indeed, the occupation of the LUMO (molecular orbital 33) of the second layer adsorption geometry is zero. [a]

Proceeding with the *standing* adsorption geometry, we see that its LUMO is approximately half filled (0.8 of 2 possible electrons or equivalently 41 %). In contrast to that, the LUMO of the *flat* adsorption geometry is nearly filled (1.6 electrons, 79 %).

An interesting effect can be seen in the molecular orbitals 20 to 30. These orbitals are fully occupied if the molecule is in the gas phase as well as when the molecule is adsorbed in the *second* layer. When the molecule is adsorbed in the *first* layer, however, the molecule hybridizes with the substrate and electrons move from the metal substrate to the LUMO of the adsorbate.

---

[a] The *flat-on-flat* bilayer show a very minor filling of the LUMO: 0.14 of 2 possible electrons or equivalently 7 %

As a counter-effect electrons of lower lying molecular orbitals are pushed back into the substrate – an effect known as back-donation.

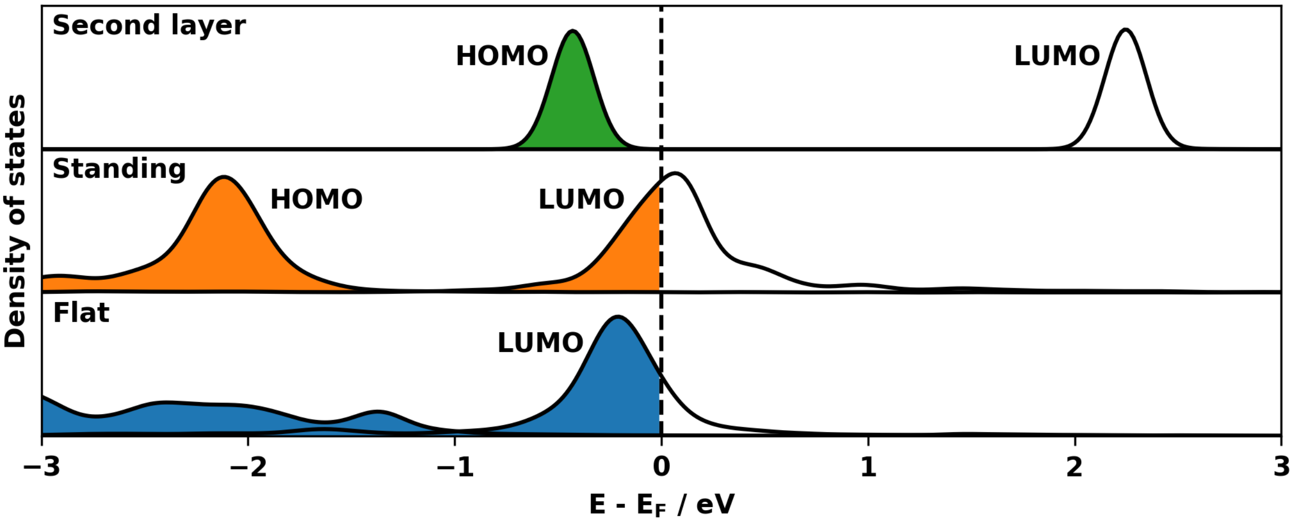

Figure 5.4.: MODOS of a TCNE molecule in different adsorption geometries. (a) flat-lying geometry, (b) standing geometry, and (c) adsorbate geometry in the second layer (on top of a first layer of upright-standing molecules).

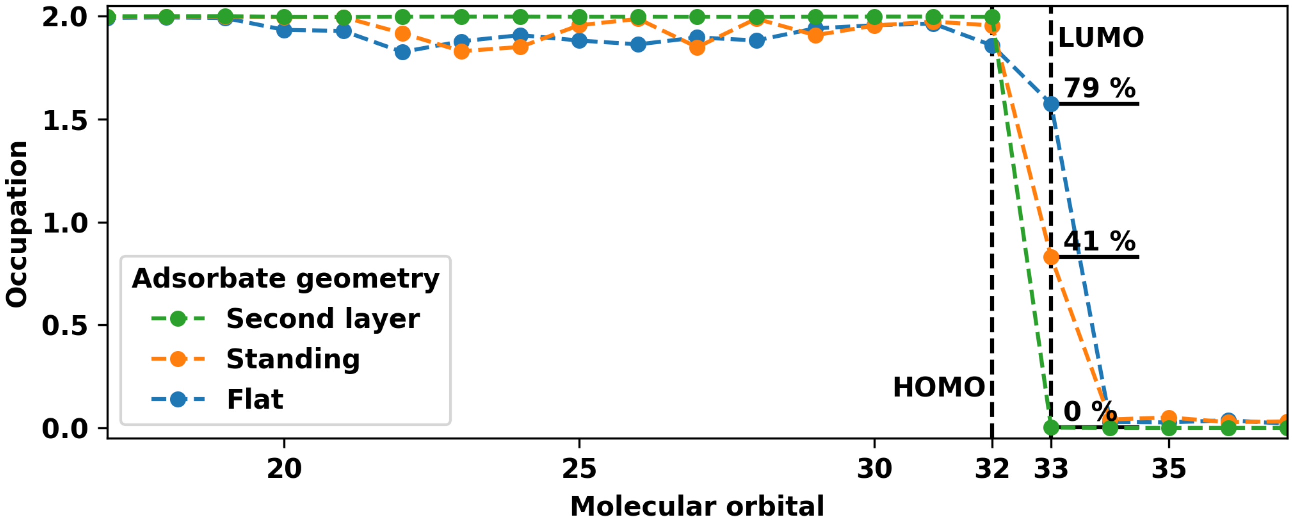

Figure 5.5.: Occupation of molecular orbitals for different adsorption geometries. The percentage is given relative to the maximum occupation of two electrons per molecular orbital.

Conclusion and Outlook

In this work, the SAMPLE approach was applied to predict surface polymorphs of TCNE on Cu(111). The main question to be answered was: Which surface polymorphs form when the coverage is increased above the full-monolayer coverage of flat-lying TCNE molecules? Two scenarios arise: Either a densely-packed monolayer consisting of upright-standing adsorbates forms or a second layer forms on top of the first one. To answer the question, we first had to find all monolayer polymorphs including densely-packed ones.

As densely-packed polymorphs show particularly high pairwise interaction energies, their prediction necessitates sophisticated selection of the training data that constitutes the input for the chosen machine learning approach. To that end, we implemented training set selection strategies based on several different optimality criteria. Subsequently, we benchmarked these strategies on a test system.

With that, we then employed the SAMPLE approach to predict the adsorption energies of an exhaustive set of discretized monolayer polymorphs (named *configurations*). With those predicted energies and the help of ab initio thermodynamics, we are able to find those polymorphs that are expected to form in thermodynamic equilibrium. Consequently, we created a surface phase diagram for a wide range of partial gas pressures and temperatures.

Based on these monolayer polymorphs we created several bilayer polymorph candidates and optimized them using local geometry optimization. Including these bilayer polymorphs into the surface phase diagram revealed that densely-packed monolayers consisting of upright-standing TCNE molecules are energetically more favorable than bilayer polymorphs of the same total coverage. We, therefore, predict that a phase transition from flat-lying to upright-standing molecules occurs in the first adsorbate layer when the coverage is increased above that of a full monolayer of flat-lying TCNE molecules. Only at even lower temperature and/or higher partial gas pressure – when the packing density of the upright-standing molecules becomes too high – bilayer polymorphs form.

To verify our prediction, we compared them to an existing experiment. Therefore, we calculated the vibrational modes and frequencies as well as evaluated the molecular density of states (MODOS).

In accordance with the experiment, we found neutral, singly-, and doubly charged species of TCNE in the adsorbate polymorphs. Furthermore, we found these species in the same succession as in the experiment: doubly charged TCNE below the full monolayer coverage of flat-lying TCNE ($\Theta = 2.2\ N_{ads}/nm^2$), singly charged TCNE at twice that coverage and neutral TCNE at even higher coverage.

In contrast to the experimental findings, we did *not* see any polymorph that contains singly charged TCNE in the second adsorbate layer. Consequently, we do *not* corroborate the interpretation that charge transfer to the second layer occurs in this system. Rather, we found the singly charged species of TCNE in the form of upright-standing TCNE molecules in the first adsorbate layer. With these findings, we predict the existence of a phase transition in the first adsorbate layer from flat-lying molecules to upright-standing ones.

During this study, some starting points for future improvements became clear:
Currently, the structure search for bilayers is based on several guess polymorphs which are used as starting points for geometry optimizations. Although this was a useful approach to answer the specific questions raised above, it is a step that could be replaced by a more systematic approach. Namely, one could build an exhaustive set of second layer polymorphs by repeating the local adsorption geometry approach as proposed in chapter 4.3.

Furthermore, the current feature vector that describes pairwise interactions only depends on the relative position of molecules. In order to distinguish pairs of flat-lying molecules from upright-standing pairs and mixed pairs, we had to introduce an additional dimension that completely separates these feature spaces (see chapter 2.3.3). This approach works well as long as the orientation of the adsorbate molecules is restricted to two discrete possibilities (*lying* and *standing*). However, as soon as we include the possibility of continuously tilted adsorbates a full decorrelation between them would not make sense anymore. Instead one would have to include a continuous feature that accounts for the tilting angle of the adsorbate molecule with respect to the surface.

Appendix

A. Upright-Standing Adsorbates: parallel vs. perpendicular

As discussed in section 2.2.1, there are in total three types of local adsorption geometries: flat-lying ones, upright-standing ones with the bond of the two central carbon atoms parallel to the substrate surface, and upright-standing ones where that bond is oriented perpendicular to the surface (see Figure 2.1). As the TCNE molecule is approximately square, both upright-standing orientations result in a similar footprint on the substrate surface (as illustrated in Figure 1.2). Furthermore, it turned out that all local adsorption geometries with the C=C-parallel orientation have significantly stronger adsorption energies than their C=C-perpendicular counterparts. As illustrated in Figure A1, each of the *parallel* orientated local adsorption geometries leads to an adsorption energy that is at least $\Delta E = 270\ meV$ lower than the *perpendicular* oriented counterparts. With all those facts, there is no obvious reason why any energetically favorable polymorph should include upright-standing C=C-parallel TCNE instead of upright-standing C=C-perpendicular TCNE.

Still, there is one possibility that needs to be ruled out before we can safely neglect the *perpendicular* orientation: If the intermolecular interactions of adsorbates in the *perpendicular* orientation are highly favorable compared to the interactions of *parallel* adsorbates this could over-compensate for the weaker single-body interactions. To check the two-body interactions, a series of geometries containing all three combinations of these two upright-standing types are created.

Figure A2 shows the first series of such calculations. We use those two local adsorption geometries of the *parallel* and *perpendicular* orientation that have the smallest energetic difference – i.e. number 6 and 7 in Figure A1. Furthermore, we choose the positions of those molecules on the substrate lattice such that they stand in direct opposition – i.e. all other pairs can be expected to have even lower pair interactions. With these two choices, these pairs constitute the edge case with the minimal difference in one-body adsorption energy and high pair interaction energy. If the energy of none of these pairs can be decreased by using *perpendicular* oriented molecules, the same should be true for all other possible pairs.

Below the illustrations, the energetic difference between each configuration and the *parallel-parallel* configuration is listed. The index GO in ΔE_{GO} indicates that these energies were obtained after a local geometry optimization was applied to the depicted geometries. In contrast to that, ΔE_{SP} is obtained without geometry optimization (*single point*). Both configurations that contain molecules with the central C=C bond *perpendicular* to the surface are significantly higher in energy then the pair where both molecules are in the *parallel* orientation. Consequently, even if we account for pairwise interactions the *perpendicular* oriented TCNE do not become favorable.

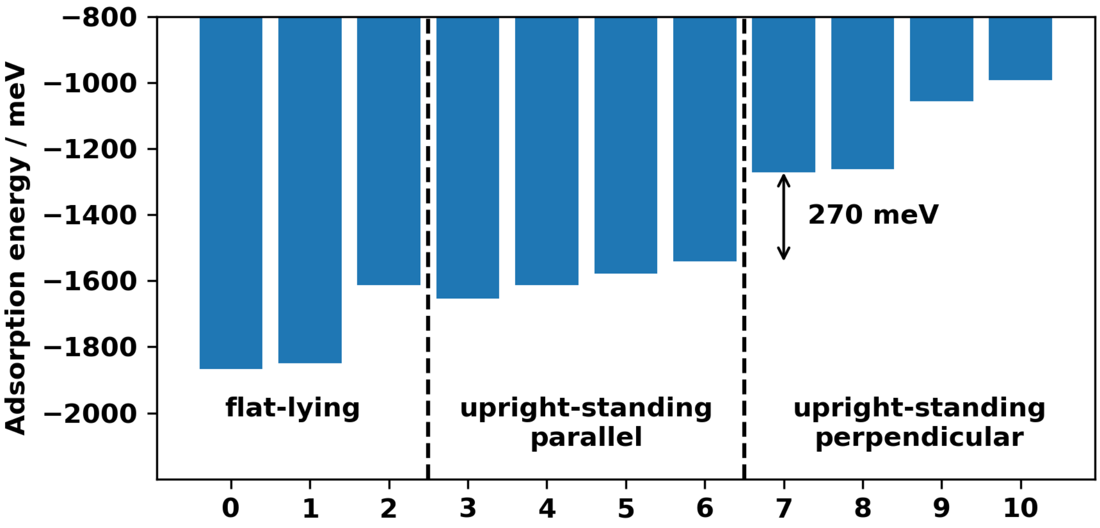

Figure A1.: Adsorption energy of all local adsorption geometries
Annotated is the smallest energetic difference between the two types upright-standing molecules. The order of the local adsorption geometries corresponds to that in Figure 2.1

To close even the last possible loophole, we created another series of calculations where the molecules are even closer. The distance between the local adsorption geometries cannot be chosen continuously as they are per definition restricted to specific positions relative to the substrate lattice. In this case, both local adsorption geometries are centered directly above a copper atom (*top* position). Therefore, the only pair that is closer than the one above is one were both molecules sit on neighboring substrate atoms (Figure A3. In these pairs the minimal distance between the molecules is only $d \approx 2.2$ Å. This distance is already smaller than the minimal distance of the predicted configurations ($d_{min} = 2.6$ Å). Furthermore, here we can only list *single point* energies as this distance is so small that during geometry optimization the molecules would simply slide into different local adsorption geometries to increase the intermolecular distance. Even in this extreme case, the pair interactions are too weak to make any of the pairs that include a *perpendicular* oriented TCNE molecule energetically favorable compared to the *parallel-parallel* pair.

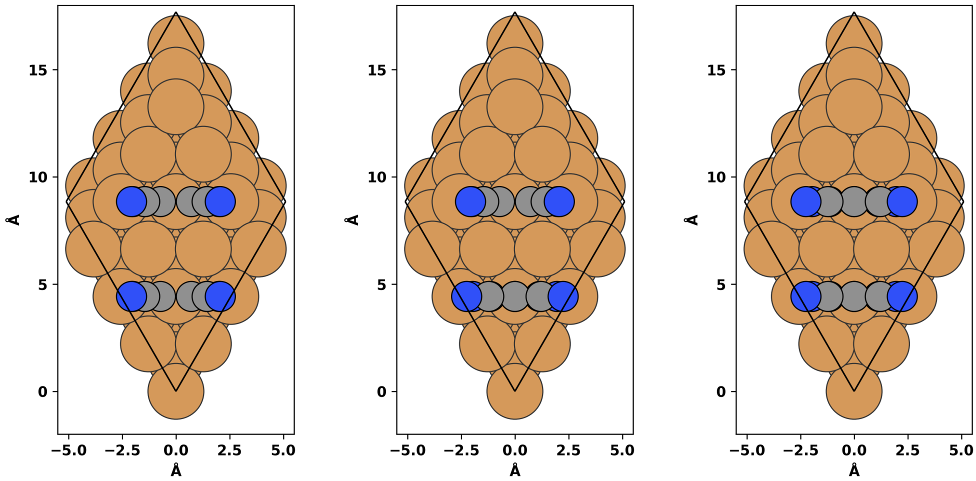

name	parallel-parallel	parallel-perpendicular	perpendicular-perpendicular
ΔE_{SP}	(reference)	$+180\ meV$	$+390\ meV$
ΔE_{GO}	(reference)	$+170\ meV$	$+360\ meV$

Figure A2.: Energetic difference between pairs of the two types of upright-standing local adsorption geometries calculated in the periodic unit cell indicated by the black rhombus.
Minimal intermolecular distance: $d \approx 4.4$Å

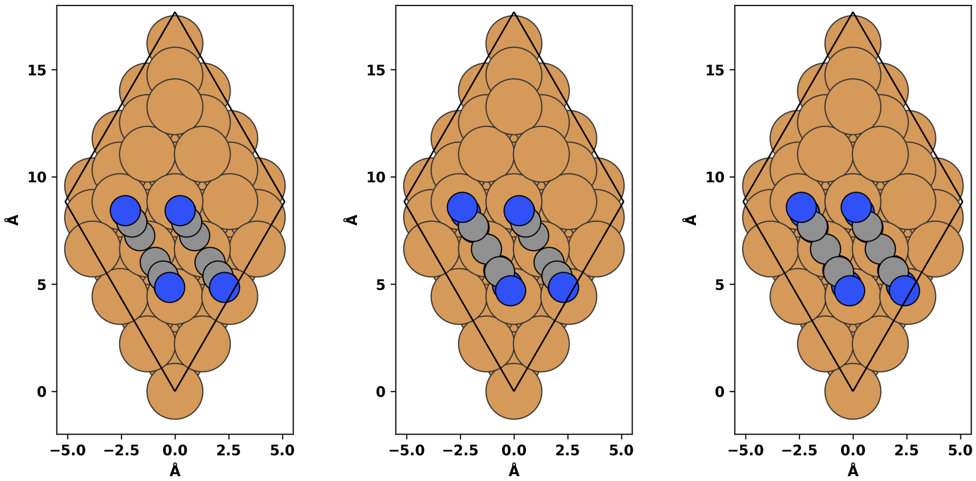

name	parallel-parallel	parallel-perpendicular	perpendicular-perpendicular
ΔE_{SP}	(reference)	$+130\ meV$	$+460\ meV$

Figure A3.: Energetic difference between pairs of the two types of upright-standing local adsorption geometries calculated in the periodic unit cell indicated by the black rhombus.
Minimal intermolecular distance: $d \approx 2.2$Å

B. Hyperparameter Optimization

To find an optimized set of hyperparameters, sweeps over a wide range of values are done for each hyperparameter (see Figure A4). For each value in these sweeps Bayesian linear regression is fitted with training sets of varying size, where each training set is chosen D-optimally using the current set of hyperparameters.

Two of these hyperparameters (*Intermolecular distance cutoff* d_{max} and *model uncertainty* σ_{model}) show only weak changes in the prediction accuracy on variation of the parameter. Therefore, σ_{model} is chosen such that it meets the convergence criterion that is used for the DFT settings.

The *Intermolecular distance cutoff* is chosen such that the intermolecular interactions have decayed to a few meV at this distance.

For the *decay power* n all values ≥ -2 lead to good prediction accuracies whereas more negative values (resulting in a steeper decay) lead to far worse predictions.

The *feature threshold* Δf yields best results for $\Delta f \leq 0.02$. For large training sets a smaller *feature threshold* could deliver even slightly better performance, however, at the cost of increased computational cost since smaller Δf results in a higher number of features.

We observe that larger *real space decay length* τ result in a better prediction accuracy for the largest training set ($n = 300$) while all smaller training sets yield better predictions for smaller τ. We, therefore, use the compromise of $\tau = 3.0$ Å.

Similarly, for the *feature correlation length* ξ a value of 1.0 has been chosen as it yields good predictions in conjunction with large training sets and, at the same time, avoids the steep increase in prediction error that occurs for smaller values of ξ with small training sets.

Finally, the *prior two-body uncertainty* σ_{2body}^{max} shows a rather flat behavior with slightly better results for larger values. Since larger values for σ_{2body}^{max} correspond to a broader prior probability density distribution for the two-body interactions, we rather use a more conservative guess of $\sigma_{2body}^{max} = 1.0\ eV$.

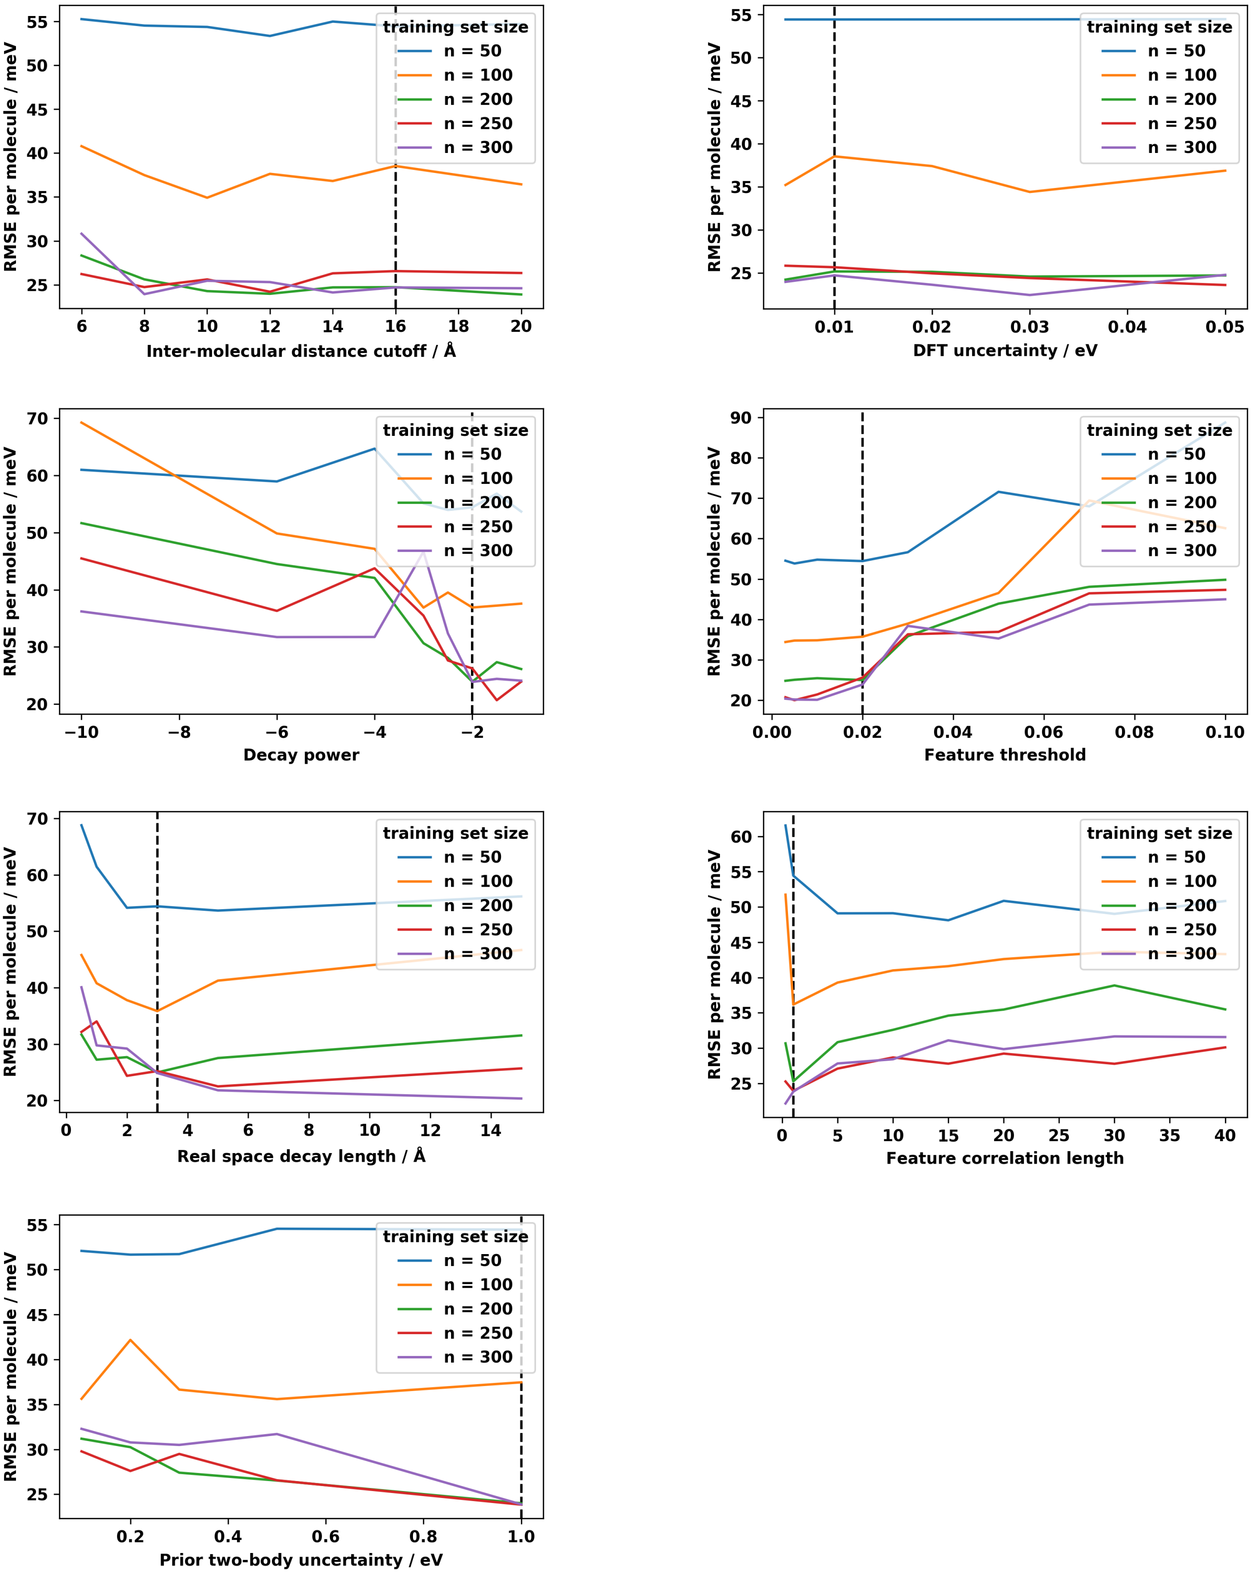

Figure A4.: Hyperparameter optimizations for D-optimally selected training sets increasing size. The dashed black line indicates the chosen parameters.

C. Additional Vibration Calculations

As described in chapter 5.2.1, discrepancies are found between specific calculated vibrational frequencies and experimentally observed frequencies. The frequencies of the two studied second-layer adsorbates are determined as $\tilde{\nu}_{cc} = 1482\ cm^{-1}$ and $\tilde{\nu}_{cc} = 1512\ cm^{-1}$, respectively. Both values fit well with the frequency that is calculated for the neutral TCNE molecule in the gas phase ($\tilde{\nu}_{cc} = 1491\ cm^{-1}$). In contrast to that, the experimentally observed frequency that is suspected to correspond to this adsorption geometry is significantly larger: $\tilde{\nu}_{cc} = 1565\ cm^{-1}$, $\Delta\tilde{\nu}_{cc} \approx (50-80)\ cm^{-1}$

Table A.1 summarizes the experimental data and opposes it to literature values as well as our calculated frequencies. From the first two columns, we see that the frequencies measured in the studied experiment match the literature values for the $\tilde{\nu}_{c-c}$ frequencies of TCNE in different charge states withing $\Delta\tilde{\nu}_{cc} \approx 10\ cm^{-1}$. These literature values also fit the calculated frequencies for TCNE in different charge states – except in the neutral form.

Experimental		Literature		Calculated (gas phase, FHI-*aims*, PBE)	
TCNE form	$\tilde{\nu}_{cc}$ / cm^{-1}	TCNE form	$\tilde{\nu}_{cc}$ / cm^{-1}	TCNE form	$\tilde{\nu}_{cc}$ / cm^{-1}
Low coverage	1275	$[TCNE]^{2-}$	1260	$TCNE^{2-}$	1273
Medium coverage	1375	$[TCNE]_2^{2-}$	1364, 1385	$TCNE^{1-}$	1375
High coverage	1565	TCNE	1570	TCNE	1491

Table A.1.: Experimental vibration frequencies of the C=C stretching mode measured via EELS in different adsorbate geometries as well as literature values taken from [51] and the corresponding frequencies calculated in FHI-*aims*.

To rule out problems that are specific to the used DFT code or the chosen exchange-correlation functional, we conduct additional vibration calculations in *Gaussian16*[52] using a 6-311G** basis set and the *PBE* and *B3LYP* functionals.

To compensate systematic errors, scaling factors – that are specific for the chosen functional and basis function – are commonly used. For full transparency, we list both the raw data, as well as scaled frequencies (Table A.3). For the calculation based on Pople type basis sets tabulated scaling factors are available (Ref. [53]). As there are no scaling factors specifically for the *6-311++G*** basis set, we used the equivalent factor for this basis set without the diffuse functions (*6-311G***).

XC functional	Basis functions	Scaling factor
PBE	$6-311G^{**}$	0.991
B3LYP	$6-311G^{**}$	0.967

Table A.2.: Scaling factors for vibrational calculations in *Gaussian16*; taken from Ref. [53]

For the numerical vibration calculations of FHI-*aims*, no tabulated correction factors are available. Empiric knowledge by Simon Erker (University of Technology, Graz) points to a scaling factor of $\eta_{cc} \approx 1.02$ for calculations concerning the C=C bond.

TCNE form	FHI-*aims* PBE raw data $\tilde{\nu}_{cc}$ / cm^{-1}	FHI-*aims* PBE scaled $\tilde{\nu}_{cc}$ / cm^{-1}	Gaussian PBE raw $\tilde{\nu}_{cc}$ / cm^{-1}	Gaussian PBE scaled $\tilde{\nu}_{cc}$ / cm^{-1}	Gaussian B3LYP raw $\tilde{\nu}_{cc}$ / cm^{-1}	Gaussian B3LYP scaled $\tilde{\nu}_{cc}$ / cm^{-1}
$TCNE^{2-}$	1273	1298	1250	1238	1280	1237
$TCNE^{1-}$	1375	1403	1382	1370	1424	1377
TCNE	1491	1521	1515	1501	1580	1528

Table A.3.: Calculated vibration frequencies of the C=C stretching mode of TCNE in the gas phase with different charge states. *'scaled'* indicates this use of scaling factors. Calculations done in FHI-*aims* use numerical atomic orbitals as listed in Appendix D while Gaussian calculations are conducted using the *6-311++G*** basis set.

Applying the assumed scaling factor to the data calculated in FHI-*aims* (second column) shifts the frequency of the neutral molecule into the correct direction. However, at the same time this scaling shifts the frequencies of the other two charge states away from the experimental values. Looking at the frequencies calculated in Gaussian using the *PBE* functional, we see that the scaled frequency for $TCNE^{2-}$ is too low by about 40 cm^{-1}, the frequency for $TCNE^{1-}$ fits nicely and the frequency for neutral TCNE is still too low by about 60 cm^{-1}. When we switch to the *B3LYP* functional the first two frequencies stay roughly the same. Interestingly, thou, the frequency of the neutral TCNE molecule shifts by 27 cm^{-1} towards the experimental value. Using this calculation scheme all three calculated frequencies end up within $\Delta\tilde{\nu} = 40\ cm^{-1}$ with respect to the experimental values.

However, benchmark papers name typical errors of $\Delta\tilde{\nu} = 45\ cm^{-1}$ for *PBE* and $\Delta\tilde{\nu} = 85\ cm^{-1}$ for *B3LYP* (both on a 6-31++G* basis set, averaged over 32 small molecules of the Gaussian G2 test set; see Ref. [54]). Similarly, Ref. [55] reports errors of about 4 % (corresponding to $\Delta\tilde{\nu} = 60\ cm^{-1}$ at $\tilde{\nu} = 1500\ cm^{-1}$). Consequently, the error that is observed in the vibrational frequencies of the neutral molecule is not unusual, but rather the nearly perfect agreements of the other calculated frequencies with the experimental values are coincidental.

D. DFT Settings

For all single point calculations (i.e. no geometry optimizations) settings are used as displayed in Listing 5.1. All surface calculations are done with seven layers of Cu. To reduce the computational costs while still describing processes like charge transfer adequately two different levels of quality are used for the basis sets of the substrate atoms. The uppermost three layers, which interact directly with the adsorbates, are described using more rigorous settings (see Listing 5.2). On the other hand, the lower lying substrate layers mainly serve as an electron reservoir and, hence, can be described using more rough settings (Listing 5.3). For a more detailed presentation of this approach, please see Ref. [56]. The number of both types of substrate layers – as well as all other significant settings – have been converged to an uncertainty in the range of $\Delta E = 10\ meV$. For the k-points Monkhorst-Pack grids [57] with a maximal k-point spacing of $\Delta k = \frac{2\pi}{80}$ are used.

Listing 5.1: Settings used in the single point DFT calculations

```
# General Settings:
xc  pbe
spin  none
charge  0
relativistic  atomic_zora scalar
occupation_type gaussian 0.1
k_points_external k_list.in

# Convergence Criteria:
sc_accuracy_rho 1e-2
sc_accuracy_etot  1e-5
sc_accuracy_forces  1e-3
sc_iter_limit 200

# Mixer:
preconditioner  kerker 1.5
charge_mix_param  0.05

# Other Settings:
vdw_correction_hirshfeld  .true.
RI_method lvl_fast
use_dipole_correction .true.
compensate_multipole_errors .true.
collect_eigenvectors  .false.
vdw_pair_ignore Cu Cu
```

Listing 5.2: Basis set settings for species Cu

```
species Cu

hirshfeld_param 59 10.9 2.4
nucleus 29
mass 63.546
l_hartree 6
cut_pot 4.6 2.0 1.0
basis_dep_cutoff 1e-3
radial_base 53 7.0
radial_multiplier 1

# Angular grid divisions
angular_grids specified
division 0.3478 50
division 0.6638 110
division 0.9718 194
division 1.1992 302
division 1.5920 434
outer_grid 434

# Minimal basis functions
valence 4 s 1.
valence 3 p 6.
valence 3 d 10.
ion_occ 4 s 0.
ion_occ 3 p 6.
ion_occ 3 d 9.

# Additional basis functions
```

```
ionic 4 p auto
hydro 4 f 7.4
hydro 3 s 2.6
hydro 3 d 5
```

Listing 5.3: Basis set settings for species Cu_reallylight

```
species Cu_reallylight

hirshfeld_param 59 10.9 2.4
nucleus 29
mass 63.546
l_hartree 4
cut_pot 3.5 1.5 1.0
basis_dep_cutoff 1e-3
radial_base 53 5.0
radial_multiplier 1

# Angular grid divisions
angular_grids specified
division 0.3478 50
division 0.6638 110
division 0.9718 194
outer_grid 194

# Minimal basis functions
valence 4 s 1.
valence 3 p 6.
valence 3 d 10.
ion_occ 4 s 0.
ion_occ 3 p 6.
ion_occ 3 d 9.

# Additional basis functions
ionic 4 p auto
```

Listing 5.4: Basis set settings for species C

```
species C

nucleus 6
mass 12.0107
l_hartree 6
cut_pot 4.0 2.0 1.0
basis_dep_cutoff 1e-4
radial_base 34 7.0
radial_multiplier 2

# Angular grid divisions
angular_grids specified
division 0.2187 50
division 0.4416 110
division 0.6335 194
division 0.7727 302
division 0.8772 434
outer_grid 434

# Minimal basis functions
```

```
valence 2 s 2.
valence 2 p 2.
ion_occ 2 s 1.
ion_occ 2 p 1.

# Additional basis functions
hydro 2 p 1.7
hydro 3 d 6
hydro 2 s 4.9
hydro 4 f 9.8
hydro 3 p 5.2
hydro 3 s 4.3
hydro 5 g 14.4
hydro 3 d 6.2
```

Listing 5.5: Basis set settings for species N

```
species N

nucleus 7
mass 14.0067
l_hartree 6
cut_pot 4.0 2.0 1.0
basis_dep_cutoff 1e-4
radial_base 35 7.0
radial_multiplier 2

# Angular grid divisions
angular_grids specified
division 0.1841 50
division 0.3514 110
division 0.5126 194
division 0.6292 302
division 0.6939 434
outer_grid 434

# Minimal basis functions
valence 2 s 2.
valence 2 p 3.
ion_occ 2 s 1.
ion_occ 2 p 2.

# Additional basis functions
hydro 2 p 1.8
hydro 3 d 6.8
hydro 3 s 5.8
hydro 4 f 10.8
hydro 3 p 5.8
hydro 1 s 0.8
hydro 5 g 16
hydro 3 d 4.9
```

List of Figures

Bibliography

[1] Guido van Rossum and Jelke de Boer. “Interactively Testing Remote Servers Using the Python Programming Language”. In: *CWi Quarterly* 4.4 (1991), pp. 283–303.

[2] David Ascher et al. “An Open Source Project Numerical Python David Ascher Paul F. Dubois Konrad Hinsen Jim Hugunin Travis Oliphant with Contributions from the Numerical Python Community.” In: (2001).

[3] J. D. Hunter. “Matplotlib: A 2D Graphics Environment”. In: *Computing In Science & Engineering* 9.3 (2007), pp. 90–95. DOI: `10.1109/MCSE.2007.55`.

[4] Ask Hjorth Larsen et al. “The Atomic Simulation Environment—a Python Library for Working with Atoms”. In: *Journal of Physics: Condensed Matter* 29.27 (2017), p. 273002.

[5] Donald Ervin Knuth and Duane Bibby. *The Texbook*. Vol. 3. Addison-Wesley Reading, 1984.

[6] Jutta Rogal and Karsten Reuter. “Ab Initio Atomistic Thermodynamics for Surfaces: A Primer”. In: *Experiment, Modeling and Simulation of GasSurface Interactions for Reactive Flows in Hypersonic Flights* Educational Notes RTO-EN-AVT-142, Paper 2 (2019), pp. 2–1 –2–18.

[7] Elizabeth C Beret and Luca M Ghiringhelli. “Ab Initio Atomistic Thermodynamics”. 2011. URL: `https://th.fhi-berlin.mpg.de/th/Meetings/DFT-workshop-Berlin2011/presentations/2011-07-18_Beret_Ghiringhelli.pdf`.

[8] Veronika Obersteiner et al. “Structure Prediction for Surface-Induced Phases of Organic Monolayers: Overcoming the Combinatorial Bottleneck”. In: *Nano Letters* 17.7 (2017), pp. 4453–4460. ISSN: 1530-6984. DOI: `10.1021/acs.nanolett.7b01637`.

[9] Michael Scherbela et al. “Charting the Energy Landscape of Metal/Organic Interfaces via Machine Learning”. In: *Physical Review Materials* 2.4 (2018), p. 043803. DOI: `10.1103/PhysRevMaterials.2.043803`.

[10] Lukas Hörmann et al. “Surface Structure Search Enabled by Coarse Graining and Statistical Learning”. In: (2018). arXiv: `1811.11702 [cond-mat]`. URL: `http://arxiv.org/abs/1811.11702` (visited on 01/10/2019).

[11] Veronika Obersteiner. “Computational Modeling of Organic-Inorganic Nanomaterials”. Dissertation. Graz University of Technology, 2017.

[12] Michael Scherbela. “Structure Prediction at Organic/Inorganic Interfaces Using Machine Learning”. Diploma thesis. Graz University of Technology.

[13] Wulf Erley and Harald Ibach. “Spectroscopic Evidence for Surface Anion Radical Formation of Tetracyanoethylene Adsorbed on Copper(111) at 100 K: A High-Resolution Electron Energy Loss Study”. In: *The Journal of Physical Chemistry* 91.11 (1987), pp. 2947–2950. ISSN: 0022-3654. DOI: `10.1021/j100295a059`.

[14] Ludwig Boltzmann. *Vorlesungen über die Principe der Mechanik.* Vol. 1. Leipzig: J.A. Barth, 1897. 628 pp.

[15] Yuhan Zhang et al. "Probing Carrier Transport and Structure-Property Relationship of Highly Ordered Organic Semiconductors at the Two-Dimensional Limit". In: *Physical Review Letters* 116.1 (2016), p. 016602. DOI: `10.1103/PhysRevLett.116.016602`.

[16] R. Fletcher. *Practical Methods of Optimization; (2Nd Ed.)* New York, NY, USA: Wiley-Interscience, 1987. ISBN: 978-0-471-91547-8.

[17] Adriana Supady, Volker Blum, and Carsten Baldauf. "First-Principles Molecular Structure Search with a Genetic Algorithm". In: *Journal of Chemical Information and Modeling* 55.11 (2015), pp. 2338–2348. ISSN: 1549-9596. DOI: `10.1021/acs.jcim.5b00243`.

[18] Jörg Behler and Michele Parrinello. "Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces". In: *Physical Review Letters* 98.14 (2007), p. 146401. DOI: `10.1103/PhysRevLett.98.146401`.

[19] Volker Blum et al. "Ab Initio Molecular Simulations with Numeric Atom-Centered Orbitals". In: *Computer Physics Communications* 180.11 (2009), pp. 2175–2196. ISSN: 0010-4655. DOI: `10.1016/j.cpc.2009.06.022`.

[20] Karsten Reuter and Matthias Scheffler. "Composition, Structure, and Stability of RuO2(110) as a Function of Oxygen Pressure". In: *Physical Review B* 65.3 (2001), p. 035406. DOI: `10.1103/PhysRevB.65.035406`.

[21] Donald A. McQuarrie. *Statistical Mechanics.* Harper & Row New York, 1976.

[22] Christoph J. Cramer. *Essentials of Computational Chemistry.* John Wiley & Sons, 2004. ISBN: 978-0-470-09182-1.

[23] Ask Hjorth Larsen et al. "The Atomic Simulation Environment—a Python Library for Working with Atoms". In: *Journal of Physics: Condensed Matter* 29.27 (2017), p. 273002. URL: `http://stacks.iop.org/0953-8984/29/i=27/a=273002`.

[24] Frank Jensen. *Introduction to Computational Chemistry.* 2nd ed. OCLC: ocm70707839. Chichester, England ; Hoboken, NJ: John Wiley & Sons, 2007. 599 pp. ISBN: 978-0-470-01186-7 978-0-470-01187-4.

[25] E. Schrödinger. "Quantisierung Als Eigenwertproblem". In: *Annalen der Physik* 384.4 (1926), pp. 361–376. ISSN: 1521-3889. DOI: `10.1002/andp.19263840404`.

[26] M. Born and R. Oppenheimer. "Zur Quantentheorie Der Molekeln". In: *Annalen der Physik* 389.20 (1927), pp. 457–484. ISSN: 1521-3889. DOI: `10.1002/andp.19273892002`.

[27] P. Hohenberg and W. Kohn. "Inhomogeneous Electron Gas". In: *Physical Review* 136 (3B 1964), B864–B871. DOI: `10.1103/PhysRev.136.B864`.

[28] W. Kohn and L. J. Sham. "Self-Consistent Equations Including Exchange and Correlation Effects". In: *Physical Review* 140 (4A 1965), A1133–A1138. DOI: `10.1103/PhysRev.140.A1133`.

[29] John P. Perdew et al. "Prescription for the Design and Selection of Density Functional Approximations: More Constraint Satisfaction with Fewer Fits". In: *The Journal of Chemical Physics* 123.6 (2005), p. 062201. ISSN: 0021-9606, 1089-7690. DOI: `10.1063/1.1904565`.

[30] A. D. Becke. "Density-Functional Exchange-Energy Approximation with Correct Asymptotic Behavior". In: *Physical Review A* 38.6 (1988), pp. 3098–3100. DOI: `10.1103/PhysRevA.38.3098`.

[31] Chengteh Lee, Weitao Yang, and Robert G. Parr. “Development of the Colle-Salvetti Correlation-Energy Formula into a Functional of the Electron Density”. In: *Physical Review B* 37.2 (1988), pp. 785–789. DOI: `10.1103/PhysRevB.37.785`.

[32] John P. Perdew, Kieron Burke, and Matthias Ernzerhof. “Generalized Gradient Approximation Made Simple”. In: *Physical Review Letters* 77.18 (1996), pp. 3865–3868. DOI: `10.1103/PhysRevLett.77.3865`.

[33] Axel D. Becke. “Density-functional Thermochemistry. IV. A New Dynamical Correlation Functional and Implications for Exact-exchange Mixing”. In: *The Journal of Chemical Physics* 104.3 (1996), pp. 1040–1046. ISSN: 0021-9606. DOI: `10.1063/1.470829`.

[34] Jianwei Sun, Adrienn Ruzsinszky, and John P. Perdew. “Strongly Constrained and Appropriately Normed Semilocal Density Functional”. In: *Physical Review Letters* 115.3 (2015), p. 036402. DOI: `10.1103/PhysRevLett.115.036402`.

[35] Jianmin Tao et al. “Climbing the Density Functional Ladder: Nonempirical Meta-Generalized Gradient Approximation Designed for Molecules and Solids”. In: *Physical Review Letters* 91.14 (2003), p. 146401. ISSN: 0031-9007. DOI: `10.1103/PhysRevLett.91.146401`. pmid: `14611541`.

[36] Carlo Adamo and Vincenzo Barone. “Toward Reliable Density Functional Methods without Adjustable Parameters: The PBE0 Model”. In: *The Journal of Chemical Physics* 110.13 (1999), pp. 6158–6170. ISSN: 0021-9606. DOI: `10.1063/1.478522`.

[37] R. S. Mulliken. “Electronic Population Analysis on LCAO–MO Molecular Wave Functions. I”. In: *The Journal of Chemical Physics* 23.10 (1955), pp. 1833–1840. ISSN: 0021-9606. DOI: `10.1063/1.1740588`.

[38] Alexandre Tkatchenko and Matthias Scheffler. “Accurate Molecular Van Der Waals Interactions from Ground-State Electron Density and Free-Atom Reference Data”. In: *Physical Review Letters* 102.7 (2009), p. 073005. DOI: `10.1103/PhysRevLett.102.073005`.

[39] Lukas Hörmann. “Overcoming the Configurational Explosion for Surface Structure Search by Coarse Grained Modeling and Bayesian Learning”. Graz University of Technology.

[40] Matthias Rupp et al. “Fast and Accurate Modeling of Molecular Atomization Energies with Machine Learning”. In: *Physical Review Letters* 108.5 (2012). ISSN: 0031-9007, 1079-7114. DOI: `10.1103/PhysRevLett.108.058301`.

[41] Bradley Jones and Peter Goos. “I-Optimal Versus D-Optimal Split-Plot Response Surface Designs”. In: *Journal of Quality Technology* 44 (2012), pp. 85–101. DOI: `10.1080/00224065.2012.11917886`.

[42] nist_engineering_statistics_handbook. *5.5.2. What Is a Computer-Aided Design?* URL: `https://www.itl.nist.gov/div898/handbook/pri/section5/pri52.htm` (visited on 07/31/2018).

[43] A. C. Atkinson and A. N. Donev. *Optimum Experimental Designs*. Oxford Statistical Science Series. Oxford, New York: Oxford University Press, 1992. 344 pp. ISBN: 978-0-19-852254-6.

[44] Fedorov, Valerii. *Theory of Optimal Experiments*. Academic Press New York and London, 1972.

[45] Michael Scherbela et al. “Leaving the Valley: Charting the Energy Landscape of Metal/Organic Interfaces via Machine Learning”. In: (2017). arXiv: `1709.05417 [cond-mat]`. URL: `http://arxiv.org/abs/1709.05417` (visited on 10/12/2017).

[46] Wulf Erley. “Reflection-Absorption Infrared Spectroscopy of Tetracyanoethylene Adsorbed on Copper(111): Observation of Vibronic Interaction”. In: *The Journal of Physical Chemistry* 91.24 (1987), pp. 6092–6094. ISSN: 0022-3654, 1541-5740. DOI: `10.1021/j100308a007`.

[47] Wulf Erley. “Reflection-Absorption Infrared Spectroscopy of Adsorbates on a Cu(111) Single Crystal Surface”. In: *Journal of Electron Spectroscopy and Related Phenomena* 44.1 (1987), pp. 65–78. ISSN: 0368-2048. DOI: `10.1016/0368-2048(87)87008-1`.

[48] Marcus D. Hanwell et al. “Avogadro: An Advanced Semantic Chemical Editor, Visualization, and Analysis Platform”. In: *Journal of Cheminformatics* 4.1 (2012), p. 17. ISSN: 1758-2946. DOI: `10.1186/1758-2946-4-17`.

[49] C. J. Nelin, P. S. Bagus, and M. R. Philpott. “The Nature of the Bonding of CN to Metals and Organic Molecules”. In: *The Journal of Chemical Physics* 87.4 (1987), pp. 2170–2176. ISSN: 0021-9606. DOI: `10.1063/1.453142`.

[50] Romaner Lorenz. “Modelling of Organic Semiconductors and Their Interaction with Metallic Surfaces”. Dissertation. Graz University of Technology, 2007.

[51] Joel S. Miller. “Tetracyanoethylene (TCNE): The Characteristic Geometries and Vibrational Absorptions of Its Numerous Structures”. In: *Angewandte Chemie International Edition* 45.16 (2006), pp. 2508–2525. ISSN: 1521-3773. DOI: `10.1002/anie.200503277`.

[52] M. J. Frisch et al. “Gaussian~16 Revision B.01”. In: (2016). Gaussian Inc. Wallingford CT.

[53] Russell D. Johnson III. *NIST Computational Chemistry Comparison and Benchmark Database NIST Standard Reference Database Number 101*. 2018. URL: `http://cccbdb.nist.gov/`.

[54] Kevin E. Riley, Bryan T. Op’t Holt, and Kenneth M. Merz. “Critical Assessment of the Performance of Density Functional Methods for Several Atomic and Molecular Properties”. In: *Journal of chemical theory and computation* 3.2 (2007), pp. 407–433. ISSN: 1549-9618. DOI: `10.1021/ct600185a`. pmid: `19002267`.

[55] Shawn T. Brown, Jonathan C. Rienstra-Kiracofe, and Henry F. Schaefer. “A Systematic Application of Density Functional Theory to Some Carbon-Containing Molecules and Their Anions”. In: *The Journal of Physical Chemistry A* 103.20 (1999), pp. 4065–4077. ISSN: 1089-5639. DOI: `10.1021/jp984354c`.

[56] Andreas Jeindl. “Influence of Molecule Size on Surface Polymorph Formation: An Ab-Initio Study with Machine Learning”. Diploma thesis. Graz University of Technology, 2018.

[57] Hendrik J. Monkhorst and James D. Pack. “Special Points for Brillouin-Zone Integrations”. In: *Physical Review B* 13.12 (1976), pp. 5188–5192. DOI: `10.1103/PhysRevB.13.5188`.