Deep Learning Methods for Catalyst Surface and Interface Structure Analysis

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Department of Chemical Engineering

Junwoong Yoon

B.S. Chemical Engineering, University of California, Berkeley M.S. Machine Learning, Carnegie Mellon University

> Carnegie Mellon University Pittsburgh, PA

> > May 2022

 \bigodot Junwoong Yoon, 2022 All rights reserved.

Abstract

The increase in global energy demand and raised environmental concerns have motivated the design of novel materials for energy-related applications. However, the design of ever-complicating materials for emerging energy technologies is currently bottlenecked by limited resources to understand complex surface and interface structures and property relationships. In the first part of this thesis, we develop a tandem framework that combines a molecular thermodynamic theory and molecular dynamics simulations in an attempt to investigate solid interfacial phenomena and to discuss how deep learning methods can improve the framework as a next step. In the second part, we develop a set of deep learning methods that solve various materials and catalyst design problems including property, structure, and stability analysis. We present a graph neural networks architecture to learn the optimal representations of heterogeneous catalysis systems for the accurate prediction of adsorption/binding energies. Then we extended the approach to approximate ground-state structures of the catalysis systems by incorporating differentiable optimization methods into the graph neural networks architecture. We further develop a general deep reinforcement learning framework to identify the metastability of alloy catalyst surfaces by exploring possible surface reconstructions and their associated kinetic barriers under reaction conditions. With these advanced data-driven methods that understand the surface and interfacial phenomena, we open up new avenues for accelerated materials and catalyst discovery.

Thesis Supervisor: Zachary W. Ulissi

Title: Associate Professor of Chemical Engineering

Acknowledgments

First and foremost I wish to express my gratitude for the support of my advisor, Zachary Ulissi. He is a great mentor who guided and supported me throughout my PhD. He provided me with insightful comments and helpful suggestions on any kinds of problems, issues, and concerns I had. He also allowed me to have freedom to explore research directions, and to choose my own projects. This greatly helped me in developing and acquiring skill sets that are essential for my future career. He is also an amazing leader for our group who creates a positive, collaborative, and supportive group environment.

I gratefully acknowledge funding packages from the Department of Chemical Engineering at Carnegie Mellon University that provided full financial support throughout my PhD. I am also grateful for the support of Department of Energy, National Energy Research Scientific Computing Center, and NVIDIA Corporation. They provided funding and computing resources for the completion of my research projects.

I feel lucky and blessed that I had many supportive collaborators. It has been a rewarding experience to work with great people on various projects. In particular, I would like to thank Seoin Back, Kevin Tran, Mingjie Liu, and Muhammed Shuaibi who I closely worked with. As we all shared an office, we were able to easily interact, discuss, and collaborate. I also have to be grateful to Zhonglin Cao, Rajesh Raju, Robert Burnley, and Yuyang Wang who have been incredibly helpful in completing a long, collaborative project.

I am also thankful to my thesis committee members, Andrew Gellman, Amir Barati Farimani, and Chrysanthos Gounaris. They provided me with valuable advice and suggestions for my projects. Particularly, Andrew and Amir were great mentors and principal investigators who led one of my research projects along with my advisor. The biweekly discussion with those mentors have been very helpful in completing the project. Another group of people I have to acknowledge are my friends and colleagues who I am fortunate to have with me. Kexin Zhu, Seoin Back, Yonwup Song, Mingjie Liu, Muhammed Suaibi, Jenny Zhan, Yilling Yang, Aini Palizhati, Zhitao Guo, and many others. All the things we've done and shared will become an unforgettable part of my memory at CMU.

Finally, this long journey would not have been possible without the support of my family, including my mom Jongboon Park, my father Wonyong Yoon, and my sister Ina Yoon. Being a part of this family, and their existence were really the best support and the best motivation for me. And I am always grateful for their love.

Contents

1	Intr	oducti	ion	23
	1.1	Surfac	e and Interface for energy technologies	23
	1.2	Parad	igm of Materials Science	25
	1.3	Overv	iew of Deep Learning in Materials Science	28
	1.4	Thesis	Objectives	29
Pa	art I:	Mode	ling and Simulation for solid interface	33
2	Capturing Structural Transitions in Surfactant Adsorption Isotherms			ıs
	at Solid/Solution Interfaces			35
	2.1	Abstra	act	35
	2.2	Introd	uction	36
	2.3	Metho	ds	40
		2.3.1	Molecular Thermodynamic Theory (MTT)	40
		2.3.2	Determining parameters for validation of the MTT model:	
			Lennard-Jones sphere surfact ant model on a graphene sheet $% \mathcal{A}$.	43
		2.3.3	Determining parameters for aqueous $SDS/graphene$ interface .	46
	2.4	Result	s and discussion	50
		2.4.1	Validation of the model	50
		2.4.2	Real surfactant system - aqueous SDS/graphene interface	51

	2.5	Conclusions	58	
Pa	art II	I: Deep Learning for Design of Materials and Catalyst Surfaces	59	
3	3 Convolutional Neural Network of Atomic Surface Structures To			
	Pre	dict Binding Energies for High-Throughput Screening of Cat-		
	alys	sts	61	
	3.1	Abstract	62	
	3.2	Introduction	62	
	3.3	Representation of Surface Structures	65	
	3.4	Graph Convolutional Neural Networks	69	
	3.5	Results	76	
		3.5.1 Binding Energy Prediction	76	
		3.5.2 Atomic Contributions to Binding Energies	76	
		3.5.3 Automated Failed DFT Calculation Detection	77	
	3.6	Conclusion	80	
	3.7	Dataset Preparation and Methods	84	
4	Dif	${ m ferentiable}$ Optimization for the Prediction of Ground State Struc-	-	
	ture	es (DOGSS)	87	
	4.1	Abstract	87	
	4.2	Introduction	88	
	4.3	Graph Neural Network	90	
		4.3.1 Representation of Surface Structures	90	
		4.3.2 Convolutional Layers - Message Passing Scheme	91	
		4.3.3 Differentiable Optimization	93	
	4.4	Dataset Preparation	95	
	4.5	Practical Motivation	95	
	4.6	Results	97	

		4.6.1	Loss (per-atom distances)	97
		4.6.2	Properties of Predicted Structures	101
		4.6.3	Reduction in Actual DFT Optimization Steps	101
	4.7	Concl	usion \ldots	103
-	Ъ	•		
9	Dee	ep rein	forcement learning for predicting kinetic pathways to sur	r-
	face	e recon	struction in a ternary alloy	107
	5.1	Abstra	act	107
	5.2	Introd	luction	108
	5.3	Metho	pds	112
		5.3.1	Reinforcement Learning Background	112
		5.3.2	Actor-Critic TRPO	115
		5.3.3	State Representation	116
		5.3.4	Action Space	116
		5.3.5	Reward Function	118
		5.3.6	Experiments	119
		5.3.7	Surface Optimization Baseline	119
	5.4	Result	ts and Discussion	122
		5.4.1	Baseline Surface Configurations	122
		5.4.2	DRL Training Summary	124
		5.4.3	Surface Configuration Exploration	125
		5.4.4	Reconstruction Pathway and Transition State	128
		5.4.5	Transition State Verification	128
	5.5	Concl	usion	131
6	Cor	nclusio	n and Future Work	133
Bi	blio	graphy		139
	Diolography 100			

List of Figures

- 2-1 Atomistic simulation snapshots for the adsorption of aqueous SDS surfactants on the graphene surface. The SDS surfactants direct the formation of structured aggregates from monolayer at a low surface concentration to surface micelle at high surface concentration. The hydrophilic surfactant headgroups are composed of sulfur (yellow) and oxygen (red) atoms, and the hydrophobic tailgroups are composed of carbon (green) and hydrogen (white) atoms. Water molecules are not shown for clarity.

- 2-3 (a) Free energy change for a Lennard-Jones sphere surfactant as a function of the height above the graphene surface obtained using ABF simulation. $\Delta \mu^0$ is determined as the difference in the free energy of the surfactant when it is in the bulk solution (~1.2 nm) and when it is at the interface (~0.33 nm). (b) Free energy required to concentrate surfactants in the initial state as a function of the surface concentration obtained using our simulation model (scatter) and the benchmark (line).

- 2-5 The surface concentration of surfactants as a function of the bulk mole fraction in the simplified surfactant system obtained using our MD/MTT model (scatter) and the benchmark.

- 2-7 A joint plot that shows the SDS surface concentrations corresponding to each aqueous SDS-graphene system with 30 SDS (blue) and 36 SDS (orange), and the average number of SDS molecules (black) found in each peak (defined in Figure 6) as a function of the bulk mole fraction. The two red vertical dashed lines represent the corresponding bulk mole fractions of the CSAC (left) and the CMC (right). Four structural phase regions are highlighted with different colors and labeled as phases a, b, c, and d. The simulation snapshots for the structure of the SDS molecules in each phase for the system with 30 SDS molecules are shown above the joint plot. Water molecules are not shown for clarity.

3-1A graphical representation of converting an atomic structure containing \mathbf{n} atoms into numerical inputs for the convolutional neural network. (a) 9 basic atomic properties are presented by one-hot encoding[226] to prepare the atomic feature vectors. (b) Neighbor information was encoded using the solid angle (Ω) based on the Voronoi polyhedron. The grey skeleton indicates the Voronoi polyhedron of Cu atom, and the solid angle between Fe atom and the shared plane of Cu and Fe polyhedrons is marked. Other nearest neighbor atoms were omitted for simplicity. Color code for atoms: Cu (blue), Fe (brown), C (grey), O (red). (c) The connectivity distances from adsorbate to all atoms in surface structures are counted. A side view of the surface structure with adsorbate CO molecule, surface atoms up to the second layer and their corresponding connectivity distances are presented as an example. (d). A schematic representation of three atomic structural data that could be used in the CNN. "Initial geometry", "Final geometry", and "Final adsorbate coordination geometry" are based on initial structures, final structures, and initial structures with adsorbate connectivity distances, respectively. The numbers in the DFT part (solid box) represent the initially set or relaxed atomic distances, and the integer numbers in the CNN part (dotted box) are the connectivity distances. More details can be found in the main text. 69

3-2 A process of high-throughput workflow for the catalyst screening. . . 70

3-3	Schematic illustration of the convolutional neural network on top of	
	the graph representation of atomic surface structures. Dark red and	
	red blocks correspond to atomic and neighbor feature vectors as in	
	Figure 3.3. Pink blocks correspond to the connectivity distance vec-	
	tors used as a filter to exclude atoms that are too far from adsorbates	
	(higher than the connectivity distance of 2). More details are elabo-	
	rated in the main text. \ldots	71
3-4	Convolutional neural network architecture used in this work. The	
	values in parenthesis correspond to the number of parameters for CO	
	binding energy prediction (Table S2)	73
3-5	Mean absolute error (MAE) of test set with respect to the number of	
	training atomic structures for binding energy predictions of (a) CO	
	and (b) H. The horizontal dashed lines correspond to the standard	
	deviation of DFT binding energies. As a comparison, we added the	
	prediction error from our previous study (green and red star). $[205]$	
	Two-dimensional histograms of DFT-calculated and CNN-predicted	
	binding energies of (c) CO and (d) H	75
3-6	Graphic representations of atomic contributions to CO binding ener-	
	gies at (a) top and (b) bridge site of Cu (211) surface. Darker spheres	
	indicate higher contributions. Only atoms within the connectivity dis-	
	tance of 2 are considered and transparent spheres indicate atoms with	
	connectivity distances higher than 2	78

3-7	Top view of CO adsorption on $Cu_3Al(211)$ surface. Darker spheres in-	
	dicate higher contributions. Dark brown spheres are Cu atoms directly	
	interacting with CO adsorbate (distance $= 1$), light grey spheres are	
	Cu atoms not interacting with the CO (distance $= 2$). Lilac spheres	
	are Al (distance $= 2$). Transparent spheres indicate atoms with con-	
	nectivity distances higher than 2	79
3-8	Three types of outliers detected during the analysis. Surface ID con-	
	sists of MPID, miller index, shift value and orientation (True and False	
	correspond to top and bottom of surface structures, respectively). $\ . \ .$	81
3-9	Two iterations of the outlier detection for CO binding energy prediction.	82
3-10	Two iterations of the outlier detection for H binding energy prediction.	82
4-1	Workflow to accelerate electronic structure methods using DOGSS as	
	a pre-processing model. (1) DOGSS predicts the ground-state struc-	
	ture by relaxing the initial structure using the learned harmonic force	
	field parameters. (2) and (3) electronic structure methods employ the	
	predicted structure as a starting configuration, thereby converge to	
	the ground-state in fewer optimization steps.	89
4-2	Structure of the graph neural network with embedded differentiable	
	optimization. Initial node \boldsymbol{v}^0 and edge \boldsymbol{u}^0 features are fed to T con-	
	volutional layers. Two separate hidden networks take the updated	
	features to learn properties of distinct springs connecting every atom	
	and neighbor atom pairs: spring constants ${m k}$ and equilibrium spring	
	distances r_e . With these learned harmonic force field parameters, har-	
	monic PES is computed and minimized using the gradient descent to	
	approximate the ground-state structures.	91

- 4-3 Histograms showing distributions of the individual per-atom distances $||x x_{min}^*||_2$ for surface atoms in each of the three datasets: (left) bare surfaces, (center) H adsorption, and (right) CO adsorption. For all three datasets, the distributions for atoms in the initial structures are shown in orange and the distributions for atoms in the predicted ground-state structures by DOGSS are shown in blue. The insets of the center and the right figures show the per-atom distance distributions for adsorbate atoms in H and CO adsorption datasets respectively with visualized example configurations.

5 - 1	(a) Overview of CatGym framework. At each timestep t , the state
	space s_t consists of features of surface structure c_t , including its cur-
	rent energy, fingerprint, and atom coordinates. The state space along
	with reward r_t , which is calculated by critic network based on current
	energy, are fed into the actor network, and the actor network decides
	which action to take. The action will modify c_t to new surface struc-
	ture c_{t+1} . This process repeats until the episode is over. (b) Four
	types of action the actor can choose and their corresponding effects
	on the energy. If the agent choose to move an atom, it has to fur-
	ther pick which atom to move and the distance (within the range of
	-0.1Å to 0.1Å) in the x, y and z direction. (c) Evolution of the re-
	construction pathways to a minimum configuration generated by DRL
	agent trained with CatGym. ΔE represents the relative energy of the
	catalyst structure at each timestep with respect to the initial state
	energy
5-2	The potential energy changes and the changes in the E_{diff} and tem-
	peratures during MH simulation $\ldots \ldots 121$
5-3	$\label{eq:configurations} Configurations of the Ni(green)-Pd(blue)-Au(yellow) \ ternary \ catalyst.$
	(a) Initial catalyst structure, (b)-(e) global minimum, second and third
	lowest energy configurations, and global maximum configuration ob-
	tained via brute-force method, respectively; and (f) global minimum
	configuration obtained via minima hopping (MH) method. Atoms
	in the bottom two layers with crosses represent fixed atoms that are
	constrained from moving
5-4	Average reward vs. episode across 10 different seeds. Red line repre-
	sents the average reward, and the blue shadow represents the standard
	deviation. $\ldots \ldots 124$

5 - 5	Projection of Ni-Pd-Au configurations identified by (a) MH and (b)	
	DRL method in the reduced dimensional space using the first two	
	principal components that capture the most structural variance in the	
	all possible surface arrangements. Numbers in red indicate unique	
	clusters determined by Kmeans. Boundaries between the clusters are	
	drawn using sold lines. (c) Example Ni(green)-Pd(blue)-Au(yellow)	
	configurations for each of the clusters explored by DRL method, clus-	
	ters 0, 1, 2, 5, 6, 7, 8, 9, and 10. Average relative energies (ΔE) of	
	transition states (blue) and minima (black) for each of the clusters are	
	labeled below the cluster numbers (red).	125
5-6	Projection of all possible Ni-Pd-Au surface configurations (total 560	
	configurations) identified by brute-force permutation method onto the	
	reduced dimensional space using PCA. Numbers in red represent dis-	
	tinct clusters determined by K-means clustering algorithm, and their	
	boundaries are drawn using solid lines in the PCA plot. Example	
	Ni(green)-Pd(blue)-Au(gold) surface configuration for each of the clus-	
	ters are also shown	127
5-7	Example reconstruction pathways to local minima generated by the	
	DRL agent. Ni(green)-Pd(blue)-Au(gold) configurations of the tran-	

DRL agent. Ni(green)-Pd(blue)-Au(gold) configurations of the tran-
sition state and the minimum for each of the pathways are also illus-
trated next to the pathway plot. Energy values in this figure (ΔE)
represent the relative energies of the each of the configuration with
respect to the initial state energy

List of Tables

- 3-1 Hyperparameters of convolutional neural network optimized by Sigopt. 72
- A comparison of our DFT data and the number of parameters of 3-2 our model to literature values, where the goal was to predict binding energies. We counted all unique components (element, stoichiometry, space group, facet) in literature and our DFT data. For the previous studies that did not explicitly mention the number of total DFT data or parameters, we estimated based on the presented results. The number of parameters indicate how complicated the model is. Specific parameter components of our model are summarized in Table S2. . . 74Properties used in node feature vector \boldsymbol{v}_i 92 4-1 Example OFM representations for Pt and Au atoms. 4 - 2924 - 394Summary of the ground-state structure prediction performance of DOGSS 4-4 and SchNet on test sets in the three different datasets. 96 Performance of DOGSS using the original 80/10/10 split with three 4-5different types of potential: Harmonic, Morse, and Lennard-Jones (LJ) potentials with the gradient descent. 99

- 4-6 Generalization performance of DOGSS using different data splitting methods: Original split (80/10/10), Leave-"Ni"-out (leaving all structures containing Ni out of training sets and testing DOGSS only on the remainder), and Leave-"Zn" and "Co"-out (leaving all structures containing Zn and Co out of training sets and testing DOGSS only on the remainder). We selected Ni, Zn, and Co elements in order to maintain similar train/val/test sizes over different data split. 102

Chapter 1

Introduction

1.1 Surface and Interface for energy technologies

Solutions to both raised concerns about climate change and rising global energy demand are critical to the future prosperity of a modern world. The increased use of renewable energies from abundant sources of energy such as solar and wind is a promising solution to address the world's rising energy needs while reducing climate change. However, they have an intermittency problem in which no solar power is available at night while winds can stop suddenly.[149, 169, 240] Efficient energy storage and converting technologies are thus the key to solving the intermittency problem by continuously providing renewable power with its full potential even when the energy sources are not available.

Various emerging devices for energy conversion and storage have shown great potential to make a significant impact on the advancement of energy technologies. For example, solid-state batteries can be alternative energy storage devices to current lithium-ion batteries in terms of safety and energy density.[1, 52] Fuel cells are promising energy conversion devices that have high theoretical energy density and efficiency but low environmental effects.[156] However, these emerging devices are currently to some extent bottlenecked by one or several of the key materials for energy conversion and storage processes. Performance of the solid-state batteries is currently hampered due to the bottleneck of the solid electrolyte material and other materials for mitigating interfacial instabilities while the high cost and limited durability of electrocatalysts block the use of fuel cells.[95, 109, 221] The Discovery of novel materials with desired properties thus can bring a breakthrough in emerging energy technologies.

Surface and interface designs can be an efficient strategy to enhance the performance of energy technologies as they are closely related to the properties of novel energy materials. For example, the surface of graphene can be functionalized or modified to manipulate its physical and chemical properties and improve its performance in targeted applications. Graphene is a two-dimensional material that exhibits unique structure and outstanding properties that make it excellent in nanotechnologies and applications. In particular, for energy applications, functionalized graphene can potentially benefit various energy devices including batteries, fuel cells, and other energy conversion and storage devices. [57, 66, 156] The surface and interface of heterogeneous catalysts are also essential to their performance as they are often considered to be active sites for energy-related catalytic reactions such as oxygen reduction reactions occurring at the cathodes in fuel cells and electrochemical CO2 reduction reactions that convert CO2 into valuable chemicals or fuels. [62, 96, 222, 240] Therefore, developing materials science tools for understanding surface and interfacial phenomena can unlock the full potential of those novel materials that will play a central role in energy technologies.

1.2 Paradigm of Materials Science

The development of materials science can be represented by four paradigms. The first paradigm is purely empirical and is relying on the intuitive observation experience with no scientific justification. A few centuries ago, the foundation of theoretical models characterized by the formulation of various physics laws in the form of mathematical equations shifted the paradigm. Thermodynamics provides theoretical intuition about possible existing phases, stabilities, and other thermodynamic properties of materials.[68, 135] Kinetics tells about the effects of various factors including concentration, temperature, pressure, catalyst, etc on reaction rates to describe the whole material process. [135] The recent development of quantum mechanics further provides tools to understand the relationship between atomic structure and material properties. Theoretical models enable the design of new materials with a fundamental understanding of materials science, but these models easily become too complex for many real-world problems where analytical solutions typically do not exist.

In the field of computational materials science, the invention of computers a few decades ago along with the theoretical advances of quantum and classical models have allowed simulations of complex practical problems based on the theoretical models of the second paradigm, leading to the third paradigm of materials discovery. At this period, density functional theory (DFT) and molecular dynamics (MD) simulations emerged as the most popular computational methods. DFT calculates materials properties by approximating and solving the Schrodinger equation based on the first principles and enables computational high-throughput screening of materials. MD simulations allow atoms and molecules to interact for a fixed period of time to study the dynamic evolution of the system by analyzing the physical movements. The success of the computational methods has made tremendous contributions to innovations in materials.[59, 88, 89, 114, 142, 182, 185]

The third scientific paradigm has guided materials discovery for decades. Mean-

while, materials scientists begin to look for ever-complicating materials for emerging and more specific applications such as catalyst development for hydrogen and methane conversions for energy harvesting and storage.[240] Nevertheless, the design of such complex materials using the computational screening method is greatly limited due to the computational cost that cannot deal with a massive search space arising from the complex materials structure and property relations. Consequently, the need for more efficient design tools has been emphasized in the field of computational chemistry and materials science.

Recently, the trend of materials research has shifted towards the stage of big datadriven science also known as the fourth paradigm of materials science. [5, 129, 133] The mass data generated by experiments and computational simulation methods has enabled the use of data-driven techniques such as machine learning in materials research, thereby opening up new avenues for accelerated materials discovery. Machine learning comprises models that learn from existing data. Unlike traditional computational approaches that employ hard-coded algorithms provided by human experts, machine learning models are capable of making predictions by learning or acquiring information from the data itself. [152] The training of machine learning models can largely be divided into supervised learning and unsupervised learning depending on the type of the available data. In supervised learning, the data consists of sets of input and associated output known as a label. The goal of this algorithm is to predict output labels for unforeseen data by learning the relationship between the input and its label from the training data. When the available data does not contain labels, unsupervised learning can be used to cluster the unlabeled data by identifying hidden patterns.

There are mainly three steps in the construction of a machine learning system in materials science: data preparation, model selection, and model evaluation.[131, 171, 174] Since data is a key ingredient to train machine learning models, it is essential to collect, preprocess, and featurize the data as a first step. In materials research, the original data are generated by experiments and computational simulations. These data may contain incomplete, flawed, and noisy samples, and thus, they should be preprocessed to handle those outliers. Furthermore, it is essential to perform proper feature engineering tasks including feature extraction, feature selection, and feature construction to prepare datasets that are understandable and informative.[67] Finally, the processed data should be partitioned into a training dataset and a test dataset for model training and evaluation.

The second step in the machine learning pipeline is model selection. Based on a task of an application, it is critical to select an appropriate machine learning algorithm as each algorithm has its own usage and characteristics. Machine learning algorithms can typically be categorized into four tasks: regression, classification, clustering, and probability estimation. In materials research, regression, classification, and clustering algorithms are commonly used in material property prediction.[31] Probability estimation algorithms including generative and probabilistic models are typically used to generate new materials structures with desired properties. Therefore, with a clear understanding of a task, one should select the most suitable model from a set of available machine learning algorithms.

Lastly, it is critical to determine how well the model performs not only on the training data but also on unseen data. To measure the model performance, one should compare the model's predictions to the true labels using a proper evaluation metric. Common evaluation metrics are mean squared error (MSE) and mean absolute error (MAE) for regression tasks, and cross entropy, precision, and recall for classification tasks. These measurements guide the model training and help in designing optimal model architectures for the best outcomes.

1.3 Overview of Deep Learning in Materials Science

For decades, machine learning systems required a carefully designed feature extractor that uses human domain expertise to transform the raw data into a proper feature representation. Many traditional machine learning models in computational chemistry rely on several hand-crafted representations.[14, 17, 18, 180] These humanengineered representations encode interactions between atoms to provide local structural environments of individual atoms using an invariant set of atom-centered radial and angular symmetry functions. However, these features are not transferable from one system to another, and their applicability is also limited to systems containing up to three or four elements due to the complexity of configuration space that is exponentially growing as the number of elements increases. These limitations of human-designed features have motivated the adoption of deep learning in the computational chemistry field.

Deep learning is a class of machine learning algorithms that allows a model to be fed with raw data and to automatically learn feature representations needed for the task of interest.[6, 20, 71] With the ability of learning representations at multiple levels of abstraction, deep learning methods have outperformed traditional machine learning methods in image processing, speech recognition, and natural language processing tasks.[44, 220, 235]

In computational materials science fields, graphs have become a popular representation of the chemical structure. Graphs are a general language for describing entities with relations or interactions, e.g., molecules, mathematical objects, social networks, transportation paths, etc.[49, 170, 223, 233] A graph consists of two main components: a finite set of vertices (or nodes) and a finite set of directed or undirected edges, each connecting a pair of vertices. This special data structure offers a unified representation for materials by encoding fundamental properties of atoms into node features and encoding atom connections into edge features.[223] Using all the information about the molecular graph consisting of raw features, deep network models extract optimal representations for a task of interest which are also called embeddings. These embeddings are optimized via message passing layers where features gradually contain more information on the surrounding environment.[174] This framework is called graph neural network and has been widely used in materials science applications including property prediction for solid materials[223] and structure prediction for catalysts[233].

A more sophisticated deep learning framework that has not been well studied in the computational chemistry field is deep reinforcement learning. Deep reinforcement learning is a combination of reinforcement learning and deep learning in which a policymaker or an agent learns strategies to solve an optimization problem by interacting with an environment.[8] Unlike supervised or unsupervised learning algorithms that learn from fixed training datasets, a reinforcement learning algorithm is dynamically learned by adjusting actions based on continuous feedback from the environment to maximize a reward. DRL has been successfully used in various applications including control[127], autonomous driving[107], games[158], etc.

1.4 Thesis Objectives

In this thesis, we explore several interesting problems related to the surface and interface science. Each chapter of this thesis focuses on one different problem, with the goal of developing a framework and providing solutions and future guidance for the specific problem. We divide the thesis into two parts based on the type of problems and methods discussed in each part.

Before we dive into deep learning methods, Part 1 of this thesis demonstrates an

example of how the theoretical models and computational simulation methods can be used for understanding interfacial phenomena, and it will discuss what are the possible areas that can be improved by the use of deep learning techniques.

• Chapter 2 presents a combined molecular thermodynamic theory (MTT) and molecular dynamics (MD) simulations to predict non-ideal adsorption at a solid-solution interface. Simulation techniques for fluid-fluid interfaces are well developed, but these techniques cannot be directly applied to hard interfaces such as two-dimensional layered materials or nanoparticle surfaces. The framework proposed in this chapter provides atomistic insights into the non-ideal behavior of surfactants by capturing structural phases of the surfactants at the hard interface.

Part 2 of this thesis aims to develop a set of deep learning methods that solve various solid materials and heterogeneous catalyst design problems. There can be different machine learning solutions for those problems, but the deep learning methods presented in this part of the thesis are state-of-the-art and actively used for advancing the catalyst design process.

- Chapter 3 presents a graph neural network model to predict binding energies for high-throughput screening of inorganic catalysts. Binding energies can be used as descriptors to predict catalytic properties, but high-throughput calculations using density functional theory have been limited to simple catalytic systems due to its computational cost. The deep learning model discussed in this chapter could accelerate the screening and broaden the search space by providing binding energies at ab-initio level accuracy orders of magnitude faster.
- Chapter 4 presents a deep learning method that combines graph neural network and differentiable optimization into a single framework to directly predict the ground state structure of inorganic catalyst surfaces. Ground state structures

are a prerequisite for most computational materials science analyses since they reveal chemical and physical properties that are crucial for designing materials. This framework suggests a breakthrough to accelerate the identification of ground states of inorganic structures that was bottlenecked by the expensive computing cost of electronic structure methods.

• Chapter 5 presents a deep reinforcement learning framework for predicting kinetic pathways to surface reconstruction in a ternary alloy catalyst. Under realistic reaction conditions, catalyst surfaces are often reconstructed and their catalytic properties are also changed. The reinforcement agent explores possible surface reconstructions and transition states to identify the metastability of the alloy catalyst surface.

Part I Modeling and Simulation for interface structure

Chapter 2

Capturing Structural Transitions in Surfactant Adsorption Isotherms at Solid/Solution Interfaces

This work originally appeared as: Junwoong Yoon and Zachary W. Ulissi. Capturing structural transitions in surfactant adsorption isotherms at solid/solution interfaces. Langmuir 36, 3, 819-826 (2020). It has been edited to include the supporting information in-line.

2.1 Abstract

Although adsorption isotherms of surfactants are critical in determining the relationship between interfacial properties and structures of surfactants, providing quantitative predictions of the isotherms remains challenging. This is especially true for adsorption at hard interfaces such as on 2D layered materials or on nanoparticles where simulation techniques developed for fluid-fluid interfaces that dynamically change surface properties by adjusting unit cells do not apply. In this work, we predict nonideal adsorption at a solid-solution interface with a molecular thermodynamic theory (MTT) model that utilizes molecular dynamics (MD) simulations for the determination of free energy parameters in the MTT. Furthermore, the MD/MTT model provides atomistic insights into the non-ideal behavior of surfactants by capturing structural phases of the surfactants at the interface. Our approach captures structural transitions from the ideal state at low concentrations and then to the critical surface aggregation concentration (CSAC) and finally through the critical micelle concentration (CMC). We validate our model against the original MTT model by comparing predicted adsorption isotherms of a simplified surfactant system from both approaches. We further substantiate the applicability of our model in complex systems by providing adsorption isotherms in an aqueous sodium dodecyl sulfate(SDS)graphene system, in good agreement with experimental observations of the CSAC for the same system.

2.2 Introduction

The adsorption of surfactant at interfaces has received considerable attention in both academic research and industrial production because it can improve applications of interfacial phenomena in biotechnology, energy science, catalysis, and advanced materials. [93, 188, 207] Interface modification via surfactant adsorption alters interfacial properties such as colloidal stability, interlayer distance, and surface tension and energy, which affect the performance of the interfacial applications.[83, 93, 140, 181, 186] Surfactants adsorb to the interfacial surfaces and form dense layers or aggregate structures based on the surface concentration Γ , the adsorbed amount of surfactants per unit area at the interface.[83, 93, 181] While understanding the struc-


Figure 2-1: Atomistic simulation snapshots for the adsorption of aqueous SDS surfactants on the graphene surface. The SDS surfactants direct the formation of structured aggregates from monolayer at a low surface concentration to surface micelle at high surface concentration. The hydrophilic surfactant headgroups are composed of sulfur (yellow) and oxygen (red) atoms, and the hydrophobic tailgroups are composed of carbon (green) and hydrogen (white) atoms. Water molecules are not shown for clarity.

tural behavior of surfactants at the interface is of great importance in determining the interfacial properties, the concentration-dependence of these properties is usually studied through qualitative molecular simulations.

The structures of surfactant aggregates on interfaces are closely related to the adsorption isotherm $\Gamma(c)$, Γ at a given the bulk concentration c.[181] Figure 2-1 shows the current experimental understanding of the adsorption of aqueous sodium dodecyl sulfate (SDS) surfactants at graphene surface.[93, 208] This includes structural transitions due to different levels of the surface concentration. Although the adsorption isotherm is critical in determining the relationship between the interfacial properties and structural behavior of surfactant, providing quantitative predictions of the adsorption isotherm remains challenging. Therefore, developing appropriate techniques for the atomic-scale studies of surfactant adsorption at interfaces is a prerequisite for the design of surfactant molecules for which the adsorption isotherms and structural behavior are unknown.

Advanced experimental techniques such as neutron reflectivity[122, 162] and transmission electron microscopy[117] can probe the micellar structures of adsorbed surfactant molecules along adsorption isotherms. However, these techniques have several limitations including difficult sample preparation, high cost, and insufficient resolution to deliver atomistic-level details of surfactant structure.[237] The reported adsorption isotherms from experimental measurements are often inconsistent with one another or produce large variations due to the complex nature of experimental settings and measurements.[28, 181]

Molecular Dynamics (MD) simulations have been used for the quantitative prediction of properties of molecular systems while delivering atomistic-level insights into the molecular structure and dynamics. When using MD simulations to investigate the effect of surfactant adsorption on interfacial properties, proper modeling of surfactant surface concentration is required.[37, 115, 206] Unfortunately, conventional MD simulations are currently limited by simulation time scales that are much shorter than the time required for the surfactant-bulk interface to reach equilibrium.[188] As such, there is a practical motivation for developing techniques to predict both the surfactant adsorption isotherm and the structure of surfactant molecules using MD simulations.

Alternative approaches such as Monte Carlo simulations can also be considered. In Monte Carlo simulations, molecules are inserted and deleted to calculate equilibrium chemical potential using two simulation configurations representing the surface and the bulk phases separately.[147, 231] Yoo et al. [231] have developed a discrete fractional component Monte Carlo method in the Gibbs ensemble framework, which provides adsorption isotherms for surfactants in an air/water interface. This approach works well in the dilute concentration but an extension of this work to high concentration regimes requires better modeling techniques that can capture accurate surfactant-surfactant interactions in the bulk. This method is also computationally unfeasible to study solid/liquid systems where the unit cell is not adjustable and the molecules are modeled on atomistic length scales.

Molecular Thermodynamic Theory (MTT) can be a useful tool for calculating adsorption isotherms because it is capable of capturing both the ideal and non-ideal free energy contributions of surfactants. Sresht et al. [186] have developed MTT models combined with MD simulations to provide a link between the surface concentration and bulk concentration of surfactant to make quantitative predictions of the interfacial tension of an air-water system without conducting any experimental measurements. This model uses MD simulations to determine the MTT parameters for the calculation of the adsorption isotherms based on several underlying assumptions that simplify the surfactant system: surfactant molecules in this system only form a single monolayer at the interface, and individual surfactant molecules in the monolayer are modeled as hard disks with short-ranged attractive interactions. This approach can reproduce experimental adsorption isotherms for polyethylene glycol in an air/water system, but the required assumptions may not be applicable to more complex surfactant systems especially when non-ideal surfactant interactions and aggregates are present. Another MD/MTT framework developed by Jaishankar et al. [104] could predict adsorption isotherms for surfactant on solid surfaces that also agree well with experimental data. However, this model is also only able to consider the formation of a monolayer of surfactants on solid surfaces while ignoring micelle formation or self-aggregation on surfaces.

In this work, we predict the surface-bulk equilibrium concentrations by developing a new MD/MTT model. We choose MD/MTT model originally developed by Sresht et al. [186] as a benchmark. Similar to the benchmark, our model determines all required MTT parameters directly from MD simulations. We relax the simplifying assumptions in the benchmark that model adsorbed surfactants as hard disks and limit surfactant aggregates to a monolayer by developing an all-atomistic MD simulation technique that allows complex non-ideal surfactant interactions and aggregates. Our MD/MTT model predicts the adsorption isotherms for a range of bulk concentrations below the critical micelle concentration (CMC) while capturing structural transitions. We demonstrate that our MD/MTT model can reproduce the adsorption isotherms produced by the benchmark MD/MTT model and further provide adsorption isotherms in aqueous SDS-graphene systems while capturing structural transitions. Finally, we verify that our model is able to predict the consistent adsorption isotherms for the SDS-graphene systems with a different number of SDS molecules.

2.3 Methods

2.3.1 Molecular Thermodynamic Theory (MTT)

In the benchmark model developed by Sresht et al. [186], an expression for the chemical potential μ^{σ} of a surfactant molecule at the interface in the air/water system is constructed as follows,

$$\frac{\mu^{\sigma}}{k_B T} = \frac{\mu^{\sigma, id}}{k_B T} + -\ln(1-\eta) + \frac{3\eta - 2\eta^2}{(1-\eta)^2} + 2B\Gamma$$
(2.1)

based on the following assumptions: (i) there is a single monolayer of surfactant molecules at the interface, (ii) individual surfactant molecules in the monolayer is modeled as hard disks, (iii) and the surfactant tails stick out into the air phase even at low concentrations. Here, $\mu^{\sigma,id}$ is the ideal chemical potential of a surfactant molecule at the surface; *B* is a pairwise coefficient for van der Waals attraction between adsorbed surfactants; Γ is the surface concentration of the adsorbed surfactant molecules; and $\eta = \Gamma A_c$ is a packing fraction (ratio of the cross-sectional area A_c of an adsorbed surfactant to the available surface area per surfactant) of the surfactant molecules adsorbed on the interface. Contrary to this original expression for



Figure 2-2: Schematic diagram of the MD simulation model for the determination of $G(\Gamma)$, the difference between chemical potential of an SDS at initial state μ^{σ,Γ_0} and the chemical potential at final state $\mu^{\sigma,\Gamma}$. The initial state is considered as being at an ideal condition where the non-ideal interactions among the SDS surfactants are absent $(f(\Gamma_0) = 0)$. The final state is where the surface concentration is high in order to result in the aggregate formation of SDS surfactants. Water molecules are not shown for clarity.

the chemical potential of a surfactant molecule at the surface, we define the μ^{σ} as follows without those assumptions listed above,

$$\mu^{\sigma} = \mu^{\sigma, id} + f(\Gamma) \tag{2.2}$$

where $f(\Gamma)$ represents the non-ideal free energy contribution of the adsorbed surfactant as a function of the surface concentration. $f(\Gamma)$ provides a more general expression of the non-ideal free energy contribution because it does not involve any assumptions used to develop the analytic expression for μ^{σ} in the benchmark model, eq 1.

As in the benchmark model, we define the reference state as the surfactant system at infinite dilution where the surface pressure is negligible (~ 0 mN/m). We can calculate the ideal chemical potential of an adsorbed surfactant $\mu^{\sigma,id}$ using the reference-state chemical potential of the surfactant $\mu^{\sigma,0}$ and an entropic term as follows,

$$\mu^{\sigma,id} = \mu^{\sigma,0} + k_B T \ln(\eta).$$
(2.3)

Assuming that there is no significant interactions among surfactant molecules in dilute bulk solution (below the CMC), the chemical potential of the surfactant in the bulk phase can be expressed as

$$\mu^{b} = \mu^{b,0} + k_B T \ln(x), \qquad (2.4)$$

where $\mu^{b,0}$ is the surfactant bulk chemical potential at the reference state at infinite dilution and x represents the bulk mole fraction of the surfactants. At thermodynamic equilibrium where $\mu^{\sigma} = \mu^{b}$, the bulk mole fraction can be determined as a function of the surfactant surface concentration as follows,

$$\ln(x) = \frac{\Delta\mu^0}{k_B T} + \frac{G'(\Gamma)}{k_B T} + \ln(\eta_0).$$
 (2.5)

To solve this equilibrium equation, two parameters need to be determined using MD simulations: (1) $\Delta \mu^0 = \mu^{\sigma,0} - \mu^{b,0}$, the difference between the reference state

chemical potential of a surfactant molecule at the interface $(\mu^{\sigma,0})$ and the reference state chemical potential of the surfactant molecule in the bulk solution phase $(\mu^{b,0})$; and (2) $G(\Gamma)$, the change in the chemical potential of a surfactant at initial (ideal) state and the potential at final state as described in Figure 2-2. $G(\Gamma)$ represents the free energy required to concentrate surfactant molecules at the initial state where the non-ideal free energy contribution is negligible into the final state where the surface concentration is high enough to induce non-ideal surfactant interactions and aggregates. To verify whether the initial state is truly ideal, we estimate the nonideal free energy contribution of the surfactant at the interface at the initial state $f(\Gamma_0)$. Then we define a corrected free energy parameter $G'(\Gamma) = G(\Gamma) - f(\Gamma_0)$ that relaxes the non-ideal contribution that can possibly be present in the initial state. By simultaneously solving $G(\Gamma)$ and **eq 5** using determined $\Delta\mu^0$, $f(\Gamma_0)$, and η_0 (packing fraction at the initial state as shown in Figure 2), we can predict the equilibrium surfactant surface concentrations for a range of bulk mole fraction.

2.3.2 Determining parameters for validation of the MTT model: Lennard-Jones sphere surfactant model on a graphene sheet

The NAMD simulation package was employed for the determination of parameters in the MTT model. [165] All simulation systems for the validation of the MTT model were constructed with simulation boxes of dimensions $5.1 \times 5.1 \times 5.0$ nm and each contains a 5.0×5.0 nm graphene sheet and Lennard-Jones sphere surfactants that represent hard disk surfactants in the benchmark model. Carbon atoms in the rigid graphene were represented using the all-atomistic CHARMM36 force field [24] and the Lennard-Jones sphere surfactants were modeled using CHARMM22 force field[138, 139]. Simulations were performed under the canonical NVT ensemble while maintaining a constant temperature at 300K using Langevin dynamics[55].

Determining $\Delta \mu^0$

To construct a reference state at infinite dilution, a single Lennard-Jones sphere surfactant was placed 0.3 nm above the graphene sheet. The Adaptive Biasing Force (ABF)[43, 84] method was then applied to a collective variable, the distance between the surfactant and the graphene, to calculate the potentials of mean force (PMF) acting between the surfactant and the graphene as a function of their separation. The ABF method allows the free energy surface along the reaction coordinate of interest to be estimated by continuously updating a running average of the force acting along the coordinate. The resulting PMF from this ABF simulation represents the free energy changes of the surfactant as a function of the height of the Lennard-Jones sphere surfactant from the graphene surface. A time step of 1.0 fs was used throughout the simulation for 10.0 ns. Figure 2-3(a) shows that the surfactant has the lowest free energy at 0.33 nm above the graphene and the gradient of the free energy becomes close to zero at a height greater than about 1.0 nm from the graphene surface. The free energy of adsorption of an SDS at the reference state $\Delta \mu^0$ was estimated at -4.8 kcal/mol based on the difference between the free energy of the surfactant at the graphene interface where the surfactant has the lowest free energy and the free energy of the surfactant in the bulk phase.

Determining $G(\Gamma)$

25 Lennard-Jones sphere surfactants were placed uniformly on the graphene sheet at a height of 0.3 nm to create an initial state with low surface concentration, approximately 1.4 nm². For simplification purpose, we assumed non-ideal free energy contribution of the adsorbed Lennard-Jones sphere surfactant at the initial state are negligible $(f(\Gamma_0) \approx 0)$. An external harmonic position restraint $k(x - x_0)^p$ was



Figure 2-3: (a) Free energy change for a Lennard-Jones sphere surfactant as a function of the height above the graphene surface obtained using ABF simulation. $\Delta \mu^0$ is determined as the difference in the free energy of the surfactant when it is in the bulk solution (~1.2 nm) and when it is at the interface (~0.33 nm). (b) Free energy required to concentrate surfactants in the initial state as a function of the surface concentration obtained using our simulation model (scatter) and the benchmark (line).

applied to the surfactants to keep them on the graphene surface. This constraint reproduces thermodynamic equilibrium by preventing the surfactants from leaving the graphene surface due to high surface concentration. Force constant k was set to 100 $kcal/mol/nm^2$ and the exponent p was set to 2. These position restraint parameters were set to form a single monolayer even at the high surface concentration in order to mimic the assumptions in the benchmark model. By applying the ABF to the radius of gyration of the surfactants which is a collective variable in this simulation, the PMF as a function of the collective variable could be estimated. The resulting PMF represents the free energy associated with concentrating the fixed number of Lennard-Jones sphere surfactants at the initial low surface concentration state into the final state with higher surface concentration. The simulation was run for 10 ns with a time step of 1.0 fs. The radius of gyration and the fixed number of the surfactants were used to directly compute the surface concentration. The obtained free energy change $G(\Gamma)$ as a function of the surface concentration was then compared with analytically calculated free energy profile using the benchmark model as shown in Figure 2-3(b).

2.3.3 Determining parameters for aqueous SDS/graphene interface

All simulations used for aqueous SDS/graphene system were constructed with simulation boxes of dimensions $11.5 \text{ nm} \times 11.5 \text{ nm} \times 7.4 \text{ nm}$ that contain a $10.0 \text{ nm} \times 10.0 \text{ nm}$ graphene sheet and SDS molecules. SDS molecules were represented using all-atomistic CHARMM22 and CHARM27 force field [54, 56] and graphene was constructed using CHARMM22 force field [138, 139]. The simulation box was then filled with water molecules using the three-site (TIP3P) water model which is the current default water model in the NAMD simulation package. Bonds involving hydrogen in water molecules are set to be rigid. Langevin dynamics and Nose-Hoover Langevin

piston [55, 142] were used to maintain a constant temperature at 300 K and constant pressure at 1 atm throughout the simulations.

Determining $\Delta \mu^0$

To construct the configuration for the reference state at infinite dilution, a single SDS molecule was placed in bulk solution at 1.1 nm above the center of the graphene sheet (distance between the center of mass of SDS and the graphene). The system was then filled with $\sim 17,000$ water molecules to create aqueous SDS/graphene system. An external harmonic potential with a force constant of $2,000 \text{ kcal/mol/nm}^2$ was applied to a collective variable, distance between the SDS and graphene, to initially move the SDS to a height of 2.0 nm and then pull it down toward the graphene surface for 1 ns with a time step of 2 fs. Simulation snapshots that contain positions of atoms in the system corresponding to 16 evenly spaced heights of the SDS between 0.4 and 2.0 nm above the graphene were extracted and used as the initial configurations for 16 biased-sampling simulations. During each sampling simulation, the SDS was restrained at a height corresponding to its initial configuration using a harmonic potential with a force constant of $1,000 \text{ kcal/mol/nm}^2$ for 5 ns with a time step of 1 fs. The time-series trajectory data from all the 16 biased-simulations were then used to calculate the free energy change of the SDS as a function of the distance between the SDS and graphene using the weighted histogram analysis method (WHAM) [76, 118, 119]. Figure 2-4(a) shows the free energy of adsorption of an SDS at the reference state $\Delta \mu^0$ is estimated at -10.8 kcal/mol and the SDS has the lowest free energy at 0.39 nm above the graphene sheet.

Determining $G(\Gamma)$

To construct the initial state, 30 SDS molecules were placed uniformly on the graphene sheet at a height of 1.1 nm, followed by solvation of the simulation box with $\sim 50,000$



Figure 2-4: (a) Free energy change for SDS surfactant as a function of the height above the graphene surface obtained by analyzing the histograms of the 16 sampling simulations using WHAM. $\Delta \mu^0$ is calculated as the difference in the free energy of the surfactant when it is in the bulk solution and when it is at the interface. (b) Free energy required to concentrate the SDS surfactants in the initial state as a function of the surface concentration obtained by analyzing the histograms of 28 sampling simulations using WHAM.

water molecules. To study the correct structure of SDS hemi-micelle on the graphene surface, we assumed that around 30 SDS molecules need to be placed on the surface because the hemi-micelle would require about a half of the average SDS micelle aggregation number at the CMC in water, which is 60.[195] An external harmonic position restraint was applied to the sulfur in the SDS headgroups to avoid the desorption of the SDS molecules due to high surface concentrations. The force constant k for the position restraint was set to 1.0 kcal/mol/nm² and the exponent p was set to 2. The simulation system was then equilibrated to ensure all atoms were in an energetically stable configuration where the hydrocarbon chains of the SDS molecules were lying down on the graphene surface at a fixed low surface concentration (approximately 0.4 nm^2) under the NPT ensemble for 10 ns with a time step of 2 fs. This relaxed configuration was used as the initial state where non-ideal surfactant interactions are insignificant as shown in Figure 2-2 for the following sampling simulations for the determination of $G(\Gamma)$. The average distance between SDS molecules in this initial state was estimated at 6.4 nm. We also estimated non-ideal free energy contribution of the adsorbed SDS at the initial state $f(\Gamma_0)$ as the difference between total interactions at the ideal state and the interactions at the initial state An external harmonic potential with a force constant of 2,000 kcal/mol/nm² was then applied to a collective variable, a radius of gyration of SDS molecules, to pack the relaxed SDS molecules into smaller surface area and to induce the formation of SDS aggregates by decreasing the radius of gyration for 10.5 ns with a time step of 2 fs. During this packing simulation, the radius of gyration of the SDS molecules was used to directly measure the surface area and surface concentration with the fixed number of SDS molecules. Simulation configurations corresponding to 28 evenly spaced radius of gyration of the SDS molecules between and including 1.3 nm and 4.0 nm were selected and used as initial configurations for 28 biased-sampling simulations. In each sampling simulation, the SDS molecules were restrained at each radius of gyration



Figure 2-5: The surface concentration of surfactants as a function of the bulk mole fraction in the simplified surfactant system obtained using our MD/MTT model (scatter) and the benchmark.

using a harmonic potential with a force constant of 500 kcal/mol/nm² for 7 ns with a time step of 1 fs. The free energy change $G(\Gamma)$ associated with varying the surface concentration of the SDS molecules was then calculated by applying WHAM to the time series trajectory data from the 28 sampling simulations as shown in Figure 2-4(b).

2.4 Results and discussion

2.4.1 Validation of the model

We first check the validity of our model by comparing the two adsorption isotherms calculated using the benchmark model and our model for the simplified surfactant system. As we described before, only a single monolayer of Lennard-Jones sphere surfactants is assumed to be present on a graphene surface in the simplified system. $\Delta \mu^0$ is estimated at -4.8 kcal/mol based on the free energy profile shown in Figure 2-3(a) and the resulting $G(\Gamma)$ from the simulation model described in Figure 2-2 is shown in Figure 2-3(b). By simultaneously solving **eq 5** with the determined parameters, $\Delta \mu^0$ and $G(\Gamma)$, we generate adsorption isotherm in the system. Figure 2-5 compares between the adsorption Γ as a function of the bulk mole fraction xfor the simplified surfactant determined by our model and the benchmark model. Our MD/MTT model is able to reproduce the result from the analytic equilibrium equation developed by Sresht et al. [186] in the simplified system. The ability to reproduce the analytic equation **eq 1** also demonstrates the applicability of our model to more complex systems because our model can directly provide $G(\Gamma)$ using MD simulations without considering any assumptions used in the benchmark model that prohibit non-ideal complex surfactant interactions.

2.4.2 Real surfactant system - aqueous SDS/graphene interface

Adsorption isotherm

 $\Delta \mu^0$ for an SDS molecule in an aqueous SDS/graphene system is estimated at -10.8 kcal/mol from Figure 2-4(a), and $f(\Gamma_0)$ is estimated at -0.10 kcal/mol. $G(\Gamma)$ is determined using 28 sampling simulations and plotted in Figure 2-4(b). By simultaneously solving **eq 5**, we generate the adsorption isotherm of SDS molecules on the graphene sheet. Blue scatters in a joint plot in Figure 2-7 shows the SDS surface concentration Γ as a function of the bulk mole fraction x for the system with 30 SDS molecules. Two red vertical dashed lines indicate the corresponding bulk mole fraction of the critical surface aggregation concentration (CSAC)[214] (left) and CMC



Figure 2-6: Violin plots showing the distribution of θ on a log scale as a function of the surface concentration. θ is introduced to measure how much the headgroup is lifted from the graphene surface. Each violin is fitted to a mixture of Gaussian density functions. A peak of each Gaussian density function in each violin is labeled as peak 1, 2, and 3 based on the corresponding $\log(\theta)$. Example SDS surfactant structure at each peak is shown at the right end of the violin plots. The table shows the average $\log(\theta)$ and θ at each peak.

(right). CSAC is defined as the concentration of the surfactant required for the formation and adsorption of the hemi-micelles on the surface. Using our MD/MTT model, the surface concentration at the CSAS is estimated at 2.7 nm⁻² which is equivalent to 0.37 nm² surface area per molecule. The predicted surface area per molecule at CSAC agrees well with both the experimental data [214](0.42 nm²) and MD simulations[206] (0.40 nm²). The surface concentration and the surface area per molecule at CMC were also estimated at 3.8 nm² and 0.26 nm², respectively.

Structural behavior

To study structural behavior of SDS molecules at the water/graphene interface, we measure angle θ between the SDS molecule and the graphene to see how much the headgroup is lifted as described in a cartoon in Figure 2-6 at each time step



Figure 2-7: A joint plot that shows the SDS surface concentrations corresponding to each aqueous SDS-graphene system with 30 SDS (blue) and 36 SDS (orange), and the average number of SDS molecules (black) found in each peak (defined in Figure 6) as a function of the bulk mole fraction. The two red vertical dashed lines represent the corresponding bulk mole fractions of the CSAC (left) and the CMC (right). Four structural phase regions are highlighted with different colors and labeled as phases a, b, c, and d. The simulation snapshots for the structure of the SDS molecules in each phase for the system with 30 SDS molecules are shown above the joint plot. Water molecules are not shown for clarity.

throughout all sampling simulations used for the determination of $G(\Gamma)$. θ provides information about the structural properties of SDS molecules. A small θ that is close to zero indicates that the SDS is horizontally lying down, and a large θ indicates that the SDS is tilted by the headgroup due to a higher concentration. In Figure 2-6, we also generate a violin plot that shows the distribution of θ on a log scale at different surface concentrations. Each violin represents a rotated kernel density estimation of the distribution of θ on a log scale. Here, we plot and analyze the distributions of θ on a log scale because the actual θ is too broadly distributed to be fitted to Gaussian distributions. For example, it is hard to distinguish 1.86° from 10.02° on a normal scale where the range of θ is from 0° to 90° . However, we could differentiate between those two angles by plotting them on a log scale where they are represented as 0.27 and 1.00 respectively when the range of the $\log(\theta)$ is from -1 to 2. We then fit each violin to a mixture of Gaussian density functions to extract statistical information about the structure of SDS molecules. The motivation for using a mixture of Gaussian density functions comes from the idea that it can be used to approximate any continuous distributions with multiple local peaks.

We find three peaks at 0.27 ($\theta \approx 1.86^{\circ}$), 1.00 ($\theta \approx 10.02^{\circ}$), and 1.61 ($\theta \approx 40.57^{\circ}$) as shown in the table in Figure 2-6(b). Peak 1 and 2 are observed at fairly constant angles throughout the entire range of the surface concentration while peak 3 is only observed at relatively high surface concentrations (above $\sim 3 \text{ nm}^2$). We then estimate the average number of SDS molecules at each peak where they have the corresponding θ using the total number of SDS molecules in the simulations and the estimated relative area under each peak. Figure 2-7 shows the average number of SDS molecules at each peak using black lines and the surface concentrations simultaneously as a function of the bulk mole fraction.

Based on the change in the average number of SDS molecules at each peak, we could predict that there are four structural phases during the adsorption of SDS on

the graphene surface. Those four phases (phase (a), (b), (c) and (d)) are shown as filled regions in Figure 2-7, and the corresponding structures of SDS molecules (structure (a), (b), (c) and (d)) are visualized using visual molecular dynamics (VMD)[97] and shown above the plot in Figure 7. These structural transitions agree well with the experimentally developed SDS adsorption model at the graphene sheet [93, 208].

At low bulk mole fraction (less than $x \approx 7.3 \times 10^{-7}$), most SDS molecules are found in peak 1 where they have θ values very close to zero. Initially, SDS molecules form a monolayer while they are lying down on the graphene surface in order to minimize contact between hydrophobic molecules (hydrocarbon chain of SDS and graphene) and water molecules. We observe that the monolayer becomes denser as the bulk mole fraction increases and then becomes fully packed at $x \approx 7.3 \times 10^{-7}$, as illustrated in structure (a). In the fully packed monolayer, headgroups of most SDS molecules are pointing outward because of their hydrophilic nature.

At higher bulk mole fraction as in phase (b), SDS molecules in the fully packed monolayer with headgroups that are pointing inward start to thrust their headgroups toward the aqueous solutions while keeping the tailgroups on the graphene surface in attempt to maintain the exposure of the headgroups to water at higher concentration as shown in structure (b). This behavior causes an increase in the average number of SDS molecules at a peak 2 in phase (b).

In the range of bulk mole fraction in phase (c), more SDS molecules are found in peak 2 compared to peak 1. As the SDS molecules thrust their headgroups harder due to higher concentration, we observe that some of them in peak 2 is now found in peak 3 where they have larger θ . We consider the structure of SDS when the peak 3 first appears is a hemi-micelle on the surface as illustrated in structure (d). There may be a degree of uncertainty in this observation, but we could validate the formation of surface hemi-micelle at the surface concentration (2.7 nm²) where peak 3 first appears. The corresponding bulk mole fraction ($x \approx 5.7 \times 10^{-5}$) is close to the experimentally determined CSACs (x $\approx 3.7 \times 10^{-5}$ and 5.1×10^{-5}) in the aqueous SDS/graphene system [208, 214].

In phase (d) where the bulk mole fraction is close to CMC, we observe the similar hemi-micelle structure as we found at CSAC, but it is denser due to the higher concentration. We note that the adsorption of SDS near the CMC—where the packing and adsorption of surfactants are limited in reality—may not be as accurate as we expected, because there may be significant interactions among SDS molecules in the bulk solution at the CMC as opposed to our assumptions in eq 4.

Generalization

To show how dependent the adsorption isotherm and the structural behavior are on the number of SDS molecules in the system, we perform the same calculations for determination of MTT parameters, $G(\Gamma)$, $f(\Gamma_0)$, and $\ln(\eta_0)$ in an aqueous SDS/graphene system that contains 36 SDS molecules. This is an important verification test for our MD/MTT framework to show it can extrapolate an adsorption isotherm spanning a wide range of concentrations in the bulk and on the surface. As shown in the plot in Figure 7, our MD/MTT model is able to generate consistent adsorption isotherms for both aqueous SDS-graphene system with 30 SDS molecules (blue) and the same system with 36 SDS molecules (orange). Furthermore, we are also able to observe that the adsorption isotherm for the system with 36 SDS molecules follows the structural phases and transitions previously found in the system with 30 SDS molecules. Details of structural phases and calculation of the adsorption isotherm for the system with 36 SDS molecules can be found in Figure 2-8.



Figure 2-8: (1) $G(\Gamma)$ for both systems with 30 SDS and 36 SDS. The free energy profile for the system with 36 SDS is slightly off at low surface concentrations, but it can be corrected using $f(\Gamma_0)$ and it shows the same trend as the one for the system with 30 SDS. (2) (top) Simulation snapshots for the surfactant structures at each structural phase (phases a, b, c, and d) in each system. This shows both systems have the same structural phases in the same range of bulk mole fractions and follow the same structural transitions. (bottom) adsorption isotherms obtained from MTT parameters calculated for each system with 30 SDS and 36 SDS. Red vertical dashed lines represent CSAC (left) and CMC (right).

2.5 Conclusions

In this work, we presented a new framework that combines MTT and MD simulations to predict the surface-bulk equilibrium concentrations of surfactant, thereby overcoming the computational limitations of atomistic MD simulations for the determination of adsorption isotherms. The MD/MTT framework connects the bulk concentrations of surfactant with the surface concentrations at equilibrium while the MD simulations determine the thermodynamic parameters ($\Delta \mu^o$ and $G(\Gamma)$) presented in the MTT. The simulation modeling we developed for the direct calculation of the free energy contributions of the surfactant molecules allows complex surfactant systems to be studied without simplifying the structure and properties of the systems. This simulation technique also enables our framework to understand the relationship between the adsorption isotherms and structural behavior of surfactants at interfaces. We showed that our framework could capture structural transitions during SDS adsorption at a graphene surface. The transitions agree well with the experimental observations for the same system while providing the adsorption isotherms. We demonstrate that our MD/MTT framework is robust to a number of surfactants by predicting consistent adsorption isotherms and surfactant structural phases for different numbers of SDS surfactants in the aqueous SDS-graphene system. This MD/MTT framework is not limited to one system, and can also be applied to new surfactant systems and determine the thermodynamic parameters to provide adsorption isotherms. Thus the MD/MTT framework can serve as a better tool for the design of surfactant molecules that require an understanding of the effect of the adsorption isotherms and structural behavior of the molecules on interfacial properties.

Part II Deep Learning for Catalyst Design

Chapter 3

Convolutional Neural Network of Atomic Surface Structures To Predict Binding Energies for High-Throughput Screening of Catalysts

This work originally appeared as: Seoin Back^{*}, Junwoong Yoon^{*}, Nianhan Tian, Wen Zhong, Kevin Tran, and Zachary W. Ulissi. Convolutional neural network of atomic surface structures to predict binding energies for high-throughput screening of catalysts. J. Phys. Chem. Lett. 10, 4401 (2019). It has been edited to include the supporting information in-line.

3.1 Abstract

High-throughput screening of catalysts can be performed using density functional theory calculations to predict catalytic properties, often correlated with adsorbate binding energies. However, more complete investigations would require an order of 2 more calculations compared to the current approach, making the computational cost a bottleneck. Recently developed machine-learning methods have been demonstrated to predict these properties from hand-crafted features but have struggled to scale to large composition spaces or complex active sites. Here, we present an application of deep-learning convolutional neural network of atomic surface structures using atomic and Voronoi polyhedra-based neighbor information. The model effectively learns the most important surface features to predict binding energies. Our method predicts CO and H binding energies after training with 12,000 data for each adsorbate with a mean absolute error of 0.15 eV for diverse chemical space. Our method is also capable of creating saliency maps that determine atomic contributions to binding energies.

3.2 Introduction

Understanding atomic and molecular interactions with solid surfaces is a basis for predicting catalytic properties in heterogeneous catalysis. Density functional theory (DFT) calculations have played an important role in understanding these interactions and estimating binding strengths accurately with respect to the corresponding experimental measurements (MAE = ~ 0.2 eV).[81, 89] Binding energies can be used as descriptors to predict catalytic properties such as electrochemical onset potentials[141, 163], turn over frequencies[185], product selectivity[232] based on scaling relations between binding energies of reaction intermediates,[105, 141, 210] and Brønsted-Evans-Polanyi relation between reaction energies and kinetic barriers.[26, 210]

In the past decades, high-throughput screenings of catalysts have been performed mainly using simple metal/alloy systems, and the most stable and second most stable facets such as FCC (111) and (100) have been used due to their low surface energies. [72, 148] In addition, one descriptor has been usually used to predict catalytic activities—e.g., H, OH and CO binding energies to predict theoretical overpotentials of H₂ evolution/oxidation reactions, O₂ reduction reaction and CO₂ reduction reaction, respectively [72, 73, 74, 82] To more completely study catalytic properties of materials, however, we need to reconsider those two assumptions. First, the most stable facet is not always most active. For example, concave Au nanoparticles consisting of a large portion of high index facets showed much higher catalytic activity and selectivity compared to thin film consisting of stable facets. [120] Therefore, it is essential to consider several possible facets and a number of unique active sites on those facets, 205 increasing the computational cost by a factor of 20, assuming we consider 5 facets and 4 possible active sites. Second, one may want to extend a chemical space to explore, but it has been frequently observed that the scaling relations do not hold anymore for systems other than metal/alloy systems. Examples include single metal atom embedded in two-dimensional materials, p-block atom embedded in metals, and metal dichalcogenides. [9, 34, 61, 128] Therefore, we need to consider all reaction intermediates, increasing the computational cost by a factor of $5 \sim 10$ [10, 163] Altogether, the required computational cost becomes an order of 2 more expensive than the current high-throughput screenings, making the computational cost of DFT a main bottleneck. Thus, it is essential to develop strategies to alleviate the computational cost issue.

Recent developments of machine-learning (ML) techniques and their applications to catalysis have suggested that the ML could substantially facilitate the highthroughput screenings. Based on ML predicted binding energies, one could reduce the number of DFT calculations by excluding unpromising candidates [123, 125, 205]. For example, using fingerprints based on selected chemical and physical features of active sites and their neighbor information, many artificial neural network (ANN) models have been successful in achieving an acceptable accuracy (MAE = $\sim 0.2 \text{ eV}$) with respect to the DFT calculated binding energies. Unfortunately, many models are based on electronic features of active surface atoms, such as d-band characteristics, requiring *ab-initio* level computational cost to prepare them. Therefore, ML models that use readily available data and that achieve high accuracy with the minimal number of training DFT data are strongly desirable to pursue an efficient catalyst search over the broad chemical space.

Recently, Xie and Grossman applied a deep-learning convolutional neural network (CNN) on top of a graph representation of bulk crystals to predict various properties [36, 226] and reached the MAEs of DFT calculations and experiments $(0.10 \sim 0.15 \text{ eV})$ after training with $\sim 30,000$ DFT calculated data. The graph representation of the crystals includes information of atomic features and interatomic distances, and iterative convolutions extract neighbor information to update atomic feature vector. Note that none of expensive electronic structure information was used during the process; only the crystal structures and basic atomic features were used.

In this work, we use a modified form of this CNN method to predict CO and H binding energies on diverse surfaces of pure metals, metal alloys, and intermetallic surfaces. In our method, we collect the graph connections based on Voronoi polyhedra to take into account additional adsorbate atoms. We report state-of-the-art MAEs of 0.15 eV for both CO and H binding energy predictions using 12,000 training data for each adsorbate across a far larger collection of surface composition and structure than previously possible. We further show that our method can be used to create saliency maps to determine atomic contributions to binding energies and

to automatically detect failed DFT calculations.

3.3 Representation of Surface Structures

We modified the bulk crystal prediction code to better represent surface features. In the previous effort of the CNN applications to bulk property predictions, the authors encoded atomic information and interatomic bonding interactions into node and edge vectors, respectively [226], where the node vectors include the basic atomic information and the edge vectors contain atomic indices of nearest neighbor atoms and their bonding distances from the center atom. The main difference between bulk structures and our atomic surface structures is that the latter contain adsorbate atoms, where their initial positions are arbitrarily set by users. Since a goal of our method is to predict binding energies using initial structures as inputs, the arbitrariness of the initial adsorbate position should be addressed. We note that there is no such arbitrariness in the graph representation of bulk crystals, since atomic positions change marginally during a cell relaxation of already known structures and the final relaxed structures were used for training. [226] To resolve this arbitrariness, we modified the code so that we can incorporate neighbor information based on Voronoi polyhedra as implemented in Pymatgen[157] as "VoronoiConnectivity" (Figure 3.3b). Voronoi polyhedra, also known as Wigner-Seitz cell[216], of atoms in surface structures could provide local environment information as a solid angle. Each atomic polyhedron encompasses space in which distance to the center atom is less than or equal to the distance to other atoms. Solid angles between the plane shared by polyhedra of two adjacent atoms and center atoms are calculated, and values normalized to the maximum solid angle that belongs to the center atom were used to represent local environments of the center atom. We note that larger solid angles correspond to stronger interaction between the center and the neighbor atoms. To prepare atomic

feature vectors, we used the identical tabulated atomic information as in the original code.[226]

In some cases, the binding site changes from the initially set site to a more stable site during structure relaxation—e.g., from a top site to a bridge site. Since binding site changes could affect a prediction performance, we tested this effect by including a connectivity distance information from adsorbate atoms to all atoms in the surface structure taken from the final relaxed structures (Figure 3.3). We compared the prediction performance of training with initial structures only, final structures only, and initial structures with the connectivity distance information of adsorbates obtained from the final structures. Obviously, the prediction performance of training with final structures was best (MAE = 0.13 eV) and that with initial structures and the adsorbate connectivity distance information lies in between (MAE = 0.15 eV). For the rest of this work, we present the results trained with initial structures and the adsorbate distance information.

Using adsorbate connectivity distance changes between initial and final geometries, we could tackle a few drawbacks of using initial or final geometries for training (Figure 3.3d). For example, if CO binding at bridge site is more stable than binding at top site, CO adsorbate could move from top to bridge site during the DFT relaxation. When using the initial structure for training, the model would predict the bridge site CO binding energy from the top site CO binding structure. Training with final structures results in higher accuracy, but using the final structures is not practical. With the connectivity distance change, the model predicts the bridge site CO binding energy from both top and bridge sites CO binding initial geometries. We note that we do not use atomic positions from final geometries. Practically, we would use initial structure-based model at first, which already demonstrated a reasonable prediction accuracy, and we improve the accuracy of the model by including the connectivity distance information of adsorbates from the completed DFT calculations in machine-learning augmented high-throughput DFT scheme (Figure 3-2).



d. Prediction of CNN



Figure 3-1: A graphical representation of converting an atomic structure containing \mathbf{n} atoms into numerical inputs for the convolutional neural network. (a) 9 basic atomic properties are presented by one-hot encoding [226] to prepare the atomic feature vectors. (b) Neighbor information was encoded using the solid angle (Ω) based on the Voronoi polyhedron. The grey skeleton indicates the Voronoi polyhedron of Cu atom, and the solid angle between Fe atom and the shared plane of Cu and Fe polyhedrons is marked. Other nearest neighbor atoms were omitted for simplicity. Color code for atoms: Cu (blue), Fe (brown), C (grey), O (red). (c) The connectivity distances from adsorbate to all atoms in surface structures are counted. A side view of the surface structure with adsorbate CO molecule, surface atoms up to the second layer and their corresponding connectivity distances are presented as an example. (d). A schematic representation of three atomic structural data that could be used in the CNN. "Initial geometry", "Final geometry", and "Final adsorbate coordination geometry" are based on initial structures, final structures, and initial structures with adsorbate connectivity distances, respectively. The numbers in the DFT part (solid box) represent the initially set or relaxed atomic distances, and the integer numbers in the CNN part (dotted box) are the connectivity distances. More details can be found in the main text.

3.4 Graph Convolutional Neural Networks

Once we converted atomic structures into graph representations, we concatenated atomic feature vectors of atom **i** and **j** (\mathbf{v}_i and \mathbf{v}_j) and neighbor feature vectors between atom **i** and **j** ($\mathbf{u}_{(i,j)}$), i.e., $\mathbf{z}_{(i,j)} = \mathbf{v}_i \oplus \mathbf{v}_j \oplus \mathbf{u}_{(i,j)}$. We then utilized the same



Figure 3-2: A process of high-throughput workflow for the catalyst screening.



Figure 3-3: Schematic illustration of the convolutional neural network on top of the graph representation of atomic surface structures. Dark red and red blocks correspond to atomic and neighbor feature vectors as in Figure 3.3. Pink blocks correspond to the connectivity distance vectors used as a filter to exclude atoms that are too far from adsorbates (higher than the connectivity distance of 2). More details are elaborated in the main text.

convolutional layer (eqn (1)) as described in Ref. [226].

$$\mathbf{v}_{i}^{t} = \mathbf{v}_{i}^{t-1} + \sum_{j} \sigma \left(\mathbf{z}_{(i,j)}^{t-1} \mathbf{W}_{f}^{t-1} + \mathbf{b}_{f}^{t-1} \right) \odot g \left(\mathbf{z}_{(i,j)}^{t-1} \mathbf{W}_{s}^{t-1} + \mathbf{b}_{s}^{t-1} \right)$$
(1)

where \odot denotes an element-wise multiplication, σ denotes a sigmoid function, g denotes non-linear activation functions ("Leaky ReLu" in this study), **W** and **b** denote weights and biases of the neural networks, respectively. After **R** convolutional layers, resulting vectors are then fully connected via **K** hidden layers, followed by a linear transformation to scalar values. Distance filters were applied to exclude contributions of atoms that are too far from the adsorbates. At this stage, (1) atomic contributions of each atom to binding energies could be predicted and (2) mean pooling layer is applied to predict DFT calculated binding energies (Figure 3-3).

The neural network is trained to minimize the loss function (MAE) between the DFT calculated and predicted binding energies using Adam optimizer with decoupled weight decay (L₂ Regularization coefficient: 10^{-5}) and warm restart [132]. We

Parameters	CO binding energies	H binding energies
Batch size	214	140
Learning rate	5.6×10^{-3}	9.9×10^{-4}
Epochs	150	150
atom_fea_len	46	107
h_fea_len	83	50
n_conv	8	6
n_h	4	1

Table 3-1: Hyperparameters of convolutional neural network optimized by Sigopt.

tested two pooling functions (sum and mean) and four activation functions (Sigmoid, Softplus, leaky ReLu, ReLu). We divided DFT results into 20 % and 80 % of test and training sets, respectively. To prevent the overfitting, the validation set (25 % of training sets) was used to optimize hyperparameters (Table 3-1). Based on the number of hyperparameters for CO binding prediction, we presented the convolutional neural network architecture in Figure 3-4. The prediction accuracy of the CNN models was evaluated from mean absolute error (MAE) of CNN-predicted and DFT-calculated binding energies.

We first compare how our DFT datasets and method are different from the previous approaches [3, 7, 16, 32, 48, 153] (Table 3-2). First, the DFT data used in most previous studies covered only the limited chemical space—e.g., materials consisting of up to two elements based mainly on transition metals, fixed elemental compositions for alloys (1:1, 1:3) and low-index facets of the most common crystal structures (FCC (111), FCC (100), HCP (0001)). On the other hand, our DFT datasets include materials consisting of up to four elements using 37 elements, 96 stoichiometries, 110 space groups and 41 facets. Second, the number of parameters in our CNN method is substantially more compared to the previous approaches. For binding energy predic-


Figure 3-4: Convolutional neural network architecture used in this work. The values in parenthesis correspond to the number of parameters for CO binding energy prediction (Table S2).

Table 3-2: A comparison of our DFT data and the number of parameters of our model to literature values, where the goal was to predict binding energies. We counted all unique components (element, stoichiometry, space group, facet) in literature and our DFT data. For the previous studies that did not explicitly mention the number of total DFT data or parameters, we estimated based on the presented results. The number of parameters indicate how complicated the model is. Specific parameter components of our model are summarized in Table S2.

	Element	Stoichiometry	Space group	Facet	Total DFT data	Parameters	Methods
Calle-Vallejo et al.[32]	1	1	1	22	132	2	GCN^a
Abild-Pederson et al.[3]	13	1	3	6	165	2	SR^b
Dickens et al.[48]	45	3	3	3	901	2	ED^{c}
Ma and Xin.[136]	3	1	1	8	37	$\sim 10^1$	$GCN^a + ED^c$
Noh et al.[153]	28	3	1	1	263	$10^{1} \sim 10^{2}$	ML^d
Andersen et al.[7]	9	3	1	4	884	$\sim 10^1$	ML^d
Li et al.[125]	24	4	1	1	1,032	$\sim 10^2$	ML^d
Batchelor et al.[16]	5	1	1	1	1,869	15 (O), 55 (OH)	ML^d
This work	37	96	110	41	43,247	4,938 (CO), $6,738$ (H)	ML^d

^{*a*}generalized coordination number, ^{*b*}scaling relations, ^{*c*}electronic descriptors, ^{*d*}machine-learning.

tions through scaling relations[3], generalized coordination numbers[32] or electronic descriptors[48], only two parameters (slope and bias) are optimized through a linear regression. ANNs consisting of an input layer, several hidden layers and an output layer are reported to have $\sim 10^2$ parameters.[7, 16, 125] In our deep-learning CNN, there are thousands of parameters to be optimized during the training so that it can predict binding energies on a variety of surfaces of catalysts. However, we note that simpler methods are useful in specific cases, which will be discussed in the last part.



Figure 3-5: Mean absolute error (MAE) of test set with respect to the number of training atomic structures for binding energy predictions of (a) CO and (b) H. The horizontal dashed lines correspond to the standard deviation of DFT binding energies. As a comparison, we added the prediction error from our previous study (green and red star).[205] Two-dimensional histograms of DFT-calculated and CNNpredicted binding energies of (c) CO and (d) H.

3.5 Results

3.5.1 Binding Energy Prediction

Figure 3-5 shows the performance of our method on predicting CO and H binding energies. We note that a difference in the prediction accuracy between CO and H using the small number of data could be originated from their data distribution; the standard deviations of CO and H DFT binding energies are 0.65 and 0.42 eV, respectively. In both cases, we observed a systematic decrease in the error metric (MAE) as the number of training atomic structures increased. Interestingly, we observed a near convergence of the prediction accuracy at 5,000 and 8,000 training data for CO and H, respectively, similar to our previous work. [205] For comparison, we also plotted MAE obtained using 12,000 training data from our previous work[205], where fingerprint-based surrogate models were trained. We note that exact comparison is not available since training/validation/test data splitting was not used previously, but training/test splitting was used instead. The previously reported values are, thus, the lower-limit accuracy of the model. The best MAEs we achieved for the CO and H binding energy prediction are 0.15 eV in both cases, and 86 % of test data are within the accuracy of 0.25 eV. We highlight that MAEs for CO and H binding energy prediction in this study are lower by 0.05 eV compared to our previous model. Given a reasonable accuracy of our model and facile preparation of input data, we expect this method could be applied for the high-throughput catalyst screening to remove unpromising candidates, thus, reducing the number of DFT jobs.

3.5.2 Atomic Contributions to Binding Energies

After the convolutional layers and the fully connected layers, the resulting output vectors are linearly transformed to scalar values for each atom, which are then filtered by the connectivity distance criterion. We only considered atoms with the connectivity distance up to 2 as they are expected to strongly interact with adsorbates. These values are then averaged in the pooling layer and the resulting value corresponds to the predicted binding energies. Thus, the linearly transformed scalar values represent the contribution of corresponding atoms to the predicted binding energies, and one possible interpretation of the model is a saliency map of atomic contributions to binding energies on this basis (Figure 3-6). We note that neighboring atoms of the active sites could significantly affect the electronic structures of the active sites, which our model cannot capture, and this effect can only be understood by density-of-state analysis. The prediction accuracy change was negligible (< 0.01 eV) even we included all the atoms, indicating that the first and the second nearest-neighbor of adsorbates mainly contribute to binding energies. In Figure 3-6, we present a graphic representation of each atomic contribution to binding energies using CO adsorption at top and bridge sites of Cu (211) surface as an example. We normalized the atomic contribution into [0,1] range by mapping the minimum and the maximum values of the contributions to 0 and 1, respectively. Clearly, we observe that surface atoms directly interacting with adsorbates have the highest contribution to the binding energies. We also visualized atomic contributions using the connectivity distance up to 4, and, as expected, their contributions are found to be negligible. This result rationalizes the use of the first and second nearest neighbor to extract neighbor information of active sites for the preparation of fingerprints in our previous model.[205] An example of bimetallic surface is shown in Figure 3-7. We note that the conventional DFT calculations are not capable of isolating the calculated binding energies into each atomic contribution.

3.5.3 Automated Failed DFT Calculation Detection

Another feature of our method is to automatically detect wrong DFT calculations in datasets, which may take lots of time by human analysis. Since our workflow



Figure 3-6: Graphic representations of atomic contributions to CO binding energies at (a) top and (b) bridge site of Cu (211) surface. Darker spheres indicate higher contributions. Only atoms within the connectivity distance of 2 are considered and transparent spheres indicate atoms with connectivity distances higher than 2.



Figure 3-7: Top view of CO adsorption on Cu_3Al (211) surface. Darker spheres indicate higher contributions. Dark brown spheres are Cu atoms directly interacting with CO adsorbate (distance = 1), light grey spheres are Cu atoms not interacting with the CO (distance = 2). Lilac spheres are Al (distance = 2). Transparent spheres indicate atoms with connectivity distances higher than 2.

in *GASpy* automatically performs DFT calculations (see the Computational Methods for details of calculations), there could be a few incorrect atomic structures and binding energies, where the final relaxed structures significantly differ from the initial structures due to surface reconstructions or unphysical initial atomic structures (Figure 3-8). Since our model uses neighbor information based on the Voronoi polyhedra, it cannot accurately predict binding energies of structures where atomic structural change is significant. After training, we considered data outside of $\pm 3\sigma$ as outliers and this process was repeated two times (Figure 3-9 and 3-10). Out of 20,771 and 22,361 data for CO and H binding, 622 and 628 outliers were detected.

3.6 Conclusion

The deep-learning CNN results presented in this work suggest that, given sufficient training data, it could be effective in reducing computational cost issue in the highthroughput screenings. However, we highlight that simpler methods may be more effective for specific cases or smaller datasets. Details are discussed in the following:

- For systems consisting of single metal element, generalized coordination number (GCN) is the best choice. The GCN analysis on low and high index Pt facets in various sizes of Pt nanoparticles predicted OH* binding energies very accurately (MAE ~0.056 eV) using only 20 DFT data.[32] In addition, one can easily count the GCN of active sites. However, an extension of this method to alloy or other materials has not been reported.
- Recently developed orbitalwise coordination number [136, 211] uses the geometric strain and the tabulated interatomic coupling matrix elements. It is as easy as to compute GCN, but it still requires DFT calculated geometries for a high accuracy. It has shown a reasonable prediction accuracy for CO, O and OH adsorption (root mean square error (RMSE): $0.10 \sim 0.19$ eV) for various sizes



Figure 3-8: Three types of outliers detected during the analysis. Surface ID consists of MPID, miller index, shift value and orientation (True and False correspond to top and bottom of surface structures, respectively).



Figure 3-9: Two iterations of the outlier detection for CO binding energy prediction.



Figure 3-10: Two iterations of the outlier detection for H binding energy prediction.

and shapes of pure Au and Pt nanoparticles, and transition metal core-Pt shell nanoparticles, but an extension to other systems consisting of more elements, facets and various crystal structures has not been reported.

- Scaling relations and electronic descriptors, such as d-band center, are suitable for a dozens of metal/alloy DFT data. The drawback of the scaling relation is that it is adsorbate-specific and a model should be optimized for each adsorbate of interest [3]. For d-band center analysis, one should perform DFT optimization and post density-of-states analysis to calculate d-band center values. Further, a preparation of inputs requires DFT calculations for new systems to obtain binding energies of atomic species, d-band center values or the orbitalwise coordination number.
- Artificial neural networks (ANN) have also shown a similar accuracy to the scaling relations or electronic descriptors [7, 16, 125, 153], and they are suitable for hundreds/thousands of DFT data of similar metal/alloy systems. They require a careful preprocessing of fingerprints which could considerably affect the prediction accuracy. A preparation of inputs for a binding energy prediction of new systems is usually fast as most of the fingerprints are determined beforehand.
- The deep-learning methods, CNN as in this study, could predict binding energies for a variety of surfaces without a preparation of fingerprints as they collect all the information automatically. However, the main drawback of the deep-learning is that it requires lots of data (> 10,000 DFT data). As we want to include diverse materials for the high-throughput screenings, the deep-learning methods would be ideally suited to this purpose.

In summary, we modified the original crystal graph CNN code by Xie and Grossman[226] to collect neighbor information using Voronoi polyhedra for the application in predicting binding energies on heterogeneous catalyst surfaces. Our method predicted CO and H binding energies with 0.15 eV MAE for a variety of materials which have never been considered previously. Given the reasonable accuracy of our model, we expect that binding energy prediction from our method could help to rule out unpromising candidates and reduce the number of DFT jobs in the high-throughput catalyst screening to effectively facilitate a catalyst discovery. Furthermore, our method successfully partitioned the calculated binding energies into each atomic contribution, rationalizing the use of the first and second nearest neighbor atoms for the preparation of fingerprints. It also automatically detected wrong DFT calculations generated during the high-throughput calculations, which will be useful to remove those outliers from the datasets.

3.7 Dataset Preparation and Methods

We collected DFT calculated results of CO and H binding on surfaces of pure metals, metal alloys and intermetallic alloys from our open-source GASpy database (~20,000 data for each adsorbate). In the following, the calculation procedure is elaborated and more details can be found in our previous publication.[205]

Our DFT calculated data were automatically generated as follows; 1) 1,499 bulk structures from the Materials Project[100] are relaxed. 2) Using the optimized bulk structures, unique surfaces with Miller indices between -2 and 2 are enumerated, resulting in 17,507 surfaces. 3) All unique active sites are identified, 1,684,908 sites in total, and they are uploaded to the database. 4) Using "Materials Id" in the Materials Project as an user input, all the required jobs (structure relaxation with and without adsorbates) are automatically submitted to supercomputers. 4) The calculated binding energies and initial/final structures of successfully completed jobs with a residual force less than 0.05 eV/Å are uploaded to our *GASpy* database.

We used Pymatgen[157] to enumerate various surfaces and to find unique ac-

tive sites. We used Luigi[22] and Fireworks[102] to manage high-throughput DFT calculations across many clusters. DFT calculations were performed using VASP code[111, 112] with the revised Perdew-Burke-Ernzerhof (RPBE) functional[81] and projector augmented wave (PAW) pseudopotential[113].

Chapter 4

Differentiable Optimization for the Prediction of Ground State Structures (DOGSS)

This work originally appeared as: Junwoong Yoon and Zachary W. Ulissi. Differentiable Optimization for the Prediction of Ground State Structures (DOGSS). Phys. Rev. Lett 125, 173001 (2020). It has been edited to include the supporting information in-line.

4.1 Abstract

Ground-state or relaxed inorganic structures are the starting point for most computational materials science or surface science analyses. Many of these structure relaxations represent systematic changes to the structure but there are currently no general methods to improve the initial structure guess based on past calculations. Here we present a method to directly predict the ground-state configuration using differentiable optimization and graph neural networks to learn the properties of a simple harmonic force field that approximates the ground-state structure and properties. We demonstrate this flexible open source tool for improving the initial configurations for large datasets of inorganic multicomponent surface relaxations across 32 elements, and the relaxation of adsorbates (H and CO) on these surfaces. Using these improved initial configurations reduces the expensive adsorbate-covered surface relaxations by approximately 50%, and is complementary to other approaches to accelerate the relaxation process.

4.2 Introduction

Fast and accurate optimization methods that provide ground-state inorganic structure and properties are crucial for materials science, surface science, and catalysis, among many other fields. Electronic structure methods have made remarkable progress in predicting inorganic structure-property relationships based on relaxed or ground-state structures [80, 87, 103, 154, 155, 179, 218?]. In most cases, these ground-state structures are identified by relaxing initial structures guessed by either experts or automated software packages. The structural relaxations of these initial configurations often require high computational cost that has impeded large scale explorations of new inorganic materials. Unfortunately, there are currently no general methods to predict better initial configurations for the relaxation process, leading to massive redundancy in calculations.

One approach to accelerating the relaxation is the development of machine learning potentials (MLPs) that mimic commonly-used electronic structure methods such as density functional theory (DFT) [19, 45, 46, 166, 177, 178]. MLPs approximate the true potential energy surface (PES) by fitting numerous DFT single-point energies and forces. A perfect approximation to the ground-truth DFT PES is obviously desirable, but residual force errors in the fitted force field often yield different ground-



Figure 4-1: Workflow to accelerate electronic structure methods using DOGSS as a pre-processing model. (1) DOGSS predicts the ground-state structure by relaxing the initial structure using the learned harmonic force field parameters. (2) and (3) electronic structure methods employ the predicted structure as a starting configuration, thereby converge to the ground-state in fewer optimization steps.

state structures, or worse, lead to unphysical configurations. We note that there has been only very limited success in developing MLPs that work across the elemental composition. Despite the considerable excitement about these methods for practical simulations, no MLP has been demonstrated for the specific challenge of predicting ground-state inorganic structures.

In this Letter, we present a deep learning framework, Differentiable Optimization for the prediction of Ground State Structures (DOGSS), that incorporates differentiable optimization around an end-to-end trainable Graph Neural Network (GNN) to approximate ground-state structures of inorganic materials. Over the last few years, there has been a surge of interest in GNN for investigating properties of materials due to their unique ability to model irregular structures [11, 35, 51, 78, 108, 224, 225, 227].

We propose that our GNN-based method can be used as a pre-processing tool that provides any electronic structure method with a starting configuration near the ground-state, and thereby accelerates the relaxation as illustrated in Figure 5.3.1. DOGSS first learns the optimal properties of the harmonic force fields between atoms to predict PES of the structure. The predicted harmonic PES is then minimized and the relaxed structure is generated. The harmonic force field is not evaluated on how well it fits energies and forces; instead, it is evaluated by how well the structure predicted using the force field matches the ground-state structure. Finally, the relaxation of the predicted structure using electronic structure methods can be performed in fewer optimization steps compared to the relaxation of the initial structure.

4.3 Graph Neural Network

4.3.1 Representation of Surface Structures

Similar to the GNN method developed by Xie and Grossman [224], DOGSS represents an atomic structure by an undirected graph consisting of nodes and edges. Each node i is described by a node feature v_i , a stack of one-hot encoded atomic and periodic properties of atom i (properties used to define node features can be found in Table 4-1). In addition to those property encoding, we have incorporated orbital-field matrix (OFM) representations for all element types used in this work. Example OFM representations using Pt and Au atoms are shown in Table 4-2

We define the local environment for each atom i by sampling M nearest neighbor atoms connected by distinct edges. Each edge (i,j) is represented by an edge feature $\boldsymbol{u}_{i,j}$ that encodes the interatomic distance between atom i and atom j and the sum of their covalent radii. Unlike case-by-case human-engineered atomic representations [13, 15, 17, 38, 63, 203], the molecular graphs are defined by raw node ($\boldsymbol{v} \in \mathbb{R}^{N \times F_v}$) and edge ($\boldsymbol{u} \in \mathbb{R}^{N \times M \times F_u}$) features that will be further engineered during convolutions. We note that N is the number of nodes, and F_v and F_u are lengths of node features and edge features respectively.



Figure 4-2: Structure of the graph neural network with embedded differentiable optimization. Initial node v^0 and edge u^0 features are fed to T convolutional layers. Two separate hidden networks take the updated features to learn properties of distinct springs connecting every atom and neighbor atom pairs: spring constants k and equilibrium spring distances r_e . With these learned harmonic force field parameters, harmonic PES is computed and minimized using the gradient descent to approximate the ground-state structures.

4.3.2 Convolutional Layers - Message Passing Scheme

Convolutional layers for the molecular graphs can be expressed as a message-passing scheme [58, 69, 174]. For each convolution during T iterations, messages exchanged between node i and node j along the edges (i, j) are generated. These messages capture the pair-wise interactions between atoms within the local chemical environment. The node feature v_i and edge feature $u_{i,j}$ are then separately updated using the node messages ϕ_v and edge messages ϕ_u respectively. We update the node feature v_i according to

$$\phi_v^{t+1} = s(\boldsymbol{z}_{i,j}^t \boldsymbol{W}_{s,v}^t) \odot \sigma(\boldsymbol{z}_{i,j}^t \boldsymbol{W}_{f,v}^t)$$
(4.1)

$$\boldsymbol{v}_{i}^{t+1} = s(\boldsymbol{v}_{i}^{t} + \sum_{j} \phi_{v}^{t+1}),$$
(4.2)

Property	Unit	Range	Number of categories
Group number ^{a}	-	1,2,, 18	18
Period number ^{a}	-	1,2,,9	9
\mathbf{E} lectronegativity ^{<i>a</i>}	-	0.5-4.0	10
Covalent radius ^{a}	pm	25-250	10
Valence $electrons^a$	-	1,2,, 12	12
First ionization $energy^a$	eV	1.3-3.3	10
Electron affinity ^{a}	eV	-3-3.7	10
Block^a	-	s, p, d, f	4
Atomic volume ^{a}	$\rm cm^3/mol$	1.5-4.3	10
Orbital ^b	-	$s^1, s^2,, f^{13}, f^{14}$	36

Table 4-1: Properties used in node feature vector v_i .

^aProperties from CGCNN representation [224].

^bProperties from orbital-field matrix (OFM) representation [33].

Table 4-2: Example OFM representations for Pt and Au atoms.

Element	$\mathbf{s^1}$	s^2	$\mathbf{p^1}$	$\mathbf{p^2}$	\mathbf{p}^{3}	\mathbf{p}^4	\mathbf{p}^{5}	\mathbf{p}^{6}	d^1	d^2	d^3	d^4	d^5	d^6	d^7	d^8	d^9	d^{10}	$\mathbf{f^1}$	f^2	f ³	f^4	f ⁵	f^6	f^7	f^8	f9	f ¹⁰	f ¹¹	f^{12}	f^{13}	f ¹⁴
Pt	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1
Au	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1

where σ and s denote a sigmoid function and a softplus activation function respectively, and \odot denotes element-wise multiplication. $\boldsymbol{z}_{i,j}^t \in \mathbb{R}^{N \times M \times F_z}$ represents a local feature $(\boldsymbol{z}_{i,j}^t = \boldsymbol{v}_i^t \oplus \boldsymbol{v}_j^t \oplus \boldsymbol{u}_{i,j}^t)$ that is used with weights $\boldsymbol{W}_{s,v}^t \colon \mathbb{R}^{F_z} \to \mathbb{R}^{F_v}$ and $\boldsymbol{W}_{f,v}^t \colon \mathbb{R}^{F_z} \to \mathbb{R}^{F_v}$ to learn a self-feature $s(\cdot)$ and a filter $\sigma(\cdot)$ respectively. Once messages ϕ_v^{t+1} for each pair of node i and node j for $(t+1)^{\text{th}}$ convolution are generated, they are aggregated (summed) to compute an overall message for node i from its local environment. The node feature \boldsymbol{v}_i^t is then updated using the aggregated message and a softplus activation. Similarly, the edge feature $\boldsymbol{u}_{i,j}$ can be updated as,

$$\phi_u^{t+1} = s(\boldsymbol{z}_{i,j}^t \boldsymbol{W}_{s,u}^t) \odot \sigma(\boldsymbol{z}_{i,j}^t \boldsymbol{W}_{f,u}^t)$$
(4.3)

$$\boldsymbol{u}_{i,j}^{t+1} = s(\boldsymbol{u}_{i,j}^t + \phi_u^{t+1}), \tag{4.4}$$

where $\boldsymbol{W}_{s,u}^t \colon \mathbb{R}^{F_z} \to \mathbb{R}^{F_u}$ and $\boldsymbol{W}_{f,u}^t \colon \mathbb{R}^{F_z} \to \mathbb{R}^{F_u}$ are self-feature generating weights and filter generating weights respectively for edge update. Unlike node update, each edge feature is directly updated based on the pre-aggregated pair-wise messages ϕ_u^{t+1} during each $(t+1)^{\text{th}}$ convolution.

4.3.3 Differentiable Optimization

After T convolutions, both updated node and edge features are fed to two distinct hidden networks to learn the harmonic force field parameters: (i) pair-wise equilibrium spring distances for atom-neighbor pairs $\mathbf{r}_e \in \mathbb{R}^{N \times M \times 1}$ and (ii) spring constants $\mathbf{k} \in \mathbb{R}^{N \times M \times 1}$ for springs connecting each atom-neighbor pair. Once the force field parameters are learned by the hidden networks, the harmonic PES is constructed as

$$\hat{E}(\boldsymbol{r}_{o}(\boldsymbol{x}_{o}), \boldsymbol{r}_{e}, \boldsymbol{k}) = \sum_{i} \sum_{j} \boldsymbol{k}_{i,j} (\boldsymbol{r}_{o,(i,j)}(\boldsymbol{x}_{o}) - \boldsymbol{r}_{e,(i,j)})^{2}, \qquad (4.5)$$

followed by an energy minimization based on the following optimization problem:

$$\boldsymbol{x}_{pred} = \operatorname*{arg\,min}_{\boldsymbol{x}_o \notin \boldsymbol{x}_{con}} \hat{E}(\boldsymbol{r}_o(\boldsymbol{x}_o), \boldsymbol{r}_e, \boldsymbol{k}), \qquad (4.6)$$

where \boldsymbol{x}_{pred} is a set of atomic positions at local minima in the harmonic PES, \boldsymbol{x}_{con} is a set of positions of constrained atoms that are not allowed to move, and $\boldsymbol{r}_o(\boldsymbol{x}_o) \in \mathbb{R}^{N \times M \times 1}$ represents pair-wise distances between atoms and their neighbor atoms in the initial structures.

The argmin operation over the unconstrained surface atoms is optimized using the gradient descent algorithm,

$$\boldsymbol{x}_{o}^{p+1} = \boldsymbol{x}_{o}^{p} + \eta \nabla_{\boldsymbol{x}_{o}^{p}} \hat{E}(\boldsymbol{r}_{o(\boldsymbol{x}_{o}^{p})}, \boldsymbol{r}_{e}, \boldsymbol{k}).$$

$$(4.7)$$

After the gradient descent has converged to local minima in P steps, the model outputs the updated atomic structures \boldsymbol{x}_{pred} which we consider as the approximated ground-state structures. To measure the closeness between x_{pred} and the groundstate structures x_{min}^* , we define a loss function for the GNN model as follows:

$$L = \frac{1}{N} \sum_{i} \|\boldsymbol{x}_{i,pred} - \boldsymbol{x}_{i,min}^*\|_2.$$
(4.8)

This loss represents the average per-atom distance that the unconstrained atoms in the predicted structures need to move to arrive at the ground-states. Initial loss between \boldsymbol{x}_o and \boldsymbol{x}_{min}^* was also computed to be compared with the prediction loss. In case an atom crosses the boundary of its unit cell to reach the ground-state, the loss invokes the minimum-image distance convention that considers the nearest periodic image of the ground-state atom. The entire network was trained using endto-end backpropagation with AdamW optimizer [132] to minimize the prediction loss. Hyperparameters were optimized using SigOpt [40], a model tuning software package, and the optimized hyperparameters are summarized in Table 4-3.

Hyperparameters	Bare-Surfaces	H adsorption	CO adsorption
Batch size	17	18	37
Embedding layer	236	103	125
Number of convolutional layers	12	12	13
Length of spring constant vector ${\pmb k}$	4	18	117
Length of spring equilibrium constant vector \boldsymbol{r}_e	4	18	117
Depth of hidden network for ${\pmb k}$	1	16	1
Depth of hidden network for \boldsymbol{r}_e	1	1	1
Learning rate	0.0393415	0.003770	0.0145320

 Table 4-3: Hyperparameters optimized using SigOpt [40].

4.4 Dataset Preparation

To demonstrate the effectiveness of our proposed framework, we predict the groundstate structures of diverse surfaces of pure metals, metal alloys, and intermetallic alloys covering 32 different elements, with and without the presence of adsorbates (H and CO) collected from our open-source database, Generalized Adsorption Simulation for Python (*GASpy*) [204]. We first relaxed bulk structures from the Materials Project [101] using DFT. Unique surfaces with Miller indices between -2 and 2 were enumerated from the relaxed bulk structures, then were relaxed. For each relaxed surface, the adsorbate was placed at a certain height above the site of interest, and then the structure was relaxed using DFT. More details about the data preparation can be found in our previous publication.[204] From the GASpy database, we only sampled structures that guaranteed each atom at least 12 nearest neighbors within 7 Å radius. Further, structures that were already at the ground-state or caused major reconstructions were also filtered out. Finally, we collected 6,821 DFT-relaxed diverse bare metal surfaces, 11,411 surfaces with H, and 8,250 surfaces with CO.

Each of the three datasets was randomly split into train/validation/test (80/10/10) sets for training and testing DOGSS on structures covering 32 surface elements. We also sampled about 4-6 intermediate structures from the DFT relaxation of each of the training structures to train force fields using SchNet [178].

4.5 Practical Motivation

After training DOGSS, we computed prediction loss on the test sets to evaluate how close the predictions are to the ground-state structures. The same test structures were also relaxed using the trained SchNet force fields to generate ground-state structures and the same loss was used to compute the prediction loss on those relaxed structures for comparison. We summarize the performance of DOGSS and SchNet

Dataset	Sample	Size	$L_{ m surf}$	(Å)	$L_{ m ads}$	(Å)	$L_{ m total}~({ m \AA})$		
	Training	Test	Initial	Pred.	Initial	Pred.	Initial	Pred.	
Bare Suface									
DOGSS	$5,\!524$	683	0.1337	0.0635	N/A	N/A	0.1337	0.0635	
SchNet	$21,\!034$	683	0.1337	1.1449	N/A	N/A	0.1337	1.1449	
H adsorption									
DOGSS	9,242	$1,\!142$	0.0855	0.0508	0.7841	0.3301	0.1358	0.0709	
SchNet	37,346	$1,\!142$	0.0855	0.6104	0.7841	1.0594	0.1358	0.6428	
CO adsorption									
DOGSS	$6,\!682$	825	0.0844	0.0548	0.7286	0.3237	0.1706	0.0908	
SchNet	43,787	825	0.0844	0.4239	0.7286	0.9116	0.1706	0.4891	

Table 4-4: Summary of the ground-state structure prediction performance ofDOGSS and SchNet on test sets in the three different datasets.

for approximating the ground-state structures for the three different datasets in Table 4-4. We include initial losses to quantitatively compare the relative closeness between the predicted structures and the ground-state structures. We further separate the total loss (L_{total}) into the loss for the surface atoms (L_{surf}) and the loss for the adsorbate atoms (L_{ads}) if applicable. For all three datasets, the force field method using SchNet could not achieve enough prediction accuracy and resulted in different ground-state structures. We note that we provided only limited examples (4-6 examples of DFT relaxation of each structure) for training SchNet on our datasets. SchNet previously demonstrated accurate force field predictions for small organic molecules using 1,000 to 50,000 training examples and a separate model for each molecule. However, training a separate force field model using the same amount of training examples for each of the unique inorganic multicomponent structures in our datasets is computationally infeasible. By comparison, DOGSS only requires



Figure 4-3: Histograms showing distributions of the individual per-atom distances $||x - x_{min}^*||_2$ for surface atoms in each of the three datasets: (left) bare surfaces, (center) H adsorption, and (right) CO adsorption. For all three datasets, the distributions for atoms in the initial structures are shown in orange and the distributions for atoms in the predicted ground-state structures by DOGSS are shown in blue. The insets of the center and the right figures show the per-atom distance distributions for adsorbate atoms in H and CO adsorption datasets respectively with visualized example configurations.

the initial and the final relaxed structures for directly predicting ground-state structures instead of using thousands of example structures to fit every part in the PES. Besides, one DOGSS model can be used to approximate ground-state structures of diverse materials instead of training a separate force field model for each material.

4.6 Results

4.6.1 Loss (per-atom distances)

Distributions of individual per-atom distances in the predicted structures (blue) by DOGSS and the initial structures (orange) for the three different datasets are presented in Figure 4-3. We could clearly observe that the surface atoms in the predicted structures for the bare surface dataset are more closely located to the ground-state positions compared to the atoms in the initial structures. Quantitatively, DOGSS could reduce the initial loss for the surface atoms in this dataset by 54%. We acknowledge that the amount of reduction in the loss does not guarantee the same amount of reduction in the DFT relaxation steps. However, we will shortly demonstrate that the ability of our method to reduce the loss can effectively attain a reduction in the computational cost required for the actual DFT relaxations.

For H and CO adsorption datasets, DOGSS could reduce the total losses for both datasets by approximately 47%. We note that selecting initial positions of adsorbates currently relies on experts' guesses, and the insets in the center and the right figures in Figure 4-3 show this randomness in adsorbates' positions in the initial structures (orange). Also, most of the metal surfaces in both datasets were initially relaxed before the adsorbates were added, but the presence of adsorbates near the surfaces changes the local environment of the surface atoms and their positions. About 35-40% reduction in the losses for surface atoms in both datasets shows that DOGSS was able to capture the surface rearrangements induced by the presence of adsorbates. More importantly, DOGSS could approximate ground-state positions of adsorbate atoms by approximately 58% and 55% respectively.

We studied the impact of different potentials such as the Morse potential and the Lennard-Jones (LJ) potential on the performance of DOGSS, and Table 4-5 summarizes the results. In general, the harmonic potential with the gradient descent we proposed provides a better approximation of the ground-state structures. The Morse potential and the LJ potential encounter the vanishing gradient problem that causes a bottleneck in the energy minimization using the simple gradient descent algorithm when atoms are too far apart from their neighbor atoms.

Additional experiments were also conducted to investigate DOGSS's ability to

Dataset	$L_{ m s}$	urf	L_{i}	ads	$L_{ m total}$			
Dataset	Initial	Pred.	Initial	Pred.	Initial	Pred.		
Bare Suface								
- Harmonic	0.1379	0.0631	N/A	N/A	0.1379	0.0631		
- Morse	0.1379	0.0647	N/A	N/A	0.1379	0.0647		
- LJ	0.1379	0.0636	N/A	N/A	0.1379	0.0636		
H-adsorption								
- Harmonic	0.0867	0.0527	0.8121	0.3324	0.1380	0.0725		
- Morse	0.0867	0.0556	0.8121	0.3525	0.1380	0.0766		
- LJ	0.0867	0.0607	0.8121	0.4220	0.1380	0.0863		
CO-adsorption								
- Harmonic	0.0844	0.0548	0.786	0.3237	0.1706	0.0908		
- Morse	0.0844	0.0720	0.786	0.4630	0.1706	0.1243		
- LJ	0.0844	0.0611	0.786	0.3547	0.1706	0.1004		

Table 4-5: Performance of DOGSS using the original 80/10/10 split with three different types of potential: Harmonic, Morse, and Lennard-Jones (LJ) potentials with the gradient descent.



Figure 4-4: Comparison of properties of the initial structures (orange) and the predicted ground-state structures (blue) by DOGSS for the three datasets: (left) bare surfaces, (center) H adsorption, and (right) CO adsorption. (a) Average absolute differences between total energies of the initial structures or the predicted structures and the ground-state energies. (b) Average maximum inter-atomic forces in each of the initial structures and the predicted ground-state structures. Error bars represent the 95% confidence interval.

generalize to new, unseen, inorganic structures that contain new elements using leave-one-out and leave-two-out data splitting methods. Table 4-6 summarizes the performance of DOGSS on new structures. The prediction loss slightly increases as DOGSS gets tested on structures containing more number of new elements, but in general, DOGSS could generate reasonably approximated ground-state structures of those unseen structures. DOGSS could reduce the initial loss of H and CO adsorbate atoms by approximately 40-45% and the initial loss of surface atoms in the bare surface dataset by approximately 30-40% depending on the number of new elements in the test structures.

4.6.2 Properties of Predicted Structures

Further, in Figure 4-4 we compare the relative properties of the initial structures in the test sets and the corresponding predicted structures by DOGSS to the groundstate structures. Figure 4-4 (a) shows that DOGSS reduces the difference in the total energies between the initial structures and the ground-state structures $|\Delta E_{\text{initial}}|$ for the three datasets by a factor of 5.63 (bare surfaces), 9.41 (H adsorption), and 7.07 (CO adsorption). We find that the total energy differences between the predicted structures and the ground-state structures $|\Delta E_{\text{pred}}|$ in all three datasets are close to 0.1 eV. Similarly, Figure 4-4 (b) shows that the maximum inter-atomic forces $|F_{\text{max}}|$ in the initial structures are reduced by a factor of 2.80 (bare surfaces), 8.72 (H adsorption), and 2.28 (CO adsorption) after pre-processing the initial structures using DOGSS. Given these comparisons with the initial structures, DOGSS's predicted structures and their properties are fairly close to the ground-states.

4.6.3 Reduction in Actual DFT Optimization Steps

To show the feasibility of our proposed workflow as illustrated in Figure 5.3.1, we performed DFT relaxations using randomly sampled predicted structures from test

Table 4-6: Generalization performance of DOGSS using different data splitting methods: Original split (80/10/10), Leave-"Ni"-out (leaving all structures containing Ni out of training sets and testing DOGSS only on the remainder), and Leave-"Zn" and "Co"-out (leaving all structures containing Zn and Co out of training sets and testing DOGSS only on the remainder). We selected Ni, Zn, and Co elements in order to maintain similar train/val/test sizes over different data split.

Dataset	$L_{ m s}$	urf	$L_{ m a}$	ds	$L_{ m total}$		
	Initial	Pred.	Initial	Pred.	Initial	Pred.	
Bare Suface							
- Original split	0.1377	0.0635	N/A	N/A	0.1377	0.0635	
- Leave-"Ni"-out	0.1200	0.0740	N/A	N/A	0.1200	0.0740	
- Leave-"Zn" and "Co"-out	0.1363	0.0937	N/A	N/A	0.1363	0.0937	
H-adsorption							
- Original split	0.0855	0.0508	0.7841	0.3301	0.1358	0.0709	
- Leave-"Ni"-out	0.0710	0.0638	0.7800	0.4321	0.1174	0.0879	
- Leave-"Zn" and "Co"-out	0.0810	0.0824	0.8040	0.4726	0.1335	0.1108	
CO-adsorption							
- Original split	0.0844	0.0548	0.7286	0.3237	0.1706	0.0908	
- Leave-"Ni"-out	0.0919	0.0740	0.7303	0.3946	0.1719	0.1142	
- Leave-"Zn" and "Co"-out	0.0902	0.0836	0.6911	0.4408	0.1756	0.1343	

sets of the three different datasets. We used the same DFT settings as the initial structures were relaxed for a direct comparison of the number of steps in DFT relaxations until the maximum inter-atomic force reaches below 0.05 eV/Å. Table 4-7 summarizes the results from the test DFT simulations, and example structures are visualized in Figure 4-5. For bare surfaces, it takes about 10.1 ionic relaxation steps

to relax the predicted structures, which are about 30% fewer steps compared to the relaxation of the initial surfaces. For relaxations of H and CO adsorbates on the metal surfaces, DOGSS could reduce about 45% and 52% of the total DFT relaxation steps respectively. We observe that bringing randomly placed adsorbate in the initial structure to the adsorption site during DFT relaxations is a time-consuming process as the relaxations of adsorbate-covered surfaces are significantly more expensive than the relaxations of bare surfaces. Therefore, approximating ground-state positions of adsorbates prior to the relaxations could effectively reduce the expensive computational cost of DFT for investigating surface science that involves adsorptions.

 Table 4-7: The average number of steps in DFT relaxations of randomly sampled

 initial and DOGSS-predicted structures from test sets in the three different datasets.

Dataset	Sample Size	Avg. DFT Steps				
		Initial	Pred.			
Bare Surfaces	601	14.3	10.1			
H adsorption	480	27.3	15.0			
CO adsorption	540	104.8	50.8			

4.7 Conclusion

In this Letter, we present a deep learning framework that incorporates differentiable optimization within graph neural networks to approximate ground-state structures. We demonstrate that the framework provides reliable ground-state structure predictions of diverse inorganic multicomponent surfaces with and without the presence of adsorbates (H and CO) and reduces the computational cost of DFT by allowing relaxations of the predicted structures that require fewer optimization steps. Fur-



Figure 4-5: Example structures in CO adsorption dataset. (Top) CO adsorption on Fe and As alloy surface and (bottom) CO adsorption on Ni and Si alloy surface.

ther, the framework is flexible so that one can easily modify the molecular graph representations, neural network architecture, embedded optimization method, and type of potential energy to investigate different systems and to result in different prediction performance. This method is complementary to other published methods to accelerate DFT relaxations using on-the-fly neural network potentials as it learns a better initial guess for the process.

Chapter 5

Deep reinforcement learning for predicting kinetic pathways to surface reconstruction in a ternary alloy

This work originally appeared as: Junwoong Yoon*, Zhonglin Cao*, Rajesh K Raju*, Yuyang Wang, Robert Burnley, Andrew J. Gellman, Amir Barati Farimani, and Zachary W. Ulissi. Differentiable Optimization for the Prediction of Ground State Structures (DOGSS). Mach. Learn.: Sci. Technol. 2. 045018 (2021). It has been edited to include the supporting information in-line.

5.1 Abstract

The majority of computational catalyst design focuses on the screening of material components and alloy composition to optimize selectivity and activity for a given reaction. However, predicting the metastability of the alloy catalyst surface at realistic operating conditions requires an extensive sampling of possible surface reconstructions and their associated kinetic pathways. We present CatGym, a deep reinforcement learning (DRL) environment for predicting the thermal surface reconstruction pathways and their associated kinetic barriers in crystalline solids under reaction conditions. The DRL agent iteratively changes the positions of atoms in the near-surface region to generate kinetic pathways to accessible local minima involving changes in the surface compositions. We showcase our agent by predicting the surface reconstruction pathways of a ternary Ni₃Pd₃Au₂(111) alloy catalyst. Our results show that the DRL agent can not only explore more diverse surface compositions than the conventional minima hopping method, but also generate the kinetic surface reconstruction pathways. We further demonstrate that the kinetic pathway to a global minimum energy surface composition and its associated transition state predicted by our agent is in good agreement with the minimum energy path predicted by nudged elastic band calculations.

5.2 Introduction

The performance of a heterogeneous catalyst depends on the catalyst's surface composition and structure. The discovery of novel robust catalysts for a given reaction is often achieved via surface engineering of existing catalysts. Significant attention has been drawn towards the design and development of alloy catalysts, as the synergistic effects of alloying two or more metals can provide catalytic activity, selectivity, and stability superior to their pure component counterparts.[23, 60, 167] Furthermore, alloying noble metal catalysts (Pt, Pd, Ag, Au, etc.) with low cost, highly abundant metals (Ni, Cu, Sn, Co, etc.) can function to reduce the catalyst cost in scaling up industrial level operations. The design of these catalysts is complicated by dynamic transformations of the multi-metallic atomic environment. Often, the reconstruction
of a catalyst's surface causes the real surface structure to differ from that predicted simply by cleaving the bulk crystals along a given plane. In alloy catalysts, segregation can result in the surface enrichment or depletion of one or more components in an effort to minimize the surface free energy and reduce lattice strain. Lateral rearrangements can cause a change in the surface layer periodicity. The simultaneous compositional evolution and lateral rearrangement of an alloy catalyst's surface play a critical role in determining the performance of a catalyst.

Predictions of the surface composition of alloy catalysts commonly employ thermodynamic arguments, however, kinetic information is a prerequisite for predicting metastable states that have no kinetically-feasible nearby local minima due to the high kinetic barriers or high transition state energies. Under reaction conditions, changes in temperature or pressure can cause a change in the stability of the initial thermodynamic equilibrium. The ability of a catalyst to realize the new equilibrium or new local minimum surface composition under these conditions is dictated by the barriers along the kinetic reconstruction pathways. A study of the pathway that describes the evolution of the multi-metallic catalyst surface at reaction conditions is thus necessary to understand the catalytic mechanism at the atomic scale as well as to tune the catalyst's activity and selectivity by controlling the surface reconstruction processes.

The structural transformation of catalyst surfaces has been widely studied from in-situ spectroscopic techniques such as in-situ X-ray diffraction (XRD), in-situ Xray absorption spectroscopy (XAS), in-situ x-ray photoelectron spectroscopy (XPS) and in-situ infrared spectroscopy.[21, 75, 91, 173, 196, 199, 219, 234] Other experimental techniques include in-situ scanning probe microscopy (SPM), in-situ scanning tunneling microscopy (STM), in-situ atomic force microscopy (AFM), and transmission electron microscopy (TEM).[150, 159, 164, 189, 190, 194, 197, 198, 215] Due to the prohibitive size of the design space, physical limitations, and experimental costs many studies have focused on few discrete temperatures and alloy compositions, limiting the size of the design space explorable by these techniques.

Rapid growth in the design of high performance supercomputing resources coupled with the development of modern computational methodologies primarily based on density functional theory (DFT) now complement the traditional trial and error experimental approaches to the discovery of novel materials for various catalytic applications. Computational approaches can accelerate the catalyst discovery efforts by screening component and composition spaces to eliminate material candidates with low activity and selectivity before performing experimental measurements. However, the search for both stable and high-efficient multi-metallic catalysts for a specific reaction through computational catalytic design is still a Herculean task dramatically increasing in complexity and computational cost with the number of alloy components. Moreover, material screening calculations relying on equilibrium properties ignore the possibility of surface reconstruction. The prediction of a specific surface reconstruction is a complicated process involving many structural evaluations along a hypothesized reconstruction pathway. Such an exercise may require hundreds to thousands of cpu-hours to evaluate only a single pathway of all possible surface reconstructions. Therefore, a robust exploration of the reconstruction pathways requires a technique capable of creatively generating reconstruction pathways because most of the current computational efforts in heterogeneous catalysis avoid the comprehensive exploration of reconstruction pathways required to identify metastable states.

With recent advances in machine learning (ML), researchers have successfully applied ML techniques in the prediction of materials properties [12, 124, 137, 204, 223], relaxed structures [233], and alloy properties [25, 94, 228], but have yet to predict phenomena on the timescales of surface reconstruction pathways needed to identify metastable states. This is still an extremely challenging problem for the supervised ML models that are commonly used for the direct prediction of the struc-

ture/property relationships as the identification of the metastable structures require not only the property labels, but also the information about various kinetic pathways that lead to the changes in the alloy catalyst structure and properties. ML has also established atomistic force field methods using a potential energy surface (PES) fitted to a set of DFT samples in the materials search space.[19, 45, 46, 166, 177, 178] However, the residual force errors in the fitted force field often result in structures differing from the true equilibrium states, or worse, lead to non-physical configurations. Furthermore, generating a huge amount of DFT samples throughout the energy pathways for accurate force field predictions while considering all possible catalyst surface reconstructions at a specific reaction condition is infeasible. Despite the considerable excitement about these ML methods, identifying metastable catalyst surfaces still remains untapped due to its highly complex design space.

To deal with the complexity issue, an alternative approach called deep reinforcement learning (DRL) that is capable of traversing the compositional space more efficiently has been employed. Reinforcement learning (RL) is a subfield of machine learning where a decision maker or an agent learns strategies to solve an optimization problem by iteratively interacting with an environment. DRL involves applying deep neural network in RL frameworks to generalize complex problems. Recent applications of RL in optimizing molecular structures [39, 106, 143, 184] and chemical reaction pathways [238] have demonstrated its ability to tackle complex molecular design problems.

In this work, we introduce a new framework for predicting the kinetic energy pathways to the catalyst surface segregation under reaction conditions by combining DRL with the domain knowledge of catalysis. For a given catalyst surface, a DRL agent repeatedly attempts to explore nearby local minima with changed surface compositions, and the exploration process is evaluated based on the properties of the local minima and the transition states the agent discovered. In this way, the agent iteratively builds up the knowledge that can generate the kinetic pathways to the accessible surface transformations at a given reaction temperature, and thus it can be further used to identify metastable catalyst surfaces. The new framework is called CatGym, the first open RL environment developed for reconstruction kinetics.

We showcase our framework with a demonstration of surface segregation kinetics for a ternary $Ni_3Pd_3Au_2(111)$ alloy catalyst. Pd and Pd-based catalysts are of great industrial importance, being used as hydrogen purification membranes and as a catalyst for the hydrogen evolution reaction (HER), oxygen reduction reaction (ORR), and ethanol oxidation reaction (EOR). In particular, NiPdAu catalysts are useful in hydrogen generation for fuel cells via the catalytic dehydrogenation of formic acid, as well as in the EOR in direct ethanol fuel cells. [50, 213] Researchers have found that the addition of one or more components can increase the activity of Pd-based catalysts, improve resistance to poisoning, and prevent the hydrogen embrittlement brought on by the metal/hydride phase transition. [92, 130, 134, 160, 200, 201, 236] Ternary Pd-based alloys have attracted significant attention as the addition of two components allows greater tuning of electrical and structural properties as compared to their binary counterparts. By predicting surface reconstructions and their associated activation barriers, CatGym can aid in the discovery of complex and novel multicomponent catalysts.

5.3 Methods

5.3.1 Reinforcement Learning Background

RL is a class of artificial intelligence which aims at training an artificial agent by actively interacting with the environment[191]. Following the Markov Decision Process (MDP), at each timestep t, the RL agent chooses an action a_t from the action space \mathcal{A} , given the current state s_t from state space \mathcal{S} . The environment returns

a reward, r_t , as the feedback for the state action pair (s_t, a_t) . The policy π_t is a mapping from the state space S to the action space A, namely how the agent decides the action at each timestep. An episode will be terminated once the agent achieves the goal state s_g or the number of steps reaches the set maximal. The goal of the RL agent is to learn the policy which maximizes the cumulative return $R_t = \sum_t^T \gamma^{t-1} r_t$, where γ is the discount factor which adjust the importance of future reward and T is the length of the episode. During training, the RL agent optimizes the policy π_t sequentially as it actively interacts with the environment. Recently, combination of RL and Deep Learning has made it possible to tackle the problems in very high dimensional space. One of the main drawbacks of RL is its inability to resolve the curse of dimensionality. Deep neural networks, can be used to approximate the value function in reinforcement learning. Benefiting from both reinforcement learning and deep learning, DRL has seen astonishing advancement in various fields, including superhuman-level video game control[144, 145], GO playing[183], robotics control[126, 175, 176], and chemical compound design[168, 212, 239].



Figure 5-1: (a) Overview of CatGym framework. At each timestep t, the state space s_t consists of features of surface structure c_t , including its current energy, fingerprint, and atom coordinates. The state space along with reward r_t , which is calculated by critic network based on current energy, are fed into the actor network, and the actor network decides which action to take. The action will modify c_t to new surface structure c_{t+1} . This process repeats until the episode is over. (b) Four types of action the actor can choose and their corresponding effects on the energy. If the agent choose to move an atom, it has to further pick which atom to move and the distance (within the range of -0.1Å to 0.1Å) in the x, y and z direction. (c) Evolution of the reconstruction pathways to a minimum configuration generated by DRL agent trained with CatGym. ΔE represents the relative energy of the catalyst structure at each timestep with respect to the initial state energy.

5.3.2 Actor-Critic TRPO

To efficiently optimize the policy in DRL, an actor-critic algorithm [79, 146, 192] is proposed which contains two networks: an actor network which determines actions based on a given state embedding, and a critic network which evaluates how good the chosen action is (by estimating the cumulative return). In our work, we train an actor-critic network within the customized framework (Figure 5.3.1(a)). We implement trust region policy optimization (TRPO)[175], an actor-critic network which has been prevalent in various RL applications. TRPO is effective for optimizing neural network based policies and guarantees monotonic improvement on each update. To this end, we utilize parameterized action space[53, 151], which consists of hierarchical sub-actor networks to decompose the action space.

5.3.3 State Representation

The state space S is a set of state observations that describes quantified abstraction of the perceived environment. Training an agent that discovers possible surface segregation requires representations for each of the atoms in the near-surface region that can capture dynamically changing local chemical environment. To meet such a requirement, we employ atomic symmetry functions[17] and atomic positions of the free atoms near the surface to efficiently encode the elemental and spatial atomic information. In addition, we provide properties of the surface structure such as energies and forces, so that the agent can understand how the changes in the configurations affect their properties. Finally, we also encode a binary vector that tells whether the agent has found transition states in each episode. These five state observations are then processed through the distinct multi-layer perceptrons and combined to create an embedding of states that are fed to the actor network.

5.3.4 Action Space

The action space \mathcal{A} is a set of actions that the agent can perform to interact with the environment. In general, the action space is designed to be either discrete or continuous based on the reinforcement learning problem and the algorithm used for solving it. In this work, we formulate a hierarchical hybrid discrete-continuous action space (Figure 5.3.1(b)), in which each action $a_t = (a_t^{(1)}, a_t^{(2)}, a_t^{(3)})$ comprises

- $a^{(1)} \in \{1, 2, 3, 4\}$ selects an action from the four different possible actions,
 - 1. moving individual atoms by a fixed distance towards a direction,
 - 2. finding a nearby transition state with a saddle point solver,
 - 3. triggering a local energy relaxation to minimize the current structure,
 - 4. performing short molecular dynamics simulations to perturb the free atoms at a specified reaction temperature.

- a⁽²⁾ ∈ {1,...,N} chooses one of the N free atoms on the top two surface layers to move if a⁽¹⁾ = 1.
- $a^{(3)} \in \{(x, y, z) | -0.1 \text{\AA} \le x, y, z \le 0.1 \text{\AA}\}$ that specifies the distances to move the free atom chosen by $a^{(2)}$ in x, y, and z directions, respectively.

We use Sella[90] as a saddle point solver for finding nearby first order saddle points $(a^{(1)} = 2)$. The energy relaxation uses the BFGS optimizer implemented in Atomic Simulation Environment (ASE) [?] to find nearby local minima $(a^{(1)} = 3)$, and the MD simulations $(a^{(1)} = 4)$ use Langevin dynamics at constant temperature implemented in ASE as well.

At each timestep, the agent decides how to change the structure of the catalyst surface in order to eventually generate a sequence of actions that leads to the exploration of new local minima with different surface compositions. Figure 5.3.1(c) illustrates how the sequence of actions can become the reconstruction pathway to a specific minimum configuration.

We note that distinct configurations at different timesteps may end up with the same transition states or the same local minima after performing the saddle point solver or the relaxation. The relative energies (ΔE) between timesteps 0 and 40 in "Pathway 1" plot in Figure 5.3.1(c) shows that several saddle point search (orange) and relaxation (green) actions lead to the local minima previously visited. This greatly slows down the exploration process by repeatedly leading the agent to the same states. To mitigate the issue, we introduce short MD simulations in our action space so that the agent can more effectively escape from the current minimum or transition state and traverse the potential energy surface. Long enough MD simulations alone might be able to explore new minimum states, but the limited time scale and computing resources make it infeasible. Instead, the agent in our method can decide when one of the other actions is better to explore nearby local minima by learning from its interaction with the CatGym environment.

5.3.5 Reward Function

The reward function is designed on the basis of chemical properties of the surface structure. If the agent identifies a transition state at any time during each episode, the agent is rewarded by,

$$r(s_t, a_t) = \frac{1}{\Delta E_{\text{trans}}},\tag{5.1}$$

where ΔE_{trans} is the relative energy of the transition state to the energy of the initial state. In addition, after the completion of each episode (t = T), the agent is rewarded based on whether it discovers a pathway to a state with a changed surface composition. When the agent successfully observes kinetically feasible surface segregation from the initial surface, it is rewarded by,

$$r(s_T, a_T) = \exp\left(-\frac{\Delta E}{k_B T_r}\right),$$
 (5.2)

where k_B denotes the Boltzmann constant, T_r is a reaction temperature, and ΔE is the potential energy of the final (t = T) state relative to the initial state (t = 0). These rewards aim at encouraging the agent to explore both the lower transition states and the nearby local minima resulted from surface segregation. The traversability of the transition state is also taken into consideration when evaluating the generated pathway. The higher the transition state energy is, the more difficult the surface segregation to be realized at a given temperature. Therefore, when the agent yields a high energy states that exceeds a predefined upper energy bound $(3k_BT_r$ in our case), the episode is terminated and it is penalized by,

$$r(s_T, a_T) = \frac{\Delta E}{k_B T_r}.$$
(5.3)

Intuitively, the agent is instructed to find a reconstruction pathway, ideally one leads to a new state with thermodynamically more stable surface composition ($\Delta E < 0$) through a low energy barrier (ΔE_{trans}). The potential energies can be calculated using accurate quantum chemistry simulations such as DFT, however, evaluating a large number of trial configurations is computationally expensive. Instead, we avoid this cost by employing fast Effective Medium Theory (EMT) [98, 99, 193] to estimate the energies.

5.3.6 Experiments

To demonstrate the capability of our proposed DRL framework, we conduct experiments to generate surface reconstruction pathways for a $2 \times 2 \times 4$ Ni₃Pd₃Au₂(111) alloy catalyst at a specific reaction temperature (1200 K). The top surface layer in the initial Ni-Pd-Au catalyst has a Ni:Pd:Au composition of 1:2:1 and the second layer has a composition of 2:1:1 corresponding to a total of 3 Ni, 3 Pd, and 2 Au in 8 lattice positions in the unit cell as shown in Figure 5-3(a). Atoms in the top two layers are free to move, while atoms in the bottom two layers are considered as bulk atoms and are fixed.

We set up the DRL environment under the OpenAI Gym framework [27] and use Tensorforce [116] DRL package to run the experiments. We utilize parallel environment execution to perform multiple experiments running with the same initial conditions. These parallel experiments share the agent and all other model parameters. In each episode, the agent is asked to generate a sequence of at most 500 actions with the aim of exploring the nearby local and global minima with changed surface compositions while finding the minimum energy pathways to these minima. We utilize the EMT potentials for energy and force evaluations for all kinds of actions.

5.3.7 Surface Optimization Baseline

We use a brute force method and a minima hopping (MH) simulation method [70] for locating local and global minimum surface configurations for the Ni-Pd-Au ternary system. In the brute force method, we considered all the possible arrangements of lattice atoms in the unit cell by performing a distinct permutational method in the top two surface layers. Permutation of 3 Ni, 3 Pd, and 2 Au atoms in the 8 lattice positions in the top layers leads to a total of 560 distinct arrangements. We assume that this approach would span all possible local minima that arises from the movement of free atoms in the top two surface layers.

MH is an efficient search method employed for locating the global minimum for systems with highly complex PES. For instance, MH method has been applied to determine the structure of the reconstructed chalcopyrite surfaces [202] and metastable decorate borospherene B_{40} [172]. The MH method avoids revisiting known parts of the configuration space by utilizing a feedback mechanism based on simulation history which in turn accelerates and enforces the exploration of the new regions in the configurational space. The MH algorithm consists of an inner part that performs the moves on the PES employing molecular dynamics (MD) followed by the relaxation of the current minimum and an outer part which determines the acceptance or rejection of a new minimum.

The entire MH simulation is controlled by five parameters, $\alpha_1, \alpha_2, \beta_1, \beta_2, \beta_3$. Parameters α s control how rapidly E_{diff} is increased or decreased and β s will determine the extend to which kinetic energy changed via changing the temperature of the simulation. We used the optimal values from the original paper[70] introducing MH $(\beta_1 = \beta_2 = 1/\beta_3 = 1.05 \text{ and } \alpha = 1/\alpha_2 = 1.02)$ as it has been proposed that with these parameters, MH simulations never failed to find the minimum. We used a initial E_{diff} value of 0.5 eV and a temperature of 200 K and the simulation ran for a total of 2500 steps. The changes in the potential energy during MD and EMT steps as well as the changes in the E_{diff} and temperatures are shown in Figure 5-2. We found the global minimum configuration at step number 1655 with an energy value of 3.855 eV.



Figure 5-2: The potential energy changes and the changes in the E_{diff} and temperatures during MH simulation

5.4 Results and Discussion

5.4.1 Baseline Surface Configurations



Figure 5-3: Configurations of the Ni(green)-Pd(blue)-Au(yellow) ternary catalyst. (a) Initial catalyst structure, (b)-(e) global minimum, second and third lowest energy configurations, and global maximum configuration obtained via brute-force method, respectively; and (f) global minimum configuration obtained via minima hopping (MH) method. Atoms in the bottom two layers with crosses represent fixed atoms that are constrained from moving.

Using the brute force approach, we have found the global minimum at a relative energy of -0.23 eV with respect to the initial configuration. The minimum has a 0:2:2 composition of Ni, Pd and Au on the top surface layer in the unit cell as shown in Figure 5-3(b). Moreover, several configurations with the same compositions; however, with slightly different arrangement of atoms in the top two layers were identified within the relative energy range of -0.23 to -0.19 eV and can be considered as a group of global minima energy configurations. Followed by this, we have identified a group of the second-lowest energy configurations with Ni:Pd:Au composition of 0:3:1 in the top surface layer (Figure 5-3(c)) with different configurational orientations in the relative energy range of -0.18 to -0.16 eV. A group of the third-lowest energy configurations starts at -0.11 to -0.07 eV energy range and consists of Ni:Pd:Au composition of 1:1:2 on the surface layer as illustrated in Figure 5-3(d). The minima configurations with negative relative energies consistently find surface enrichment of Au and/or Pd with respect to the initial configuration. This observation can be rationalized by arguments for segregation primarily driven by differences in the pure component surface energies. Estimates of the surface energies (γ_i) for each of Au, Pd, and Ni elements indicate $\gamma_{Au} < \gamma_{Pd} < \gamma_{Ni}$ [209]. Thus we find a global minimum in which all free (non-fixed) Au atoms segregate to the surface and a nearby local minimum in which Pd fully replaces Ni in the surface. These findings are qualitatively in agreement with experimental studies of segregation in binary alloys: Au is favored over Pd [42, 229, 230], Pd is favored over Ni [2, 47, 187], and Au is favored over Ni [29, 110, 217] in the surface layer. To reaffirm, we also located the configuration with the highest energy (i.e global maximum) which has a Ni:Pd:Au composition of 3:1:0 in the surface layer (Figure 5-3(e)) with an relative energy of 0.81 eV. In fact, the global maximum configuration corresponds to the reversal of the top two layers in the global minimum configuration.

For a direct comparison with the our DRL model, we have used MH as a standard baseline approach for the local and global minimum search. We found the global minimum configuration with an relative energy of -0.23 eV (Figure 5-3 (f)) in the MH simulation. The global minimum configuration obtained via MH simulation shows similar segregation of Pd and Au onto the top layer with 0:2:2 composition of Ni, Pd, and Au on the surface layer. However, they differ in their configurational orientations of the atoms in the top two layers. Like brute-force permutation method, MH method also found many similar configurations of 0:2:2 composition of Ni, Pd, and Au in the top layer with different arrangements with similar energies and can be collectively grouped into a cluster of global minimum configurations.

5.4.2 DRL Training Summary



Figure 5-4: Average reward vs. episode across 10 different seeds. Red line represents the average reward, and the blue shadow represents the standard deviation.

The DRL agent was trained with 10 different seeds by interacting with the Cat-Gym environment. In Figure 5-4, the averaged episodic reward collected during the training process is represented by the red curve, and the standard deviation of the rewards in 10 training is represented by blue shadow. The reward at the initial stage of training (<5000 episodes) is low and the policy is noisy because the agent is instructed to explore and has not yet learned to discover surface reconstruction pathway without exceeding the upper energy bound that results in negative rewards. After approximately 8000 episodes of training, the agent starts to receive stable positive reward between 0 and 1. The converged positive reward indicate that the agent has learned a stable policy to discover surface reconstruction pathways that lead to different surface compositions while ensuring the traversability of transition state by not exceeding the upper energy bound.



Figure 5-5: Projection of Ni-Pd-Au configurations identified by (a) MH and (b) DRL method in the reduced dimensional space using the first two principal components that capture the most structural variance in the all possible surface arrangements. Numbers in red indicate unique clusters determined by Kmeans. Boundaries between the clusters are drawn using sold lines. (c) Example Ni(green)-Pd(blue)-Au(yellow) configurations for each of the clusters explored by DRL method, clusters 0, 1, 2, 5, 6, 7, 8, 9, and 10. Average relative energies (ΔE) of transition states (blue) and minima (black) for each of the clusters are labeled below the cluster numbers (red).

5.4.3 Surface Configuration Exploration

First, we performed Principal Component Analysis (PCA) on the minima configurations generated by the brute-force permutation method to determine the reduced dimensional space that captures the most structural variance in the all possible minima configurations. Each of the configurations is represented by a combined atomic fingerprints of the atoms in the top two layers. We then projected the minima configurations onto the reduced dimensional space using the first two principal components as shown in Figure 5-6. On the same plot, we also drew the boundaries of clusters determined by the K-means algorithm. We used 16 clusters to partition the configurations plotted in the reduced dimensional space based on their structural representations and properties.

To compare the diversity of the Ni-Pd-Au catalyst surface configurations of local minima explored by our DRL method and the baseline MH methods, we projected these configurations onto the pre-determined reduced dimensional space mentioned above. Figure 5-5(a) shows that the MH method could explore surface configurations corresponding to only two different minima clusters, cluster 2 and 5, where Ni:Pd:Au compositions of 0:2:2 and 1:1:2 on the top surface layer were found, respectively. Based on the brute-force permutational analysis, cluster 2 is one of the global minima clusters (Figure 5-3(a)) and cluster 5 is one of the clusters with configurations in the third-lowest energy group (Figure 5-3(c)). On the other hand, CatGym could explore more diverse surface configurations found in eight different minima clusters as shown in Figure 5-5(b). Figure 5-5(c) shows example Ni-Pd-Au configurations in the eight different clusters visited by the DRL agent with diverse Ni:Pd:Au compositions: 0:2:2, 1:1:2, 0:3:1, 1:2:1, 2:1:1, 1:3:0, and 2:2:0. We note that the same Ni:Pd:Au composition can be found in different nearby clusters with different surface orientations, such as clusters 8 and 9 with 2:2:0 compositions, and clusters 6 and 10 with 2:1:1 compositions. We also notice that our DRL agent was unable to explore other local minima that are not nearby. In Figure 5-5(b), clusters distant from the initial structure in cluster 0 were not identified by the DRL method in this experiment. Those undiscovered clusters include different surface orientations



Figure 5-6: Projection of all possible Ni-Pd-Au surface configurations (total 560 configurations) identified by brute-force permutation method onto the reduced dimensional space using PCA. Numbers in red represent distinct clusters determined by K-means clustering algorithm, and their boundaries are drawn using solid lines in the PCA plot. Example Ni(green)-Pd(blue)-Au(gold) surface configuration for each of the clusters are also shown.

of global minima (cluster 3), second-lowest minima (cluster 4), global maxima (cluster 15), and several other configurations with surface Ni:Pd:Au compositions of 2:1:1 (cluster 11 and 12), 2:0:2 (cluster 13), 3:0:1 (cluster 14) based on Figure S2.

5.4.4 Reconstruction Pathway and Transition State

In addition to the diverse local minima configurations, our DRL method can generate reconstruction pathways to the explored minima with different surface compositions and can provide the structure and properties of not only the minima configurations, but also the transition states, thereby enabling the identification of metastable catalyst surfaces. We note that MH cannot determine the metastability of the catalyst surface because this method does not provide kinetic information about the pathways to the minima it explores. A reconstruction pathway from the initial state to the global minimum state in Figure 5-8(a) demonstrates that the DRL method can creatively generate a sequence of actions to explore the global minimum. The transition state energy or the highest energy barrier in this pathway was determined as 1.56 eV above the energy of the initial state. In Figure 5-5(c), using the same analysis, we show the average transition state energies for all local minima within each of the clusters in the PCA plot (Figure 5-5(b)). We noticed that the transition state energies systematically increase as the relative energies of their final minimum states increase. For a cluster of global minimum configurations (cluster 2), the average transition state energy is determined as 1.56 eV while the average transition state energies for clusters of the highest energy states (cluster 8 and 9) explored by our DRL method are in a range of 1.95 eV to 1.99 eV. More example pathways to other local minima can be found in Figure 5-7.

5.4.5 Transition State Verification

We further performed nudged elastic band (NEB) [85, 86] calculations to verify the structure and energy of the transition state in the path to the global minimum. NEB is a method for finding saddle points and minimum energy paths between known initial and final states, which are in this work the same initial structure and the global minimum state discovered by our DRL method. We extracted the intermediate con-



Figure 5-7: Example reconstruction pathways to local minima generated by the DRL agent. Ni(green)-Pd(blue)-Au(gold) configurations of the transition state and the minimum for each of the pathways are also illustrated next to the pathway plot. Energy values in this figure (ΔE) represent the relative energies of the each of the configuration with respect to the initial state energy.



Figure 5-8: (a) An example energy pathway to a global minimum developed by DRL method. Each data point represents the relative energy ΔE of an Ni-Pd-Au configuration generated after taking a certain action at each timestep. (b) A minimum energy pathway, created by NEB, to the same global minimum. (c) Ni(green)-Pd(blue)-Au(yellow) configurations of the initial state, transitions states built by DRL method and NEB, and the global minimum.

figurations from the reconstruction pathway developed by the DRL method (configurations at timesteps from 1 to 8 in Figure 5-8(a)), and used them as the intermediate configurations along the minimum energy path in the NEB calculation. In Figure 5-8(a), the agent discovered a transition state to the global minimum at timestep 7, and then performed a local relaxation to achieve the global minimum at timestep 9. Figure 5-8(b) shows the minimum energy path constructed by NEB between the same initial state and the same global minimum state. The energy barrier in this minimum path is 1.54 eV, which is very close to the one (1.56 eV) estimated by the DRL method. Further, Figure 5-8(c) visually verifies that the structures of the transition states identified by both the DRL method and NEB are close to each other. Both DRL and NEB find transition states in which one of the Pd (blue) atoms vacate the surface layer. This vacancy then facilitates the swapping of a surface Ni (green) and a subsurface Au (yellow) atom, followed by the return of Pd atom down to the surface to fill the vacancy, thereby resulting in the global minimum configuration with Ni:Pd:Au composition of 0:2:2 for the top surface layer.

5.5 Conclusion

In this work, we present CatGym DRL environment for studying kinetic arguments of catalyst surface segregation under reaction conditions. We aim at exploring possible surface segregation phenomena and the associated transition states to address the challenge in predicting catalyst metastability. For a given catalyst surface, the DRL agent iteratively alters the positions of atoms and learns strategies for generating kinetic pathways to nearby local minima with different surface compositions resulted from surface segregation. Trained with the TRPO algorithm and a ternary $Ni_3Pd_3Au_2(111)$ alloy catalyst, the agent in our CatGym environment not only explores more diverse local and global minima configurations compared to the baseline MH method, but also generates kinetic pathways to those configurations. We also verify that the reconstruction pathway to the global minimum surface configuration generated by the DRL agent is in a good agreement with the minimum energy path calculated using NEB. CatGym is the first general DRL approach towards the design of metastable catalysts under reaction conditions. This approach can be extended to other systems of interest possibly containing different catalyst surfaces with varying unit cell sizes, metals, oxides, and adsorbates with only a few minor modifications.

Chapter 6

Conclusion and Future Work

In the first part of this thesis, we provided an example of how the theoretical models and computer simulations, which are the second and third paradigms of materials science respectively, can be combined into a single framework to investigate interfacial phenomena at solid interfaces. To design and modify interfaces of solid materials, we needed a new tool to predict and understand interactions between the hard interfaces and surfactants because simulation techniques developed for fluid-fluid interfaces often are not applicable to solid interfaces as surfactant adsorptions behave non-ideally at solid interfaces. This is an important problem to improve applications of interfacial phenomena in energy science, catalysis, and advanced materials as the solid interface modification via surfactant adsorption can alter interfacial properties and affect the performance of the interfacial applications. In chapter 2, we modified an existing MTT model developed for fluid interfaces to make it suitable for hard interfaces and developed MD simulation techniques to directly determine the necessary thermodynamic parameters for arbitrary small molecules. The combined MD/MTT framework was capable of relating bulk and surface concentrations of surfactants at the solid interfaces, thereby overcoming the computational limitation of atomistic simulations for the determination of adsorption isotherms. The framework was also able to capture the relationship between the adsorption isotherms and the structural behavior of surfactants at solid interfaces. We further applied this framework to a well-studied SDS adsorption at a graphene surface system as a verification, and the structural transitions of SDS developed by the framework showed a good agreement with experimental observations. We believe this framework can serve as a tool for the design of both solid interfaces and surfactant molecules for various applications.

In the future, we hope to improve the framework to more accurately model the surfactant-interface systems with a presence of a larger number of surfactant molecules as the current MD simulation technique developed for this framework is limited by the size of the interface and the number of surfactants in the system. Another approach to improve the framework is the utilization of machine learning or deep learning models for the determination of MTT parameters. This approach may require initial efforts to build a surfactant-interface database using computer simulations for the model training, but the success of this work will enable a highthroughput screening against a large number of potential solid interfaces and surfactant molecules for targeted applications.

The second part, which is the main contribution of this thesis, focused on the development of deep learning methods for solid materials and heterogeneous catalyst design problems. High-throughput screening of heterogeneous catalysts requires a huge number of property calculations using expensive quantum mechanical simulations such as DFT. Binding energies are typically used as descriptors to predict important catalytic properties, and in Chapter 3 we presented a graph-based deep learning method for the prediction of the binding energies. We modified existing graph representations previously developed for representing solid materials to encode neighbor information based on Voronoi polyhedra. The graph neural network method was applied to the graph representation, then it automatically learned the optimal representation for predicting the target properties by using multiple convo-

lution layers. As a result, the network predicted binding energies of H and CO on diverse heterogeneous catalyst surfaces with higher accuracies compared to a conventional machine learning model that used human-engineered feature representation of the adsorption systems when there are sufficient numbers of training data. Further, the graph neural network framework in our work was able to extract individual atomic contributions to the predicted binding energies thereby rationalizing that the contribution decreases as atoms get farther from the binding site. We believe the framework can effectively rule out unpromising candidate catalyst surfaces by providing reliable binding energy predictions, and thus can significantly reduce the search space for high-throughput catalyst screening.

In Chapter 4, we broadened the applicability of the graph neural network presented in Chapter 3. Although the search space for high-throughput catalyst screening can be reduced by the use of machine learning methods as demonstrated in Chapter 3, accurate property calculations for the final design or verification processes using expensive electronic structure methods are inevitable. We thus incorporated a differentiable optimization into the graph neural network to build a framework for accelerating the electronic structure method computations by directly predicting the ground state or relaxed structure of heterogeneous adsorption systems. The integrated framework estimated the harmonic force field parameters using the graph neural network method and minimized the estimated harmonic PES using a gradient descent algorithm to generate the relaxed structures. The framework was capable of approximating ground state structures of the diverse inorganic catalyst surfaces with and without the presence of adsorbates. We further demonstrated that by relaxing the predicted structures instead of the initial structures, the computational cost of the DFT relaxation could be reduced by a factor of two. We believe that this deep learning framework can be used as a preprocessing tool that provides any electronic structure method with a starting configuration near the ground state, thereby accelerating the relaxation and further facilitating the high-throughput screening process.

In future work, we would like to advance the graph convolutional neural network framework in order to improve the model performance in both the property predictions covered in Chapter 3 and structure predictions from Chapter 4. For a better graph representation of the surface structures, we can try to incorporate more of physical and geometric properties into the features. A few existing approaches include labeling of surface atoms that are directly interacting with adsorbates[77] and encoding higher-order, many-body correlations such as 3-body interactions into the edge features[161]. These works suggest graph representation with more informative domain knowledge could help in improving model performance. On the model side, several advanced graph neural network architectures have been presented to propose advanced convolutional operations to extract better graph embeddings for improved model prediction accuracy. Modeling the directionality of center atoms to neighboring atoms[64] and also updating both node and edge features throughout the convolutional layers[65] could reinforce the ability to better capture the surface structure, thereby providing better embeddings for model predictions.

In Chapter 4, the simplest form of potential was used for the approximation of the ground-truth PES for simplicity. We would also like to consider more sophisticated or higher-order forms of potential to better approximate the ground-truth PES for better prediction of the ground state structure. One possible issue with the use of more complex forms of potential that needs to be addressed is a convergence issue during the energy minimization. Since they will likely construct non-convex PES which is hard to minimize using the conventional gradient descent method we used in Chapter 4. In order to enable the use of complex PES as an approximation of ground-truth PES, it will also be required to come up with advanced differentiable optimization methods.

Chapter 5 presented a DRL framework, CatGym, for identifying the metastability

of alloy catalyst surfaces under reaction conditions. Metastability is an important factor to consider for catalyst design because various reaction conditions such as temperature or pressure may cause a change in the surface structure and property via reconstruction. CatGym is the first general DRL approach to explore surface segregation and associated kinetic barriers to determine metastable catalyst surfaces. The DRL agent learned strategies for generating kinetic pathways to nearby local minima having different surface compositions by iteratively changing the positions of atoms and receiving rewards from the environment. The DRL framework was able to explore more diverse surface reconstruction pathways to local and global minima of an example ternary alloy catalyst compared to the baseline minima hopping method. We further demonstrated that the kinetic barriers of the reconstruction pathways are in good agreement with the minimum energy path determined by nudged elastic band method. This framework is also flexible to be extended to other systems of interest containing different catalyst surfaces with different unit cell sizes, metals, and adsorbates with only a few minor modifications.

In the future, we would like to extend this study to the identification of metastability of adsorbate-covered catalyst surfaces instead of bare surfaces as presented in Chapter 5. Since the stability of the alloy catalyst surface would change when the adsorbate atoms are present due to the adsorbate-induced surface reconstruction. Moreover, recently developed and presented reinforcement algorithms such as Multi-Agent training can also be implemented in an attempt to improve the training efficiency. In a Multi-Agent environment, multiple agents will cooperate with each other to reach the goal, or maximize the reward by optimizing their policy.[30] This training algorithm has been proven to be effective in recent DRL studies in different areas including transportation management[4], finance[121], communication networks[41].

Bibliography

- [1] Manthiram A., Yu X., and Wang S. Nat. Rev. Mater., 2:1 16, 2017. URL https://www.scopus.com/inward/record.uri?eid=2-s2.0-85045205827& partnerID=40&md5=e28b843242a0bbc40f39eeae87537dab. Cited by: 66. 1.1
- M. Abel, Y. Robach, J.-C. Bertolini, and L. Porte. STM comparative study of the Pd₈Ni₉₂(110) alloy surface and the Pd/Ni(110) surface alloy. Surface Science, 454-456:1-5, 2000. ISSN 0039-6028. doi: https://doi.org/10.1016/S0039-6028(00)00269-7. URL https://www.sciencedirect.com/science/article/pii/S0039602800002697. 5.4.1
- [3] Frank Abild-Pedersen, Jeff Greeley, Felix Studt, Jan Rossmeisl, TR Munter, Poul Georg Moses, Egill Skulason, Thomas Bligaard, and Jens Kehlet Nørskov. Scaling properties of adsorption energies for hydrogen-containing molecules on transition-metal surfaces. *Phys. Rev. Lett.*, 99(1):016105, 2007. 3.4, 3-2, 3.6
- [4] Jeffrey L Adler and Victor J Blue. A cooperative multi-agent transportation management and route guidance system. *Transportation Research Part C: Emerging Technologies*, 10(5):433-454, 2002. ISSN 0968-090X.
 doi: https://doi.org/10.1016/S0968-090X(02)00030-X. URL https://www. sciencedirect.com/science/article/pii/S0968090X0200030X. 6
- [5] Ankit Agrawal and Alok Choudhary. Perspective: Materials informatics and big data: Realization of the "fourth paradigm" of science in materials science.

APL Materials, 4(5):053208, 2016. doi: 10.1063/1.4946894. URL https://aip.scitation.org/doi/abs/10.1063/1.4946894. 1.2

- [6] Laith Alzubaidi, Jinglan Zhang, Amjad J. Humaidi, Ayad Al-Dujaili, Ye Duan, Omran Al-Shamma, J. Santamaría, Mohammed A. Fadhel, Muthana Al-Amidie, and Laith Farhan. Review of deep learning: concepts, cnn architectures, challenges, applications, future directions. *Journal of Big Data*, 8 (1):53, Mar 2021. ISSN 2196-1115. doi: 10.1186/s40537-021-00444-8. URL https://doi.org/10.1186/s40537-021-00444-8. 1.3
- [7] Mie Andersen, Sergey Levchenko, Matthias Scheffler, and Karsten Reuter. Beyond scaling relations for the description of catalytic materials. ACS Catal., 9: 2752–2759, 2019. 3.4, 3-2, 3.6
- [8] Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6):26–38, 2017. doi: 10.1109/MSP.2017.2743240. 1.3
- [9] Seoin Back, Juhyung Lim, Na-Young Kim, Yong-Hyun Kim, and Yousung Jung. Single-atom catalysts for co₂ electroreduction with significant activity and selectivity improvements. *Chem. Sci.*, 8(2):1090–1096, 2017. 3.2
- [10] Seoin Back, Kevin Tran, and Zachary Ulissi. Towards a design of active oxygen evolution catalysts: Insights from automated density functional theory calculations and machine learning. ACS Catal., 2019. https://doi.org/10.1021/acscatal.9b02416. 3.2
- [11] Seoin Back, Junwoong Yoon, Nianhan Tian, Wen Zhong, Kevin Tran, and Zachary W. Ulissi. Convolutional neural network of atomic surface structures to predict binding energies for high-throughput screening of catalysts. *The Journal of Physical Chemistry Letters*, 10(15):4401–4408, 2019. doi: 10.1021/ acs.jpclett.9b01428. 4.2

- [12] Seoin Back, Junwoong Yoon, Nianhan Tian, Wen Zhong, Kevin Tran, and Zachary W. Ulissi. Convolutional neural network of atomic surface structures to predict binding energies for high-throughput screening of catalysts. *The Journal of Physical Chemistry Letters*, 10(15):4401–4408, 2019. doi: 10.1021/ acs.jpclett.9b01428. URL https://doi.org/10.1021/acs.jpclett.9b01428. PMID: 31310543. 5.2
- [13] Albert P. Bartók, Mike C. Payne, Risi Kondor, and Gábor Csányi. Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.*, 104:136403, Apr 2010. doi: 10.1103/PhysRevLett. 104.136403. URL https://link.aps.org/doi/10.1103/PhysRevLett.104. 136403. 4.3.1
- [14] Albert P. Bartók, Risi Kondor, and Gábor Csányi. On representing chemical environments. *Phys. Rev. B*, 87:184115, May 2013. doi: 10.1103/PhysRevB.87. 184115. URL https://link.aps.org/doi/10.1103/PhysRevB.87.184115.
 1.3
- [15] Albert P. Bartók, Risi Kondor, and Gábor Csányi. On representing chemical environments. *Physical Review B*, 87(18), May 2013. ISSN 1550-235X. doi: 10.1103/physrevb.87.184115. URL http://dx.doi.org/10.1103/PhysRevB. 87.184115. 4.3.1
- [16] Thomas AA Batchelor, Jack K Pedersen, Simon H Winther, Ivano E Castelli, Karsten W Jacobsen, and Jan Rossmeisl. High-entropy alloys as a discovery platform for electrocatalysis. *Joule*, 3(3):834–845, 2019. 3.4, 3-2, 3.6
- [17] Jörg Behler and Michele Parrinello. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.*, 98:146401, Apr 2007. doi: 10.1103/PhysRevLett.98.146401. 1.3, 4.3.1, 5.3.3
- [18] Jörg Behler and Michele Parrinello. Generalized neural-network representation

of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.*, 98:146401, Apr 2007. doi: 10.1103/PhysRevLett.98.146401. URL https://link.aps.org/doi/10.1103/PhysRevLett.98.146401. 1.3

- [19] Jörg Behler. Constructing high-dimensional neural network potentials: A tutorial review. International Journal of Quantum Chemistry, 115(16):1032–1050, 2015. doi: 10.1002/qua.24890. 4.2, 5.2
- [20] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives, 2012. URL https://arxiv.org/abs/1206.
 5538. 1.3
- [21] Arno Bergmann and Beatriz Roldan Cuenya. Operando Insights into Nanoparticle Transformations during Catalysis. ACS Catalysis, 9(11):10020–10043, 2019. ISSN 21555435. doi: 10.1021/acscatal.9b01831. 5.2
- [22] E. Bernhardsson, E. Freider, and A. Rouhani. Luigi, a python package that builds complex pipelines of batch jobs, 2012. URL https://github.com/ spotify/luigi. (bithub, 2012). 3.7
- [23] J.C Bertolini, P Miegge, P Hermann, J.L Rousset, and B Tardy. On the reactivity of 2d pd surface alloys obtained by surface segregation or deposition technique. Surface Science, 331-333:651-658, 1995. ISSN 0039-6028. doi: https://doi.org/10.1016/0039-6028(95)00144-1. URL https:// www.sciencedirect.com/science/article/pii/0039602895001441. 5.2
- [24] Robert B. Best, Xiao Zhu, Jihyun Shim, Pedro E.M. Lopes, Jeetain Mittal, Michael Feig, and Alexander D. MacKerell. Optimization of the additive CHARMM all-atom protein force field targeting improved sampling of the backbone φ, ψ and side-chain χ1 and χ2 Dihedral Angles. Journal of Chemical Theory and Computation, 8(9):3257–3273, 2012. ISSN 15499626. doi: 10.1021/ct300400x. 2.3.2

- [25] Uttam Bhandari, Md Rumman Rafi, Congyan Zhang, and Shizhong Yang. Yield strength prediction of high-entropy alloys using machine learning. *Materials Today Communications*, 26:101871, 2021. 5.2
- [26] T Bligaard, Jens Kehlet Nørskov, Søren Dahl, J Matthiesen, Claus H Christensen, and J Sehested. The brønsted–evans–polanyi relation and the volcano curve in heterogeneous catalysis. J. Catal., 224(1):206–217, 2004. 3.2
- [27] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016. 5.3.6
- [28] S F Burlatsky, V V Atrazhev, D V Dmitriev, V I Sultanov, E N Timokhina, E A Ugolkova, S Tulyani, and A Vincitore. Surface tension model for surfactant solutions at the critical micelle concentration. Journal of Colloid and Interface Science, 393:151-160, 2013. ISSN 1095-7103. doi: 10.1016/j.jcis.2012.10.020. URL https://www.scopus.com/ inward/record.url?eid=2-s2.0-84873058501{&}partnerID=40{&}md5= 90619f6d2c8b8bd4964c129f37f3135f. 2.2
- [29] J. J. Burton, C. R. Helms, and R. S. Polizzotti. Surface segregation in alloys: LEED, Auger, and gas adsorption study of segregation of Au to the (111) surface of Ni. *The Journal of Chemical Physics*, 65(3):1089–1100, 1976. doi: 10.1063/1.433171. URL https://doi.org/10.1063/1.433171. 5.4.1
- [30] Lucian Busoniu, Robert Babuska, and Bart De Schutter. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2):156–172, 2008. doi: 10.1109/TSMCC.2007.913919.
- [31] Keith T. Butler, Daniel W. Davies, Hugh Cartwright, Olexandr Isayev, and Aron Walsh. Machine learning for molecular and materials science. *Nature*, 559 (7715):547–555, Jul 2018. ISSN 1476-4687. doi: 10.1038/s41586-018-0337-2.

URL https://doi.org/10.1038/s41586-018-0337-2. 1.2

- [32] Federico Calle-Vallejo, José I Martínez, Juan M García-Lastra, Philippe Sautet, and David Loffreda. Fast prediction of adsorption properties for platinum nanocatalysts with generalized coordination numbers. *Angew. Chem. Int. Ed.*, 53(32):8316–8319, 2014. 3.4, 3-2, 3.6
- [33] Zhuo Cao, Yabo Dan, Zheng Xiong, Chengcheng Niu, Xiang Li, Songrong Qian, and Jianjun Hu. Convolutional neural networks for crystal material property prediction using hybrid orbital-field matrix and magpie descriptors. Crystals, 9(4), 2019. ISSN 2073-4352. URL https://www.mdpi.com/2073-4352/9/4/191. 4-1
- [34] Karen Chan, Charlie Tsai, Heine A Hansen, and Jens K Nørskov. Molybdenum sulfides and selenides as possible electrocatalysts for co₂ reduction. *Chem-CatChem*, 6(7):1899–1905, 2014. 3.2
- [35] Benson Chen, Regina Barzilay, and Tommi Jaakkola. Path-augmented graph transformer network, 2019. 4.2
- [36] Chi Chen, Weike Ye, Yunxing Zuo, Chen Zheng, and Shyue Ping Ong. Graph networks as a universal machine learning framework for molecules and crystals. *Chem. Mater.*, 31(9):3564–3572, 2019. 3.2
- [37] Tao Cheng, Qing Chen, Feng Li, and Huai Sun. Classic force field for predicting surface tension and interfacial properties of sodium dodecyl sulfate. *The Journal of Physical Chemistry B*, 114(43):13736–13744, 2010. doi: 10.1021/jp107002x. URL https://doi.org/10.1021/jp107002x. 2.2
- [38] Stefan Chmiela, Alexandre Tkatchenko, Huziel E. Sauceda, Igor Poltavsky, Kristof T. Schütt, and Klaus-Robert Müller. Machine learning of accurate energy-conserving molecular force fields. *Science Advances*, 3(5), 2017. doi: 10. 1126/sciadv.1603015. URL https://advances.sciencemag.org/content/3/
5/e1603015. 4.3.1

- [39] Mads-Peter V. Christiansen, Henrik Lund Mortensen, Søren Ager Meldgaard, and Bjørk Hammer. Gaussian representation for image recognition and reinforcement learning of atomistic structure. *The Journal of Chemical Physics*, 153(4):044107, 2020. doi: 10.1063/5.0015571. URL https://doi.org/10. 1063/5.0015571. 5.2
- [40] Scott Clark and Patrick Hayes. SigOpt Web page. https://sigopt.com, 2019.
 URL https://sigopt.com. (document), 4.3.3, 4-3
- [41] J. Cortes, S. Martinez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243– 255, 2004. doi: 10.1109/TRA.2004.824698. 6
- [42] Jérôme Creuze, Hazar Guesmi, Christine Mottet, Beien Zhu, and Bernard Legrand. Surface segregation in AuPd alloys: Ab initio analysis of the driving forces. Surface Science, 639:48-53, 2015. ISSN 0039-6028. doi: https:// doi.org/10.1016/j.susc.2015.04.017. URL https://www.sciencedirect.com/ science/article/pii/S0039602815001089. 5.4.1
- [43] Eric Darve, David Rodríguez-Gómez, and Andrew Pohorille. Adaptive biasing force method for scalar and vector free energy calculations. *Journal of Chemical Physics*, 128(14):1–13, 2008. ISSN 00219606. doi: 10.1063/1.2829861. 2.3.2
- [44] Li Deng and John Platt. Ensemble deep learning for speech recognition. In Proc. Interspeech, September 2014. URL https://www.microsoft.com/en-us/research/publication/ ensemble-deep-learning-for-speech-recognition/. 1.3
- [45] Volker L. Deringer, Chris J. Pickard, and Gábor Csányi. Data-driven learning of total and local energies in elemental boron. *Phys. Rev. Lett.*, 120:156001, Apr 2018. doi: 10.1103/PhysRevLett.120.156001. URL https://link.aps.

org/doi/10.1103/PhysRevLett.120.156001. 4.2, 5.2

- [46] Volker L. Deringer, Miguel A. Caro, and Gábor Csányi. Machine learning interatomic potentials as emerging tools for materials science. Advanced Materials, 31(46):1902765, 2019. doi: 10.1002/adma.201902765. 4.2, 5.2
- [47] G. N. Derry, R. Wan, F. Strauch, and C. English. Segregation and interlayer relaxation at the NiPd(111) surface. Journal of Vacuum Science & Technology A, 29(1):011015, 2011. doi: 10.1116/1.3520434. URL https://doi.org/10.1116/1.3520434. 5.4.1
- [48] Colin F Dickens, Joseph H Montoya, Ambarish R Kulkarni, Michal Bajdich, and Jens K Nørskov. An electronic structure descriptor for oxygen reactivity at metal and metal-oxide surfaces. *Surf. Sci.*, 681:122–129, 2019. 3.4, 3-2
- [49] Yuxin Ding, Shengli Yan, YiBin Zhang, Wei Dai, and Li Dong. Predicting the attributes of social network users using a graph-based machine learning method. *Computer Communications*, 73:3–11, 2016. ISSN 0140-3664. doi: https://doi.org/10.1016/j.comcom.2015.07.007. URL https://www. sciencedirect.com/science/article/pii/S0140366415002455. 1.3
- [50] Abhijit Dutta and Jayati Datta. Outstanding catalyst performance of pdauni nanoparticles for the anodic reaction in an alkaline direct ethanol (with anionexchange membrane) fuel cell. *The Journal of Physical Chemistry C*, 116 (49):25677–25688, 2012. doi: 10.1021/jp305323s. URL https://doi.org/ 10.1021/jp305323s. 5.2
- [51] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alan Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, Advances in Neural Information Processing Systems 28, pages 2224–2232.

Curran Associates, Inc., 2015. URL http://papers.nips.cc/paper/ 5954-convolutional-networks-on-graphs-for-learning-molecular-fingerprints. pdf. 4.2

- [52] Theodosios Famprikis, Pieremanuele Canepa, James A. Dawson, M. Saiful Islam, and Christian Masquelier. Fundamentals of inorganic solid-state electrolytes for batteries. *Nature Materials*, 18(12):1278–1291, Dec 2019. ISSN 1476-4660. doi: 10.1038/s41563-019-0431-3. URL https://doi.org/10. 1038/s41563-019-0431-3. 1.1
- [53] Zhou Fan, Rui Su, Weinan Zhang, and Yong Yu. Hybrid actor-critic reinforcement learning in parameterized action space. arXiv preprint arXiv:1903.01344, 2019. 5.3.2
- [54] Scott Feller and Alexander D. MacKerell. An improved empirical potential energy function for molecular simulations of phospholipids. *The Journal of Physical Chemistry B*, 104(31):7510–7515, 2000. doi: 10.1021/jp0007843. URL https://doi.org/10.1021/jp0007843. 2.3.3
- [55] Scott E. Feller, Yuhong Zhang, Richard W. Pastor, and Bernard R. Brooks. Constant pressure molecular dynamics simulation: The Langevin piston method. *The Journal of Chemical Physics*, 103(11):4613–4621, 1995. ISSN 00219606. doi: 10.1063/1.470648. 2.3.2, 2.3.3
- [56] Scott E Feller, Klaus Gawrisch, and Alexander D MacKerell. Polyunsaturated Fatty Acids in Lipid Bilayers: Intrinsic and Environmental Contributions to Their Unique Physical Properties. Journal of the American Chemical Society, 124(2):318–326, 2002. doi: 10.1021/ja0118340. URL https://doi.org/10. 1021/ja0118340. 2.3.3
- [57] Filipe Vargas Ferreira, Luciana De Simone Cividanes, Felipe Sales Brito, Beatriz Rossi Canuto de Menezes, Wesley Franceschi, Evelyn Alves Nunes Si-

monetti, and Gilmar Patrocínio Thim. Functionalization of Graphene and Applications, pages 1–29. Springer International Publishing, Cham, 2016. ISBN 978-3-319-35110-0. doi: 10.1007/978-3-319-35110-0_1. URL https://doi.org/10.1007/978-3-319-35110-0_1. 1.1

- [58] Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In ICLR Workshop on Representation Learning on Graphs and Manifolds, 2019. 4.3.2
- [59] Giacomo Fiorin, Michael L. Klein, and Jérôme Hénin. Using collective variables to drive molecular dynamics simulations. *Molecular Physics*, 111(22-23):3345– 3362, 2013. ISSN 00268976. doi: 10.1080/00268976.2013.813594. 1.2
- [60] Jie Fu, Xiaofang Yang, Carl A. Menning, Jingguang G. Chen, and Bruce E. Koel. Composition, structure and stability of surfaces formed by ni deposition on pd(111). Surface Science, 646:56–64, 2016. ISSN 0039-6028. doi: https://doi.org/10.1016/j.susc.2015.05.026. URL https://www.sciencedirect.com/science/article/pii/S0039602815001582. Surface science for heterogeneous catalysis, a special issue in Honour of Richard Lambert. 5.2
- [61] Terry ZH Gani and Heather J Kulik. Understanding and breaking scaling relations in single-site catalysis: Methane to methanol conversion by fe^{IV}=0. ACS Catal., 8(2):975–986, 2018. 3.2
- [62] Wenpei Gao, Zachary D. Hood, and Miaofang Chi. Interfaces in heterogeneous catalysts: Advancing mechanistic understanding through atomicscale measurements. Accounts of Chemical Research, 50(4):787-795, 2017. doi: 10.1021/acs.accounts.6b00596. URL https://doi.org/10.1021/acs. accounts.6b00596. PMID: 28207240. 1.1
- [63] M. Gastegger, L. Schwiedrzik, M. Bittermann, F. Berzsenyi, and P. Marquetand. wacsf—weighted atom-centered symmetry functions as descriptors

in machine learning potentials. The Journal of Chemical Physics, 148(24): 241709, Jun 2018. ISSN 1089-7690. doi: 10.1063/1.5019667. URL http: //dx.doi.org/10.1063/1.5019667. 4.3.1

- [64] Johannes Gasteiger, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs, 2020. URL https://arxiv.org/abs/2003.
 03123. 6
- [65] Johannes Gasteiger, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional graph neural networks for molecules, 2021. URL https://arxiv.org/abs/2106.08903. 6
- [66] A. K. Geim and K. S. Novoselov. The rise of graphene. Nature Materials, 6(3): 183-191, Mar 2007. ISSN 1476-4660. doi: 10.1038/nmat1849. URL https: //doi.org/10.1038/nmat1849. 1.1
- [67] Luca M. Ghiringhelli, Jan Vybiral, Sergey V. Levchenko, Claudia Draxl, and Matthias Scheffler. Big data of materials science: Critical role of the descriptor. *Phys. Rev. Lett.*, 114:105503, Mar 2015. doi: 10.1103/PhysRevLett.114.105503.
 URL https://link.aps.org/doi/10.1103/PhysRevLett.114.105503.
- [68] J. Willard Gibbs and E. B. Wilson. The scientific papers of j. willard gibbs: Vol.
 i. thermodynamics, elementary principles in statistical mechanics and vector analysis. *American Journal of Physics*, 30(4):313–314, 1962. doi: 10.1119/1. 1942006. URL https://doi.org/10.1119/1.1942006. 1.2
- [69] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17, page 1263–1272. JMLR.org, 2017. 4.3.2
- [70] Stefan Goedecker. Minima hopping: An efficient search method for the global minimum of the potential energy surface of complex molecular systems. journal

of chemical physics, 120(21), 9911-9917. h. Journal of Chemical Physics, 120 (21):9911-9917, 2004. doi: 10.1063/1.1724816. URL ttps://doi.org/10.1063/1.1724816. 5.3.7

- [71] Garrett B. Goh, Nathan O. Hodas, and Abhinav Vishnu. Deep learning for computational chemistry. Journal of Computational Chemistry, 38 (16):1291-1307, 2017. doi: https://doi.org/10.1002/jcc.24764. URL https: //onlinelibrary.wiley.com/doi/abs/10.1002/jcc.24764. 1.3
- [72] J Greeley, IEL Stephens, AS Bondarenko, Tobias Peter Johansson, Heine Anton Hansen, TF Jaramillo, Jan Rossmeisl, INJK Chorkendorff, and Jens Kehlet Nørskov. Alloys of platinum and early transition metals as oxygen reduction electrocatalysts. *Nat. Chem.*, 1(7):552, 2009. 3.2
- [73] Jeff Greeley and Jens K Nørskov. Large-scale, density functional theory-based screening of alloys for hydrogen evolution. Surf. Sci., 601(6):1590–1598, 2007.
 3.2
- [74] Jeff Greeley and Jens K Nørskov. Combinatorial density functional theorybased screening of surface alloys for the oxygen reduction reaction. J. Phys. Chem. C, 113(12):4932–4939, 2009. 3.2
- [75] Isabel Xiaoye Green, Wenjie Tang, Matthew Neurock, and John T. Yates. Spectroscopic observation of dual catalytic sites during oxidation of CO on a Au/TiO 2 catalyst. *Science*, 333(6043):736–739, 2011. ISSN 00368075. doi: 10.1126/science.1207272. 5.2
- [76] A Grossfield. WHAM: The Weighted Histogram Analysis Method, v. 2.0.9;. URL http://membrane.urmc.rochester.edu/content/wham. http: //membrane.urmc.rochester.edu/content/wham. 2.3.3
- [77] Geun Ho Gu, Juhwan Noh, Sungwon Kim, Seoin Back, Zachary Ulissi, and Yousung Jung. Practical deep-learning representation for fast heterogeneous

catalyst screening. The Journal of Physical Chemistry Letters, 11(9):3185-3191, 2020. doi: 10.1021/acs.jpclett.0c00634. URL https://doi.org/10. 1021/acs.jpclett.0c00634. PMID: 32191473. 6

- [78] Geun Ho Gu, Juhwan Noh, Sungwon Kim, Seoin Back, Zachary Ulissi, and Yousung Jung. Practical deep-learning representation for fast heterogeneous catalyst screening. *The Journal of Physical Chemistry Letters*, 11(9):3185– 3191, 2020. doi: 10.1021/acs.jpclett.0c00634. URL https://doi.org/10. 1021/acs.jpclett.0c00634. 4.2
- [79] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. arXiv preprint arXiv:1812.05905, 2018. 5.3.2
- [80] B. Hammer, L. B. Hansen, and J. K. Nørskov. Improved adsorption energetics within density-functional theory using revised perdew-burke-ernzerhof functionals. *Phys. Rev. B*, 59:7413–7421, Mar 1999. doi: 10.1103/PhysRevB. 59.7413. URL https://link.aps.org/doi/10.1103/PhysRevB.59.7413. 4.2
- [81] BHLB Hammer, Lars Bruno Hansen, and Jens Kehlet Nørskov. Improved adsorption energetics within density functional theory using revised perdewburke-ernzerhof functionals. *Phys. Rev. B*, 59(11):7413, 1999. 3.2, 3.7
- [82] HA Hansen, C Shi, AC Lausche, AA Peterson, and JK Nørskov. Bifunctional alloys for the electroreduction of co₂ and co. *Phys. Chem. Chem. Phys.*, 18 (13):9194–9201, 2016. 3.2
- [83] Hendrik Heinz, Chandrani Pramanik, Ozge Heinz, Yifu Ding, Ratan K. Mishra, Delphine Marchon, Robert J. Flatt, Irina Estrela-Lopis, Jordi Llop, Sergio Moya, and Ronald F. Ziolo. Nanoparticle decoration with surfactants: Molecular interactions, assembly, and applications. Surface Science Reports, 72

(1):1-58, 2017. ISSN 01675729. doi: 10.1016/j.surfrep.2017.02.001. URL http://dx.doi.org/10.1016/j.surfrep.2017.02.001. 2.2

- [84] J Henin, G Fiorin, C Chipot, and M L Klein. Exploring Multidimensional Free Energy Landscapes Using Time-Dependent Biases on Collective Variables. J. Chem. Theory Comput., 6(1):35–47, 2010. doi: 10.1021/ct9004432. 2.3.2
- [85] Graeme Henkelman and Hannes Jónsson. Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points. The Journal of Chemical Physics, 113(22):9978–9985, 2000. doi: 10. 1063/1.1323224. URL https://doi.org/10.1063/1.1323224. 5.4.5
- [86] Graeme Henkelman, Blas P. Uberuaga, and Hannes Jónsson. A climbing image nudged elastic band method for finding saddle points and minimum energy paths. *The Journal of Chemical Physics*, 113(22):9901–9904, 2000. doi: 10. 1063/1.1329672. URL https://doi.org/10.1063/1.1329672. 5.4.5
- [87] Alyssa J. R. Hensley, Kushal Ghale, Carolin Rieg, Thanh Dang, Emily Anderst, Felix Studt, Charles T. Campbell, Jean-Sabin McEwen, and Ye Xu. Dft-based method for more accurate adsorption energies: An adaptive sum of energies from rpbe and vdw density functionals. *The Journal of Physical Chemistry C*, 121(9):4937–4945, 2017. doi: 10.1021/acs.jpcc.6b10187. URL https://doi. org/10.1021/acs.jpcc.6b10187. 4.2
- [88] Alyssa J. R. Hensley, Kushal Ghale, Carolin Rieg, Thanh Dang, Emily Anderst, Felix Studt, Charles T. Campbell, Jean-Sabin McEwen, and Ye Xu. Dft-based method for more accurate adsorption energies: An adaptive sum of energies from rpbe and vdw density functionals. *The Journal of Physical Chemistry C*, 121(9):4937–4945, 2017. doi: 10.1021/acs.jpcc.6b10187. URL https://doi. org/10.1021/acs.jpcc.6b10187. 1.2
- [89] Alyssa JR Hensley, Kushal Ghale, Carolin Rieg, Thanh Dang, Emily Anderst,

Felix Studt, Charles T Campbell, Jean-Sabin McEwen, and Ye Xu. Dft-based method for more accurate adsorption energies: an adaptive sum of energies from rpbe and vdw density functionals. *J. Phys. Chem. C*, 121(9):4937–4945, 2017. 1.2, 3.2

- [90] Eric D. Hermes, Khachik Sargsyan, Habib N. Najm, and Judit Zádor. Accelerated saddle point refinement through full exploitation of partial hessian diagonalization. Journal of Chemical Theory and Computation, 15(11):6536– 6549, 2019. doi: 10.1021/acs.jctc.9b00869. URL https://doi.org/10.1021/ acs.jctc.9b00869. PMID: 31614079. 5.3.4
- [91] Tirma Herranz, Xingyi Deng, Andreu Cabot, Zhi Liu, and Miquel Salmeron. In situ XPS study of the adsorption and reactions of NO and O2 on gold nanoparticles deposited on TiO2 and SiO2. *Journal of Catalysis*, 283(2):119– 123, 2011. ISSN 00219517. doi: 10.1016/j.jcat.2011.06.022. 5.2
- [92] Bruno Honrado Guerreiro, Manuel. H. Martin, Lionel Roué, and Daniel Guay. Hydrogen solubility of magnetron co-sputtered fcc and bcc pdcuau thin films. *The Journal of Physical Chemistry C*, 120(10):5297–5307, 2016. doi: 10.1021/ acs.jpcc.5b10711. URL https://doi.org/10.1021/acs.jpcc.5b10711. 5.2
- [93] Andrew G. Hsieh, Christian Punckt, Sibel Korkut, and Ilhan A. Aksay. Adsorption of sodium dodecyl sulfate on functionalized graphene measured by conductometric titration. *Journal of Physical Chemistry B*, 117(26):7950–7958, 2013. ISSN 15205207. doi: 10.1021/jp403876t. 2.2, 2.2, 2.4.2
- [94] Mingwei Hu, Qiyang Tan, Ruth Knibbe, Sen Wang, Xue Li, Tianqi Wu, Sams Jarin, and Ming-Xing Zhang. Prediction of mechanical properties of wrought aluminium alloys using feature engineering assisted machine learning approach. *Metallurgical and Materials Transactions A*, pages 1–12, 2021. 5.2
- [95] Kevin J. Huang, Gerbrand Ceder, and Elsa A. Olivetti. Manufacturing

scalability implications of materials choice in inorganic solid-state batteries. *Joule*, 5(3):564-580, 2021. ISSN 2542-4351. doi: https://doi.org/10. 1016/j.joule.2020.12.001. URL https://www.sciencedirect.com/science/ article/pii/S2542435120305699. 1.1

- [96] Weixin Huang and Wei-Xue Li. Surface and interface design for heterogeneous catalysis. *Phys. Chem. Chem. Phys.*, 21:523-536, 2019. doi: 10.1039/C8CP05717F. URL http://dx.doi.org/10.1039/C8CP05717F. 1.1
- [97] William Humphrey, Andrew Dalke, and Klaus Schulten. VMD Visual Molecular Dynamics. Journal of Molecular Graphics, 14:33–38, 1996. 2.4.2
- [98] K. W. Jacobsen, J. K. Norskov, and M. J. Puska. Interatomic interactions in the effective-medium theory. *Phys. Rev. B*, 35:7423-7442, May 1987. doi: 10.1103/PhysRevB.35.7423. URL https://link.aps.org/doi/ 10.1103/PhysRevB.35.7423. 5.3.5
- [99] K.W. Jacobsen, P. Stoltze, and J.K. Nørskov. A semi-empirical effective medium theory for metals and alloys. Surface Science, 366(2):394 402, 1996. ISSN 0039-6028. doi: https://doi.org/10.1016/0039-6028(96) 00816-3. URL http://www.sciencedirect.com/science/article/pii/0039602896008163. 5.3.5
- [100] Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, and Gerbrand Ceder. Commentary: The materials project: A materials genome approach to accelerating materials innovation. APL Mater., 1(1):011002, 2013. 3.7
- [101] Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, and Kristin A. Persson. Commentary: The materials project:

A materials genome approach to accelerating materials innovation. *APL Materials*, 1(1):011002, 2013. doi: 10.1063/1.4812323. 4.4

- [102] Anubhav Jain, Shyue Ping Ong, Wei Chen, Bharat Medasani, Xiaohui Qu, Michael Kocher, Miriam Brafman, Guido Petretto, Gian-Marco Rignanese, Geoffroy Hautier, Daniel Gunter, and Kristin A. Persson. Fireworks: a dynamic workflow system designed for high-throughput applications. *Concurrency Computat.: Pract. Exper.*, 27(17):5037–5059, 2015. ISSN 1532-0634. doi: 10.1002/cpe.3505. URL http://dx.doi.org/10.1002/cpe.3505. 3.7
- [103] Anubhav Jain, Yongwoo Shin, and Kristin A Persson. Computational predictions of energy materials using density functional theory. *Nature Reviews Materials*, 1(1):15004, 2016. ISSN 2058-8437. doi: 10.1038/natrevmats.2015.4.
 URL https://doi.org/10.1038/natrevmats.2015.4. 4.2
- [104] Aditya Jaishankar, Arben Jusufi, Jessica L. Vreeland, Shane Deighton, Joseph Pellettiere, and Alan M. Schilowitz. Adsorption of stearic acid at the iron oxide/oil interface: Theory, experiments, and modeling. *Langmuir*, 35(6):2033–2046, 2019. doi: 10.1021/acs.langmuir.8b03132. URL https://doi.org/10.1021/acs.langmuir.8b03132. 2.2
- [105] Glenn Jones, T Bligaard, Frank Abild-Pedersen, and Jens Kehlet Nørskov. Using scaling relations to understand trends in the catalytic activity of transition metals. J. Phys.: Condens. Matter, 20(6):064239, 2008. 3.2
- [106] Mathias S. Jørgensen, Henrik L. Mortensen, Søren A. Meldgaard, Esben L. Kolsbjerg, Thomas L. Jacobsen, Knud H. Sørensen, and Bjørk Hammer. Atomistic structure learning. *The Journal of Chemical Physics*, 151(5):054111, 2019. doi: 10.1063/1.5108871. URL https://doi.org/10.1063/1.5108871. 5.2
- [107] B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A. Al Sallab, Senthil Yogamani, and Patrick Pérez. Deep reinforcement learning for

autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–18, 2021. doi: 10.1109/TITS.2021.3054625. 1.3

- [108] Johannes Klicpera, Janek Gro
 ß, and Stephan G
 ünnemann. Directional message passing for molecular graphs, 2020. 4.2
- [109] Kensaku Kodama, Tomoyuki Nagai, Akira Kuwaki, Ryosuke Jinnouchi, and Yu Morimoto. Challenges in applying highly active pt-based nanostructured catalysts for oxygen reduction reactions to fuel cell vehicles. Nature Nanotechnology, 16(2):140–147, Feb 2021. ISSN 1748-3395. doi: 10.1038/s41565-020-00824-w. URL https://doi.org/10.1038/ s41565-020-00824-w. 1.1
- [110] M. Krawczyk, L. Zommer, J.W. Sobczak, A. Jablonski, M. Petit, C. Robert-Goumet, and B. Gruzza. IMFP measurements near Au-Ni alloy surfaces by EPES: indirect evidence of submonolayer Au surface enrichment. Surface Science, 566-568:856-861, 2004. ISSN 0039-6028. doi: https://doi.org/10.1016/j.susc.2004.06.021. URL https://www.sciencedirect.com/science/article/pii/S003960280400665X. Proceedings of the 22nd European Conference on Surface Science. 5.4.1
- [111] Georg Kresse and Jürgen Furthmüller. Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set. Comput. Mater. Sci., 6(1):15–50, 1996. 3.7
- [112] Georg Kresse and J Hafner. Ab initio molecular dynamics for open-shell transition metals. *Phys. Rev. B*, 48(17):13115, 1993. 3.7
- [113] Georg Kresse and D Joubert. From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B*, 59(3):1758, 1999. 3.7
- [114] Dilip Krishnamurthy, Vaidish Sumaria, and Venkatasubramanian Viswanathan. Quantifying robustness of dft predicted pathways and ac-

tivity determining elementary steps for electrochemical reactions. J. Chem. Phys., 150(4):041717, 2019. 1.2

- [115] Hubert Kuhn and Heinz Rehage. Molecular dynamics computer simulations of surfactant monolayers: monododecyl pentaethylene glycol at the surface between air and water. *The Journal of Physical Chemistry B*, 103(40):8493–8501, 1999. doi: 10.1021/jp9905950. URL https://doi.org/10.1021/jp9905950.
 2.2
- [116] Alexander Kuhnle, Michael Schaarschmidt, and Kai Fricke. Tensorforce: a tensorflow library for applied reinforcement learning. Web page, 2017. URL https://github.com/tensorforce/tensorforce. 5.3.6
- [117] Jayesh A. Kulkarni, Maria M. Darjuan, Joanne E. Mercer, Sam Chen, Roy Van Der Meel, Jenifer L. Thewalt, Yuen Yi C. Tam, and Pieter R. Cullis. On the Formation and Morphology of Lipid Nanoparticles Containing Ionizable Cationic Lipids and siRNA. ACS Nano, 12(5):4787–4795, 2018. ISSN 1936086X. doi: 10.1021/acsnano.8b01516. 2.2
- [118] Shankar Kumar, John M. Rosenberg, Djamal Bouzida, Robert H. Swendsen, and Peter A. Kollman. The weighted histogram analysis method for free-energy calculations on biomolecules. I. The method. *Journal of Computational Chemistry*, 13(8):1011–1021, 1992. ISSN 1096987X. doi: 10.1002/jcc.540130812. 2.3.3
- [119] Shankar Kumar, John M. Rosenberg, Djamal Bouzida, Robert H. Swendsen, and Peter A. Kollman. Multidimensional free-energy calculations using the weighted histogram analysis method. *Journal of Computational Chemistry*, 16 (11):1339–1350, 1995. ISSN 1096987X. doi: 10.1002/jcc.540161104. 2.3.3
- [120] Hye-Eun Lee, Ki Dong Yang, Sang Moon Yoon, Hyo-Yong Ahn, Yoon Young Lee, Hyejin Chang, Dae Hong Jeong, Yoon-Sik Lee, Mi Young Kim, and Ki Tae

Nam. Concave rhombic dodecahedral au nanocatalyst with multiple high-index facets for co_2 reduction. ACS Nano, 9(8):8384–8393, 2015. 3.2

- [121] Jae Won Lee, Jonghun Park, Jangmin O, Jongwoo Lee, and Euyseok Hong. A multiagent approach to q-learning for daily stock trading. *IEEE Transactions* on Systems, Man, and Cybernetics - Part A: Systems and Humans, 37(6): 864–877, 2007. doi: 10.1109/TSMCA.2007.904825. 6
- [122] Z. X. Li, C. C. Dong, and R. K. Thomas. Neutron reflectivity studies of the surface excess of gemini surfactants at the air-water interface. *Langmuir*, 15 (13):4392–4396, 1999. doi: 10.1021/la981551u. URL https://doi.org/10.1021/la981551u. 2.2
- [123] Zheng Li, Xianfeng Ma, and Hongliang Xin. Feature engineering of machinelearning chemisorption models for catalyst design. *Catal. Today*, 280:232–238, 2017. 3.2
- [124] Zheng Li, Siwen Wang, Wei Shan Chin, Luke E. Achenie, and Hongliang Xin. High-throughput screening of bimetallic catalysts enabled by machine learning. J. Mater. Chem. A, 5:24131-24138, 2017. doi: 10.1039/C7TA01812F. URL http://dx.doi.org/10.1039/C7TA01812F. 5.2
- [125] Zheng Li, Siwen Wang, Wei Shan Chin, Luke E Achenie, and Hongliang Xin. High-throughput screening of bimetallic catalysts enabled by machine learning. J. Mater. Chem. A, 5(46):24131–24138, 2017. 3.2, 3-2, 3.4, 3.6
- [126] T. Lillicrap, J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *CoRR*, abs/1509.02971, 2016. 5.3.1
- [127] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning, 2015. URL https://arxiv.org/abs/

1509.02971. 1.3

- [128] Hyung-Kyu Lim, Hyeyoung Shin, William A Goddard III, Yun Jeong Hwang, Byoung Koun Min, and Hyungjun Kim. Embedding covalency into metal catalysts for efficient electrochemical conversion of co₂. J. Am. Chem. Soc., 136(32):11355–11361, 2014. 3.2
- [129] Jianwen Liu, Wenzhi Luo, Lei Wang, Jiujun Zhang, Xian-Zhu Fu, and Jing-Li Luo. Toward excellence of electrocatalyst design by emerging descriptor-oriented machine learning. Advanced Functional Materials, n/a (n/a):2110748. doi: https://doi.org/10.1002/adfm.202110748. URL https: //onlinelibrary.wiley.com/doi/abs/10.1002/adfm.202110748. 1.2
- [130] Suli Liu, Hui Zhang, Xueqin Mu, and Changyun Chen. Surface reconstruction engineering of twinned pd2coag nanocrystals by atomic vacancy inducement for hydrogen evolution and oxygen reduction reactions. Applied Catalysis B: Environmental, 241:424-429, 2019. ISSN 0926-3373. doi: https: //doi.org/10.1016/j.apcatb.2018.09.067. URL https://www.sciencedirect. com/science/article/pii/S0926337318308968. 5.2
- [131] Yue Liu, Tianlu Zhao, Wangwei Ju, and Siqi Shi. Materials discovery and design using machine learning. *Journal of Materiomics*, 3(3):159–177, 2017. ISSN 2352-8478. doi: https://doi.org/10.1016/j.jmat.2017.08.002. URL https://www.sciencedirect.com/science/article/pii/S2352847817300515. High-throughput Experimental and Modeling Research toward Advanced Batteries. 1.2
- [132] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. arXiv, 2017. https://arxiv.org/abs/1711.05101. 3.4, 4.3.3
- [133] Ziheng Lu. Computational discovery of energy materials in the era of big data and machine learning: A critical review. *Materials Reports: Energy*,

1(3):100047, 2021. ISSN 2666-9358. doi: https://doi.org/10.1016/j.matre. 2021.100047. URL https://www.sciencedirect.com/science/article/ pii/S2666935821000823. Machine Learning and Artificial Intelligence for Energy Materials. 1.2

- [134] Lai-Ming Luo, Wei Zhan, Rong-Hua Zhang, Di Chen, Qing-Yun Hu, Yi-Fei Guo, and Xin-Wen Zhou. Ternary coaupd and binary aupd electrocatalysts for methanol oxidation and oxygen reduction reaction: Enhanced catalytic performance by surface reconstruction. Journal of Power Sources, 412:142–152, 2019. ISSN 0378-7753. doi: https://doi.org/10.1016/ j.jpowsour.2018.11.044. URL https://www.sciencedirect.com/science/ article/pii/S0378775318312850. 5.2
- [135] Qun Luo, Yanlin Guo, Bin Liu, Yujun Feng, Jieyu Zhang, Qian Li, and Kuochih Chou. Thermodynamics and kinetics of phase transformation in rare earth-magnesium alloys: A critical review. Journal of Materials Science Technology, 44:171-190, 2020. ISSN 1005-0302. doi: https://doi.org/10. 1016/j.jmst.2020.01.022. URL https://www.sciencedirect.com/science/ article/pii/S1005030220300220. 1.2
- [136] Xianfeng Ma and Hongliang Xin. Orbitalwise coordination number for predicting adsorption properties of metal nanocatalysts. *Phys. Rev. Lett.*, 118(3): 036101, 2017. 3-2, 3.6
- [137] Xianfeng Ma, Zheng Li, Luke E. K. Achenie, and Hongliang Xin. Machinelearning-augmented chemisorption model for co2 electroreduction catalyst screening. *The Journal of Physical Chemistry Letters*, 6(18):3528-3533, 2015. doi: 10.1021/acs.jpclett.5b01660. URL https://doi.org/10.1021/ acs.jpclett.5b01660. PMID: 26722718. 5.2
- [138] A. D. MacKerell, D. Bashford, M. Bellott, R. L. Dunbrack, J. D. Evanseck,

M. J. Field, S. Fischer, J. Gao, H. Guo, S. Ha, D. Joseph-McCarthy, L. Kuchnir,
K. Kuczera, F. T. K. Lau, C. Mattos, S. Michnick, T. Ngo, D. T. Nguyen,
B. Prodhom, W. E. Reiher, B. Roux, M. Schlenkrich, J. C. Smith, R. Stote,
J. Straub, M. Watanabe, J. Wiórkiewicz-Kuczera, D. Yin, and M. Karplus.
All-atom empirical potential for molecular modeling and dynamics studies of
proteins. *The Journal of Physical Chemistry B*, 102(18):3586–3616, 1998. doi:
10.1021/jp973084f. URL https://doi.org/10.1021/jp973084f. 2.3.2, 2.3.3

- [139] Alexander D. MacKerell, Michael Feig, and Charles L. Brooks. Extending the treatment of backbone energetics in protein force fields: Limitations of gasphase quantum mechanics in reproducing protein conformational distributions in molecular dynamics simulations. *Journal of computational chemistry*, 25 11: 1400–15, 2004. 2.3.2, 2.3.3
- [140] Armando Maestro, Eva Santini, Dominika Zabiegaj, Sara Llamas, Francesca Ravera, Libero Liggieri, Francisco Ortega, Ramón G. Rubio, and Eduardo Guzman. Particle and particle-surfactant mixtures at fluid interfaces: Assembly, morphology, and rheological description. Advances in Condensed Matter Physics, 2015:917516, 2015. ISSN 16878124. doi: 10.1155/2015/917516. 2.2
- [141] Isabela C Man, Hai-Yan Su, Federico Calle-Vallejo, Heine A Hansen, José I Martínez, Nilay G Inoglu, John Kitchin, Thomas F Jaramillo, Jens K Nørskov, and Jan Rossmeisl. Universality in oxygen evolution electrocatalysis on oxide surfaces. *ChemCatChem*, 3(7):1159–1165, 2011. 3.2
- [142] Glenn J Martyna, Douglas J Tobias, and Michael L Klein. Constant pressure molecular dynamics algorithm. J. Chem. Phys., 101(September):4177–4189, 1994. 1.2, 2.3.3
- [143] Søren A Meldgaard, Henrik L Mortensen, Mathias S Jørgensen, and Bjørk Hammer. Structure prediction of surface reconstructions by deep reinforce-

ment learning. Journal of Physics: Condensed Matter, 32(40):404005, jul 2020. doi: 10.1088/1361-648x/ab94f2. URL https://doi.org/10.1088/1361-648x/ab94f2. 5.2

- [144] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602, 2013. 5.3.1
- [145] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. nature, 518(7540):529–533, 2015. 5.3.1
- [146] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference* on machine learning, pages 1928–1937. PMLR, 2016. 5.3.2
- [147] Ahmad Moghimikheirabadi, Leonard M. C. Sagis, Martin Kroger, and Patrick Ilg. Gas-liquid phase equilibrium of a model langmuir monolayer captured by a multiscale approach. *Phys. Chem. Chem. Phys.*, 21:2295–2306, 2019. doi: 10.1039/C8CP05447A. URL http://dx.doi.org/10.1039/C8CP05447A. 2.2
- [148] Matthew M Montemore and J Will Medlin. Scaling relations between adsorption energies for computational screening and design of catalysts. *Catal. Sci. Technol.*, 4(11):3748–3761, 2014. 3.2
- [149] Aqsa Naeem and Naveed Ul Hassan. Renewable energy intermittency mitigation in microgrids: State-of-the-art and future prospects. In 2020 4th International Conference on Green Energy and Applications (ICGEA), pages 158–164, 2020. doi: 10.1109/ICGEA49367.2020.239699. 1.1
- [150] Eiichi Nakamura. Atomic-Resolution Transmission Electron Microscopic

Movies for Study of Organic Molecules, Assemblies, and Reactions: The First 10 Years of Development. *Accounts of Chemical Research*, 50(6):1281–1292, 2017. ISSN 15204898. doi: 10.1021/acs.accounts.7b00076. 5.2

- [151] Michael Neunert, Abbas Abdolmaleki, Markus Wulfmeier, Thomas Lampe, Tobias Springenberg, Roland Hafner, Francesco Romano, Jonas Buchli, Nicolas Heess, and Martin Riedmiller. Continuous-discrete reinforcement learning for hybrid control in robotics. In *Conference on Robot Learning*, pages 735–751. PMLR, 2020. 5.3.2
- [152] James A. Nichols, Hsien W. Herbert Chan, and Matthew A. B. Baker. Machine learning: applications of artificial intelligence to imaging and diagnosis. *Biophysical reviews*, 11(1):111-118, Feb 2019. ISSN 1867-2450. doi: 10.1007/s12551-018-0449-9. URL https://pubmed.ncbi.nlm.nih.gov/ 30182201. 30182201[pmid]. 1.2
- [153] Juhwan Noh, Seoin Back, Jaehoon Kim, and Yousung Jung. Active learning with non-ab initio input features toward efficient co₂ reduction catalysts. *Chem. Sci.*, 9(23):5152–5159, 2018. 3.4, 3-2, 3.6
- [154] J K Nørskov, T Bligaard, J Rossmeisl, and C H Christensen. Towards the computational design of solid catalysts. *Nature Chemistry*, 1(1):37-46, 2009. ISSN 1755-4349. doi: 10.1038/nchem.121. URL https://doi.org/10.1038/nchem.121. 4.2
- [155] Artem R Oganov, Chris J Pickard, Qiang Zhu, and Richard J Needs. Structure prediction drives materials discovery. Nature Reviews Materials, 4(5):331–348, 2019. ISSN 2058-8437. doi: 10.1038/s41578-019-0101-8. URL https://doi.org/10.1038/s41578-019-0101-8. 4.2
- [156] A.G. Olabi, Mohammad Ali Abdelkareem, Tabbi Wilberforce, and Enas Taha Sayed. Application of graphene in energy storage device – a review. *Re-*

newable and Sustainable Energy Reviews, 135:110026, 2021. ISSN 1364-0321. doi: https://doi.org/10.1016/j.rser.2020.110026. URL https://www.sciencedirect.com/science/article/pii/S1364032120303178. 1.1

- [157] Shyue Ping Ong, William Davidson Richards, Anubhav Jain, Geoffroy Hautier, Michael Kocher, Shreyas Cholia, Dan Gunter, Vincent L Chevrier, Kristin A Persson, and Gerbrand Ceder. Python materials genomics (pymatgen): A robust, open-source python library for materials analysis. *Comput. Mater. Sci*, 68:314–319, 2013. 3.3, 3.7
- [158] OpenAI, :, Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, Rafal Józefowicz, Scott Gray, Catherine Olsson, Jakub Pachocki, Michael Petrov, Henrique P. d. O. Pinto, Jonathan Raiman, Tim Salimans, Jeremy Schlatter, Jonas Schneider, Szymon Sidor, Ilya Sutskever, Jie Tang, Filip Wolski, and Susan Zhang. Dota 2 with large scale deep reinforcement learning, 2019. URL https://arxiv.org/abs/1912.06680. 1.3
- [159] Volkan Ortalan, Alper Uzun, Bruce C. Gates, and Nigel D. Browning. Towards full-structure determination of bimetallic nanoparticles with an aberrationcorrected electron microscope. *Nature Nanotechnology*, 5(12):843–847, 2010. ISSN 17483395. doi: 10.1038/nnano.2010.234. 5.2
- [160] Bowei Pan, Fuyi Chen, Bo Kou, Junpeng Wang, Quan Tang, Longfei Guo, Qiao Wang, Zhen Li, Weiqi Bian, and Jiali Wang. Unexpectedly high stability and surface reconstruction of pdauag nanoparticles for formate oxidation electrocatalysis. *Nanoscale*, 12:11659–11671, 2020. doi: 10.1039/D0NR01358G. URL http://dx.doi.org/10.1039/D0NR01358G. 5.2
- [161] Cheol Woo Park and Chris Wolverton. Developing an improved crystal graph convolutional neural network framework for accelerated materials discovery.

Phys. Rev. Materials, 4:063801, Jun 2020. doi: 10.1103/PhysRevMaterials.4.
063801. URL https://link.aps.org/doi/10.1103/PhysRevMaterials.4.
063801. 6

- [162] J. Penfold, R. K. Thomas, I. Tucker, J. T. Petkov, S. D. Stoyanov, N. Denkov, K. Golemanov, S. Tcholakova, and J. R. P. Webster. Saponin adsorption at the air-water interface-neutron reflectivity and surface tension study. *Langmuir*, 34(32):9540–9547, 2018. doi: 10.1021/acs.langmuir.8b02158. URL https://doi.org/10.1021/acs.langmuir.8b02158. 2.2
- [163] Andrew A Peterson, Frank Abild-Pedersen, Felix Studt, Jan Rossmeisl, and Jens K Nørskov. How copper catalyzes the electroreduction of carbon dioxide into hydrocarbon fuels. *Energy & Environ. Sci.*, 3(9):1311–1315, 2010. 3.2
- [164] Jonas H.K. Pfisterer, Yunchang Liang, Oliver Schneider, and Aliaksandr S. Bandarenka. Direct instrumental identification of catalytically active surface sites. *Nature*, 549(7670):74–77, 2017. ISSN 14764687. doi: 10.1038/ nature23661. URL http://dx.doi.org/10.1038/nature23661. 5.2
- [165] James C. Phillips, Rosemary Braun, Wei Wang, James Gumbart, Emad Tajkhorshid, Elizabeth Villa, Christophe Chipot, Robert D. Skeel, Laxmikant Kalé, and Klaus Schulten. Scalable molecular dynamics with NAMD. Journal of Computational Chemistry, 26(16):1781–1802, 2005. ISSN 01928651. doi: 10.1002/jcc.20289. 2.3.2
- [166] Evgeny V. Podryabinkin, Evgeny V. Tikhonov, Alexander V. Shapeev, and Artem R. Oganov. Accelerating crystal structure prediction by machinelearning interatomic potentials with active learning. *Phys. Rev. B*, 99:064114, Feb 2019. doi: 10.1103/PhysRevB.99.064114. URL https://link.aps.org/ doi/10.1103/PhysRevB.99.064114. 4.2, 5.2
- [167] Vladimir Ponec. Alloy catalysts: the concepts. Applied Catalysis A: Gen-

eral, 222(1):31-45, 2001. ISSN 0926-860X. doi: https://doi.org/10.1016/ S0926-860X(01)00828-6. URL https://www.sciencedirect.com/science/ article/pii/S0926860X01008286. Celebration Issue. 5.2

- [168] Mariya Popova, Olexandr Isayev, and Alexander Tropsha. Deep reinforcement learning for de novo drug design. *Science advances*, 4(7):eaap7885, 2018. 5.3.1
- [169] Guorui Ren, Jinfu Liu, Jie Wan, Yufeng Guo, and Daren Yu. Overview of wind power intermittency: Impacts, measurements, and mitigation solutions. Applied Energy, 204:47-65, 2017. ISSN 0306-2619. doi: https:// doi.org/10.1016/j.apenergy.2017.06.098. URL https://www.sciencedirect. com/science/article/pii/S0306261917308346. 1.1
- [170] Ferozuddin Riaz and Khidir M. Ali. Applications of graph theory in computer science. In 2011 Third International Conference on Computational Intelligence, Communication Systems and Networks, pages 142–145, 2011. doi: 10.1109/ CICSyN.2011.40. 1.3
- [171] Jose F. Rodrigues, Larisa Florea, Maria C. F. de Oliveira, Dermot Diamond, and Osvaldo N. Oliveira. Big data and machine learning for materials science. *Discover Materials*, 1(1):12, Apr 2021. ISSN 2730-7727. doi: 10.1007/ s43939-021-00012-0. URL https://doi.org/10.1007/s43939-021-00012-0. 1.2
- [172] Santanu Saha, Luigi Genovese, and Stefan Goedecker. Metastable exohedrally decorated borospherene B₄₀. Scientific Reports, 7:7618, 2017. ISSN 20452322. doi: https://doi.org/10.1038/s41598-017-06877-7. URL https: //www.nature.com/articles/s41598-017-06877-7. 5.3.7
- [173] M. Salmeron and R. Schlögl. Ambient pressure photoelectron spectroscopy: A new tool for surface science and nanotechnology. *Surface Science Reports*, 63 (4):169–199, 2008. ISSN 01675729. doi: 10.1016/j.surfrep.2008.01.001. 5.2

- [174] Jonathan Schmidt, Mário R G Marques, Silvana Botti, and Miguel A L Marques. Recent advances and applications of machine learning in solid-state materials science. *npj Computational Materials*, 5(1):83, 2019. ISSN 2057-3960. doi: 10.1038/s41524-019-0221-0. 1.2, 1.3, 4.3.2
- [175] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897, 2015. 5.3.1, 5.3.2
- [176] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347, 2017. 5.3.1
- [177] K. T. Schütt, P. Kessel, M. Gastegger, K. A. Nicoli, A. Tkatchenko, and K.-R. Müller. Schnetpack: A deep learning toolbox for atomistic systems. *Journal of Chemical Theory and Computation*, 15(1):448–455, 2019. doi: 10.1021/acs.jctc.8b00908. URL https://doi.org/10.1021/acs.jctc.8b00908. 4.2, 5.2
- [178] Kristof T. Schütt, Pieter-Jan Kindermans, Huziel E. Sauceda, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions, 2017. 4.2, 4.4, 5.2
- [179] Zhi Wei Seh, Jakob Kibsgaard, Colin F. Dickens, Ib Chorkendorff, Jens K. Nørskov, and Thomas F. Jaramillo. Combining theory and experiment in electrocatalysis: Insights into materials design. *Science*, 355(6321), 2017. ISSN 0036-8075. doi: 10.1126/science.aad4998. 4.2
- [180] Atsuto Seko, Akira Takahashi, and Isao Tanaka. Sparse representation for a potential energy surface. *Phys. Rev. B*, 90:024101, Jul 2014. doi: 10.1103/PhysRevB.90.024101. URL https://link.aps.org/doi/10.1103/ PhysRevB.90.024101. 1.3

- [181] Zhe Shen and Huai Sun. Prediction of Surface and Bulk Partition of Nonionic Surfactants Using Free Energy Calculations. *The Journal of Physical Chemistry B*, 119(51):15623-15630, 2015. ISSN 1520-6106. doi: 10.1021/acs.jpcb. 5b10239. URL http://pubs.acs.org/doi/10.1021/acs.jpcb.5b10239. 2.2, 2.2
- [182] Chih Jen Shih, Shangchao Lin, Michael S. Strano, and Daniel Blankschtein. Understanding the stabilization of single-walled carbon nanotubes and graphene in ionic surfactant aqueous solutions: Large-scale coarse-grained molecular dynamics simulation-assisted DLVO theory. *Journal of Physical Chemistry C*, 119(2):1047–1060, 2015. ISSN 19327455. doi: 10.1021/jp5093477. 1.2
- [183] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016. 5.3.1
- [184] Gregor Simm, Robert Pinsler, and Jose Miguel Hernandez-Lobato. Reinforcement learning for molecular design guided by quantum mechanics. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 8959–8969. PMLR, 13–18 Jul 2020. URL http://proceedings.mlr.press/v119/simm20b.html. 5.2
- [185] Suyash Singh, Sha Li, Ronald Carrasquillo-Flores, Ana C Alba-Rubio, James A Dumesic, and Manos Mavrikakis. Formic acid decomposition on au catalysts: Dft, microkinetic modeling, and reaction kinetics experiments. AIChE J., 60 (4):1303–1319, 2014. 1.2, 3.2
- [186] Vishnu Sresht, Eric P. Lewandowski, Daniel Blankschtein, and Arben Jusufi.

Combined Molecular Dynamics Simulation-Molecular-Thermodynamic Theory Framework for Predicting Surface Tensions. *Langmuir*, 33(33):8319–8329, 2017. ISSN 15205827. doi: 10.1021/acs.langmuir.7b01073. 2.2, 2.2, 2.3.1, 2.4.1

- [187] C.T.H. Stoddart, R.L. Moss, and D. Pope. Determination of the surface composition of palladium-nickel alloy film catalysts using Auger electron spectroscopy. *Surface Science*, 53(1):241-256, 1975. ISSN 0039-6028. doi: https:// doi.org/10.1016/0039-6028(75)90127-2. URL https://www.sciencedirect. com/science/article/pii/0039602875901272. 5.4.1
- [188] Alberto Striolo and Brian Patrick Grady. Surfactant Assemblies on Selected Nanostructured Surfaces: Evidence, Driving Forces, and Applications. Langmuir, 33(33):8099–8113, 2017. ISSN 15205827. doi: 10.1021/acs.langmuir. 7b00756. 2.2, 2.2
- [189] Dang Sheng Su, Bingsen Zhang, and Robert Schlögl. Electron microscopy of solid catalysts - Transforming from a challenge to a toolbox. *Chemical Reviews*, 115(8):2818–2882, 2015. ISSN 15206890. doi: 10.1021/cr500084c. 5.2
- [190] A Subramanian and L.D. Marks. Surface crystallography via electron microscopy. Ultramicroscopy, 98(2-4):151-157, 2004. ISSN 03043991. doi: 10.1016/j.ultramic.2003.08.008. 5.2
- [191] Richard S Sutton and Andrew G Barto. Reinforcement learning: An introduction. MIT press, 2018. 5.3.1
- [192] Richard S Sutton, David A McAllester, Satinder P Singh, Yishay Mansour, et al. Policy gradient methods for reinforcement learning with function approximation. In NIPs, volume 99, pages 1057–1063. Citeseer, 1999. 5.3.2
- [193] E. B. Tadmor, R. S. Elliott, J. P. Sethna, R. E. Miller, and C. A. Becker. The potential of atomistic simulations and the knowledgebase of interatomic models. *JOM*, 63(7):17–17, Jul 2011. ISSN 1543-1851. doi: 10.1007/

s11837-011-0102-6. URL https://doi.org/10.1007/s11837-011-0102-6. 5.3.5

- [194] Min Tang, Wentao Yuan, Yang Ou, Guanxing Li, Ruiyang You, Songda Li, Hangsheng Yang, Ze Zhang, and Yong Wang. Recent Progresses on Structural Reconstruction of Nanosized Metal Catalysts via Controlled-Atmosphere Transmission Electron Microscopy: A Review. ACS Catalysis, 10(24):14419– 14450, dec 2020. ISSN 2155-5435. doi: 10.1021/acscatal.0c03335. URL https://pubs.acs.org/doi/10.1021/acscatal.0c03335. 5.2
- [195] Xueming Tang, Peter H. Koenig, and Ronald G. Larson. Molecular dynamics simulations of sodium dodecyl sulfate micelles in water-the effect of the force field. The Journal of Physical Chemistry B, 118(14):3864-3880, 2014. doi: 10.1021/jp410689m. URL https://doi.org/10.1021/jp410689m. 2.3.3
- [196] Feng Tao, Michael E Grass, Yawen Zhang, Derek R Butcher, James R Renzas, Zhi Liu, Jen Y Chung, Bongjin S Mun, Miquel Salmeron, and Gabor A Somorjai. Reaction-Driven Restructuring of Rh-Pd and Pt-Pd Core-Shell Nanoparticles. *Science*, 322(5903):932–934, 2008. 5.2
- [197] Feng Tao, David Tang, Miquel Salmeron, and Gabor A. Somorjai. A new scanning tunneling microscope reactor used for high-pressure and high-temperature catalysis studies. *Review of Scientific Instruments*, 79(8), 2008. ISSN 00346748. doi: 10.1063/1.2960569. 5.2
- [198] Feng Tao, Sefa Dag, Lin-Wang Wang, Zhi Liu, Derek R. Butcher, Hendrik Bluhm, Miquel Salmeron, and Gabor A. Somorjai. Break-up of stepped platinum catalyst surfaces by high co coverage. *Science*, 327(5967):850–853, 2010. ISSN 00368075. doi: 10.1126/science.1182122. 5.2
- [199] Franklin Tao and Peter A. Crozier. Atomic-Scale Observations of Catalyst Structures under Reaction Conditions and during Catalysis. *Chemical Reviews*,

116(6):3487–3539, 2016. ISSN 15206890. doi: 10.1021/cr5002657. 5.2

- [200] Ana M. Tarditi and Laura M. Cornaglia. Novel pdagcu ternary alloy as promising materials for hydrogen separation membranes: Synthesis and characterization. Surface Science, 605(1):62–71, 2011. ISSN 0039-6028. doi: https:// doi.org/10.1016/j.susc.2010.10.001. URL https://www.sciencedirect.com/ science/article/pii/S0039602810003985. 5.2
- [201] Ana M. Tarditi, Carolina Imhoff, James B. Miller, and Laura Cornaglia. Surface composition of pdcuau ternary alloys: a combined leis and xps study. *Surface and Interface Analysis*, 47(7):745–754, 2015. doi: https://doi.org/10. 1002/sia.5759. 5.2
- [202] Sascha Thinius, Mazharul M. Islam, and Thomas Bredow. The structure of reconstructed chalcopyrite surfaces. Surface Science, 669:1-9, 2018. ISSN 00396028. doi: https://doi.org/10.1016/j.susc.2017.11.003. URL https:// www.sciencedirect.com/science/article/pii/S0039602817306891. 5.3.7
- [203] Aidan P. Thompson, Laura P. Swiler, Christian R. Trott, Stephen M. Foiles, and Garritt J. Tucker. Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials. *Journal of Computational Physics*, 285, 3 2015. ISSN 0021-9991. doi: 10.1016/j.jcp.2014.12.018. 4.3.1
- [204] Kevin Tran and Zachary W Ulissi. Active learning across intermetallics to guide discovery of electrocatalysts for CO2 reduction and H2 evolution. *Nature Catal*ysis, 1(9):696-703, 2018. ISSN 2520-1158. doi: 10.1038/s41929-018-0142-1.
 URL https://doi.org/10.1038/s41929-018-0142-1. 4.4, 5.2
- [205] Kevin Tran and Zachary W Ulissi. Active learning across intermetallics to guide discovery of electrocatalysts for co₂ reduction and h₂ evolution. Nat. Catal., 1(9):696, 2018. (document), 3.2, 3-5, 3.5.1, 3.5.2, 3.7
- [206] Naga Rajesh Tummala, Brian P. Grady, and Alberto Striolo. Lateral con-

finement effects on the structural properties of surfactant aggregates: SDS on graphene. *Physical Chemistry Chemical Physics*, 12(40):13137–13143, 2010. ISSN 14639076. doi: 10.1039/c0cp00600a. 2.2, 2.4.2

- [207] Evrim Umut. Surface Modification of Nanoparticles Used in Biomedical Applications. Modern Surface Engineering Treatments, 8:185–208, 2013. ISSN 9789533070865. doi: 10.5772/52807. 2.2
- [208] S. Vera-López, P. Martínez, M. P. San Andrés, A. M. Díez-Pascual, and M. Valiente. Study of graphene dispersions in sodium dodecylsulfate by steady-state fluorescence of pyrene. *Journal of Colloid and Interface Science*, 514:415–424, 2018. ISSN 10957103. doi: 10.1016/j.jcis.2017.12.052. 2.2, 2.4.2
- [209] L. Vitos, A.V. Ruban, H.L. Skriver, and J. Kollár. The surface energy of metals. Surface Science, 411(1):186-202, 1998. ISSN 0039-6028. doi: https://doi.org/ 10.1016/S0039-6028(98)00363-X. URL https://www.sciencedirect.com/ science/article/pii/S003960289800363X. 5.4.1
- [210] Shengguang Wang, V Petzold, Vladimir Tripkovic, Jesper Kleis, Jakob Geelmuyden Howalt, Egill Skulason, EM Fernandez, B Hvolbæk, Glenn Jones, and Anja Toftelund. Universal transition state scaling relations for (de) hydrogenation over transition metals. *Phys. Chem. Chem. Phys.*, 13(46):20760–20765, 2011. 3.2
- [211] Siwen Wang, Noushin Omidvar, Emily Marx, and Hongliang Xin. Overcoming site heterogeneity in search of metal nanocatalysts. ACS Comb. Sci., 20(10): 567–572, 2018. 3.6
- [212] Yuyang Wang, Zhonglin Cao, and Amir Barati Farimani. Deep reinforcement learning optimizes graphene nanopores for efficient desalination. arXiv preprint arXiv:2101.07399, 2021. 5.3.1
- [213] Zhi-Li Wang, Yun Ping, Jun-Min Yan, Hong-Li Wang, and Qing Jiang. Hy-

drogen generation from formic acid decomposition at room temperature using a niaupd alloy nanocatalyst. *International Journal of Hydrogen Energy*, 39(10):4850-4856, 2014. ISSN 0360-3199. doi: https://doi.org/10.1016/ j.ijhydene.2013.12.148. URL https://www.sciencedirect.com/science/ article/pii/S0360319913031406. 5.2

- [214] Erica J. Wanless and William A. Ducker. Organization of sodium dodecyl sulfate at the graphite-solution interface. *Journal of Physical Chemistry*, 100 (8):3207–3214, 1996. ISSN 00223654. doi: 10.1021/jp952439x. 2.4.2, 2.4.2
- [215] Stefan Wendt, Phillip T. Sprunger, Estephania Lira, G. K. H. Madsen, Zheshen Li, J. O. Hansen, Jesper Matthiesen, Asger Blekinge-Rasmussen, E. Laegsgaard, Bjørk Hammer, and Flemming Besenbacher. The Role of Interstitial Sites in the Ti3d Defect State in the Band Gap of Titania. *Science*, 320(5884): 1755–1759, jun 2008. ISSN 0036-8075. doi: 10.1126/science.1159846. URL https://www.sciencemag.org/lookup/doi/10.1126/science.1159846. 5.2
- [216] Eugene Wigner and Frederick Seitz. On the constitution of metallic sodium. Phys. Rev., 43(10):804, 1933. 3.3
- [217] Frank L. Williams and M. Boudart. Surface composition of nickel-gold alloys. Journal of Catalysis, 30(3):438-443, 1973. ISSN 0021-9517. doi: https://doi. org/10.1016/0021-9517(73)90161-9. URL https://www.sciencedirect.com/ science/article/pii/0021951773901619. 5.4.1
- [218] Scott M Woodley and Richard Catlow. Crystal structure prediction from first principles. Nature Materials, 7(12):937-946, 2008. ISSN 1476-4660. doi: 10. 1038/nmat2321. URL https://doi.org/10.1038/nmat2321. 4.2
- [219] Chung Yeh Wu, William J. Wolf, Yehonatan Levartovsky, Hans A. Bechtel, Michael C. Martin, F. Dean Toste, and Elad Gross. High-spatial-resolution mapping of catalytic reactions on single particles. *Nature*, 541(7638):511–515,

2017. ISSN 14764687. doi: 10.1038/nature20795. URL http://dx.doi.org/ 10.1038/nature20795. 5.2

- [220] Meiyin Wu and Li Chen. Image recognition based on deep learning. In 2015 Chinese Automation Congress (CAC), pages 542–546, 2015. doi: 10.1109/ CAC.2015.7382560. 1.3
- [221] Yihan Xiao, Yan Wang, Shou-Hang Bo, Jae Chul Kim, Lincoln J. Miara, and Gerbrand Ceder. Understanding interface stability in solid-state batteries. Nature Reviews Materials, 5(2):105–126, Feb 2020. ISSN 2058-8437. doi: 10.1038/ s41578-019-0157-5. URL https://doi.org/10.1038/s41578-019-0157-5.
 1.1
- [222] Chenlu Xie, Zhiqiang Niu, Dohyung Kim, Mufan Li, and Peidong Yang. Surface and interface control in nanoparticle catalysis. *Chemical Reviews*, 120(2):1184–1249, 2020. doi: 10.1021/acs.chemrev.9b00220. URL https://doi.org/10.1021/acs.chemrev.9b00220. PMID: 31580651. 1.1
- [223] Tian Xie and Jeffrey C. Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.*, 120:145301, Apr 2018. doi: 10.1103/PhysRevLett.120.145301. URL https://link.aps.org/doi/10.1103/PhysRevLett.120.145301. 1.3, 5.2
- [224] Tian Xie and Jeffrey C. Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.*, 120:145301, Apr 2018. doi: 10.1103/PhysRevLett.120.145301. URL https://link.aps.org/doi/10.1103/PhysRevLett.120.145301. 4.2, 4.3.1, 4-1
- [225] Tian Xie and Jeffrey C. Grossman. Hierarchical visualization of materials space with graph convolutional neural networks. *The Journal of Chemical Physics*, 149(17):174111, Nov 2018. ISSN 1089-7690. doi: 10.1063/1.5047803. URL

http://dx.doi.org/10.1063/1.5047803. 4.2

- [226] Tian Xie and Jeffrey C Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.*, 120(14):145301, 2018. (document), 3.2, 3.3, 3-1, 3.4, 3.6
- [227] Tian Xie, Arthur France-Lanord, Yanming Wang, Yang Shao-Horn, and Jeffrey C. Grossman. Graph dynamical networks for unsupervised learning of atomic scale dynamics in materials. *Nature Communications*, 10(1), Jun 2019. ISSN 2041-1723. doi: 10.1038/s41467-019-10663-6. URL http://dx.doi.org/ 10.1038/s41467-019-10663-6. 4.2
- [228] Jie Xiong, San-Qiang Shi, and Tong-Yi Zhang. A machine-learning approach to predicting and understanding the properties of amorphous metallic alloys. *Materials & Design*, 187:108378, 2020. 5.2
- [229] C. W. Yi, K. Luo, T. Wei, and D. W. Goodman. The composition and structure of PdAu surfaces. *The Journal of Physical Chemistry B*, 109(39):18535–18540, 2005. doi: 10.1021/jp053515r. URL https://doi.org/10.1021/jp053515r. PMID: 16853387. 5.4.1
- [230] Chunrong Yin, Zhitao Guo, and Andrew J. Gellman. Surface segregation across ternary alloy composition space: Cu_xAu_yPd_{1-x-y}. The Journal of Physical Chemistry C, 124(19):10605-10614, 2020. doi: 10.1021/acs.jpcc.0c02058.
 URL https://doi.org/10.1021/acs.jpcc.0c02058. 5.4.1
- [231] Brian Yoo, Eliseo Marin-Rimoldi, Ryan Gotchy Mullen, Arben Jusufi, and Edward J. Maginn. Discrete fractional component monte carlo simulation study of dilute nonionic surfactants at the air-water interface. *Langmuir*, 33(38):9793– 9802, 2017. doi: 10.1021/acs.langmuir.7b02058. URL https://doi.org/10. 1021/acs.langmuir.7b02058. 2.2
- [232] Jong Suk Yoo, Frank Abild-Pedersen, Jens K Nørskov, and Felix Studt. The-

oretical analysis of transition-metal catalysts for formic acid decomposition. ACS Catal., 4(4):1226–1233, 2014. 3.2

- [233] Junwoong Yoon and Zachary W. Ulissi. Differentiable optimization for the prediction of ground state structures (dogss). *Phys. Rev. Lett.*, 125:173001, Oct 2020. doi: 10.1103/PhysRevLett.125.173001. URL https://link.aps.org/doi/10.1103/PhysRevLett.125.173001. 1.3, 5.2
- [234] Shiran Zhang, Luan Nguyen, Yuan Zhu, Sihui Zhan, Chia Kuang Frank Tsung, and Franklin Feng Tao. In-situ studies of nanocatalysis. Accounts of Chemical Research, 46(8):1731–1739, 2013. ISSN 00014842. doi: 10.1021/ar300245g. 5.2
- [235] Zixing Zhang, Jürgen Geiger, Jouni Pohjalainen, Amr El-Desoky Mousa, Wenyu Jin, and Björn Schuller. Deep learning for environmentally robust speech recognition: An overview of recent developments. ACM Trans. Intell. Syst. Technol., 9(5), apr 2018. ISSN 2157-6904. doi: 10.1145/3178115. URL https://doi.org/10.1145/3178115. 1.3
- [236] Meng Zhao, J.C. Brouwer, Wim G. Sloof, and AJ Bottger. Surface segregation of ternary alloys: Effect of the interaction between solute elements. Advanced Materials Interfaces, 7:1901784, 02 2020. doi: 10.1002/admi.201901784. 5.2
- [237] X. Zheng, C. Bi, M. Brooks, and David S. Hage. Analysis of hormone-protein binding in solution by ultrafast affinity extraction: interactions of testosterone with human serum albumin and sex hormone binding globulin. Anal Chem., 25(4):368–379, 2015. ISSN 1527-5418. doi: 10.1016/j.cogdev.2010.08.003. Personal. 2.2
- [238] Zhenpeng Zhou, Xiaocheng Li, and Richard N Zare. Optimizing Chemical Reactions with Deep Reinforcement Learning. ACS Central Science, 3(12): 1337-1344, dec 2017. ISSN 2374-7943. doi: 10.1021/acscentsci.7b00492. URL https://doi.org/10.1021/acscentsci.7b00492. 5.2

- [239] Zhenpeng Zhou, Steven Kearnes, Li Li, Richard N Zare, and Patrick Riley. Optimization of molecules via deep reinforcement learning. *Scientific reports*, 9(1):1–10, 2019. 5.3.1
- [240] C. Lawrence Zitnick, Lowik Chanussot, Abhishek Das, Siddharth Goyal, Javier Heras-Domingo, Caleb Ho, Weihua Hu, Thibaut Lavril, Aini Palizhati, Morgane Riviere, Muhammed Shuaibi, Anuroop Sriram, Kevin Tran, Brandon Wood, Junwoong Yoon, Devi Parikh, and Zachary Ulissi. An introduction to electrocatalyst design using machine learning for renewable energy storage, 2020. URL https://arxiv.org/abs/2010.09435. 1.1, 1.2