

Rethinking Computational Catalyst Searches with Alchemical Perturbation
Density Functional Theory

by

Charles D. Griego

Bachelors of Science in Chemical Engineering, New Mexico Institute of Mining and
Technology, 2017

Submitted to the Graduate Faculty of
the Swanson School of Engineering in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy

University of Pittsburgh

2022

UNIVERSITY OF PITTSBURGH
SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

Charles D. Griego

It was defended on

April 1, 2022

and approved by

John A. Keith, Ph.D., Associate Professor,

Department of Chemical and Petroleum Engineering

J. Karl Johnson, Ph.D., Professor,

Department of Chemical and Petroleum Engineering

John R. Kitchin, Ph.D., Professor,

Department of Chemical Engineering, Carnegie Mellon University

Goetz Vesper, Ph.D., Professor,

Department of Chemical and Petroleum Engineering

Jill E. Millstone, Ph.D., Professor,

Department of Chemistry

Dissertation Director: John A. Keith, Ph.D., Associate Professor,

Department of Chemical and Petroleum Engineering

Copyright © by Charles D. Griego
2022

Rethinking Computational Catalyst Searches with Alchemical Perturbation Density Functional Theory

Charles D. Griego, PhD

University of Pittsburgh, 2022

The expense of quantum chemistry calculations significantly hinders endeavors to search through diverse materials space for catalysts that promote sustainable production of energy and chemicals. Motivated by this challenge, we report our study on alchemical perturbation density functional theory (APDFT), an easy and highly cost-efficient calculation scheme that enables high-throughput computational screening of hypothetical catalysts. APDFT requires just a small number of reference DFT calculations to approximate quantities such as adsorbate binding energies (BE) and reaction barriers on a large numbers of hypothetical catalyst surfaces by employing simple arithmetic manipulations to electrostatic potentials. In this dissertation, we discuss how APDFT can be used to rapidly predict catalyst descriptors from numerous atomic transmutations done to a single reference catalyst, and we address multiple factors that influence the accuracy of these predictions. We first demonstrate that first order APDFT predicts adsorbate BE on many variations of carbide, nitride, and oxide catalysts in close agreement with DFT results, and we determined that predictions based in metallic systems are most accurate. Additionally, first order APDFT reliably predicts many energy profiles and barrier heights using a single nudged elastic band calculation for CH₄ dehydrogenation on Pt(111). Machine learning models trained on correlations between APDFT errors and reference system properties produced BE prediction corrections over multiple classes of adsorbates at multiple coverages on hypothetical Pt alloys. We further uncover these correlations by revisiting multiple catalyst systems with second order APDFT approximations from VASP and CP2K data. Finally, we introduce ways to produce alchemical energy functions that help illustrate the agreement between reference system with differing characteristics and increasing orders of APDFT. We find that there are greater limitations with VASP for second order APDFT, but results from CP2K show promising advances with APDFT for widespread applications in computational materials science.

Table of Contents

Preface	x
1.0 Introduction	1
1.1 Conventional Computational Catalysis	4
1.2 Alchemical Perturbation Density Functional Theory	8
2.0 Benchmarking Computational Alchemy for Carbide, Nitride, and Oxide Catalysts	18
2.1 Introduction	18
2.2 Computational Methodology	19
2.3 Results and Discussion	20
2.4 Conclusion	24
3.0 Assessing APDFT Predicted Reaction Pathways and Barriers	26
3.1 Introduction	26
3.2 Results and Discussion	27
3.3 Conclusion	30
4.0 Machine Learning Corrected Alchemical Perturbation Density Functional Theory for Catalysis Applications	31
4.1 Introduction	31
4.2 Computational Methodology	32
4.3 ML Model	35
4.4 Results and Discussion	38
4.5 Conclusion	42
5.0 Second Order APDFT: Methods and Applications for Heterogeneous Catalysis	44
5.1 Introduction	45
5.2 Taylor Series Approximation with Alchemical Derivatives	48
5.3 Modelling Non-integer Nuclear Charges	50

5.3.1	VASP	51
5.3.2	CP2K	51
5.4	Modelled Materials	52
5.5	Computational Details	53
5.6	Results and Discussion	55
5.6.1	Overall Performance of APDFT	55
5.6.2	Alchemical Binding Energy Curves	58
5.7	Conclusion	64
6.0	Conclusions	66
6.1	Future Work	67
Appendix A. Supporting Information for Machine Learning Corrected APDFT		70
A.1	Numbers of Systems Evaluated	70
A.2	Feature Vector Construction	71
A.3	Model Validation	73
A.4	NH _x Parity Plots	75
A.5	OH _x and OOH Parity Plots	76
A.6	Error Analysis by Alloy Type	77
Appendix B. Supporting Information for Second Order APDFT Studies . .		78
Bibliography		87

List of Tables

1	Summary of machine learning model evaluation and selection based on the CV MAEs for each dataset.	38
2	Number of systems evaluated for each adsorbate class and partition of data sets	71
3	Counts of hypothetical alloy types in test data sets	77

List of Figures

1	Heatmap of BEs for multiple configurations of CO and H adsorption on Pd-skinned Pd ₃ X alloys (X = Co, Ni, Cu, Ag)	6
2	Thermodynamic cycle depicting atomic transmutations to the states in a BE calculation	11
3	Illustrations of steps to compute corrections to BEs based on APDFT.	14
4	Parity plot of APDFT BE predictions of OH adsorption on Pt alloys	16
5	Parity plots comparing the accuracy of APDFT BE predictions against DFT for TiC, TiN, and TiO surfaces	21
6	Parity plots depicting the accuracy of APDFT BE predictions for ZnO, SnO ₂ , and TiO ₂ with DOS plots	22
7	Parity plots depicting the accuracy of APDFT BE predictions for TiO ₂ pristine, Pt@2L, and Pt@1L with DOS plots	24
8	Snapshots of the dehydrogenation of CH ₄ adsorbed on Pt	28
9	Energy profiles and a parity plot of energy barriers for the CH ₄ dehydrogenation mechanism on alloys of Pt for APDFT predictions	29
10	Bar plots of APDFT prediction MAE with respect to ΔZ , N_T , and θ	34
11	MAE with respect to hydrogenation of the central atom for CH _x , NH _x , and OH _x adsorbate BE.	35
12	Machine learning workflow for obtaining corrected alchemy BE predictions.	36
13	Learning curves comparing ML model performance between APDFT and DFT datasets	39
14	Parity plots of APDFT BE predictions with and without ML corrections	40
15	Feature importance plot showing the significance of different parameters for the machine learning model	41
16	Visualization of the distribution of hypothetical alloy types in the test data set and how well ML improves APDFT predictions for those alloy types	42

17	MAE of APDFT-1 and APDFT-2 predictions from VASP and CP2K data . . .	56
18	Alchemical binding energy curves of H on selected transmuted variations of Pt (111)	58
19	Alchemical binding energy curves of H on selected transmuted variations of TiO ₂ -based systems from VASP data	60
20	Alchemical binding energy curves of H on selected transmuted variations of TiO ₂ -based systems from CP2K data	62
21	Fingerprinting scheme to construct feature vectors that describe dopant locations.	72
22	Fingerprinting scheme that describes dopant locations in Pt slab models with adsorbates.	73
23	Parity plot of ΔBE of NH _x adsorbates on alloys of transmuted Pt predicted with APDFT and ML-corrected APDFT compared to DFT	75
24	Parity plot of ΔBE of OH _x adsorbates on alloys of transmuted Pt predicted with APDFT and ML-corrected APDFT compared to DFT	76
25	Alchemical binding energy curves of H on all transmuted variations of Pt (111) (VASP)	78
26	Alchemical binding energy curves of H on all transmuted variations of TiO ₂ (VASP)	79
27	Alchemical binding energy curves of H on all transmuted variations of (Pt@1L)-TiO ₂ (VASP)	80
28	Alchemical binding energy curves of H on all transmuted variations of (Pt@2L)-TiO ₂ (VASP)	81
29	Alchemical binding energy curves of H on all transmuted variations of Pt (111) (CP2K)	82
30	Alchemical binding energy curves of H on all transmuted variations of TiO ₂ (CP2K)	83
31	Alchemical binding energy curves of H on all transmuted variations of (Pt@1L)-TiO ₂ (CP2K)	84
32	Alchemical binding energy curves of H on all transmuted variations of (Pt@2L)-TiO ₂ (CP2K)	85
33	Alchemical energy curves of $\Delta Z = \pm 1, 2$ transmutations to dimers	86

Preface

First I would like to thank my advisor Dr. John Keith. Joining his group was one of the most influential decisions in my life. I knew from the day in his office when he presented a plan for me to explore alchemy, that this project and his mentorship were a perfect fit. Over the years, he highlighted my strongest qualities and helped me lean into those strengths, and I will always be grateful for the time he took to talk me through the moments when I felt stuck or doubted myself.

I would like to thank my committee members Dr. Karl Johnson, Dr. John Kitchin, Dr. Goetz Vesper, and Dr. Jill Millstone. Dr. Johnson was my mentor during a summer REU in 2016, and he showed me how intriguing and challenging it could be to explore open-ended research questions. We reconnected at a conference later that year, when I was searching for grad schools, and he helped me decide to pursue Pitt. I must also recognize Dr. Kitchin for all the knowledge and advice that I was privileged to receive.

From my time at New Mexico Tech, I owe many thanks to Dr. Pabitra Choudhury, who saw my passion for coding and showed me the world of computational research. He also played a major role in my journey to Pittsburgh by encouraging me to apply to an REU and work with Dr. Johnson, his former postdoc advisor. I can't forget about Dr. Thomas C. Bickel, who inspired me as a teacher and scientist while putting up with my esoteric humor. I'm forever grateful for the best of friends I made at Tech, Andrew Petersen, Joe Kerwin, Aaron Melad, Emily Stirrup, Danielle Richards, Gabe Montoya, and Thalia Natzic.

I would like to acknowledge all of the former and current members of the Keith group. I want to first thank Dr. Mitchell Groenenboom and especially Dr. Karthikeyan Saravanan for helping me transition smoothly into the group. I'm thankful for the time I spent in the group with Dr. Yasemin Basdogan, who became a great friend that always looked out for me as I navigated this part of my life.

I'm incredibly grateful for Alex Maldonado, who became one of my best friends within days of meeting, and I have to thank him for always being so selfless and supportive and for sharing his time with me, whether it working together, going to lunch, or just hanging

out and talking about TV shows. I also want to thank many others that I met through the Keith group, Dr. Emily Eikey, Lingyan Zhao, Tae Hoon Choi, Barbaro Zulueta, Dominick Filonowich, and Chinmay Matre.

Finally, I want to thank the love of my life, my fiancé Margaret Herzog. We met for our first date on Walnut Street in the summer of 2019, just a few blocks away from the apartment where we shared practically all of our days together during the height of a pandemic. While it was a very challenging time for all of us, I am thankful for every single day I spent with her, developing an unimaginable bond that solidified us as a pair devoted to supporting and caring for each other for the rest of our lives. We now live a comfortable life with our cats Spooky and Tia in a new apartment that is still just a short walk from where we met, and we are very excited to get married and spend the next chapters of our lives together.

And of course, I cannot end this without acknowledging the love and support from my Mom and Dad, my brother Reggie, my sister Ali, and the rest of my family in New Mexico. I miss them all the time, but despite the distance, they always supported me through my PhD and I know I will always have that, regardless of where I go or what I do, to achieve my dreams.

1.0 Introduction

Computational catalysis plays a growingly important role in guiding the design of new and improved materials for sustainable catalysis.[1, 2, 3, 4] Many efforts in this field involve screening materials with high level quantum chemistry calculation schemes to find innovative catalysts that promote high chemical activity.[5, 6, 4, 7] Descriptors of catalyst activity, such as an adsorbate binding energy (BE) on a catalyst surface, are obtained using Kohn-Sham density functional theory (DFT) calculations because they are both reasonably accurate and computationally feasible for studying systems that might contain up to about 200 atoms. However, even relatively efficient methods such as DFT are too computationally costly to systematically model many different adsorbates at different coverages on different facets and compositions of hypothetical catalysts.

To obtain useful and insightful predictions, computational screening studies must be reproducible while also i) determining important active sites that are stable under specified environmental conditions on large numbers of material compositions and ii) elucidating important elementary reaction steps with barrier heights that are needed for kinetic modeling. To date, most computational catalysis studies can address one of these points at a time,[5, 4, 6] but new advances are needed to make it possible to address both points at the same time. Even though the computational cost of DFT calculations continues to decrease to allow even more calculations possible in the future, one should not expect that the two points above can be addressed at the same time using DFT without invoking severe approximations.

There have been efforts to address such points by developing approximate models such as cluster expansions,[8, 9, 10] coordination number models,[11, 12, 13] as well as an assortment of machine learning models[14, 15, 16, 17, 18, 6, 19] that all can be trained from quantum mechanics calculations but leveraged to make predictions on systems outside of the training set. An alternative approach that technically requires no a priori model training is alchemical perturbation density functional theory (APDFT).[20, 21, 22, 23, 24] The method has previously been called “computational alchemy”, but more recently termed APDFT to

better distinguish it to other forms of “computational alchemy”.[25] APDFT is a means to predict a hypothetical energy contribution based on a perturbation on a reference system. In our APDFT scheme, the hypothetical energy contribution is based on an approximate Taylor series expansion truncated to a number of terms that contain alchemical derivatives, which are differences in atomic electrostatic potentials. When given a known adsorbate BE as well as the alchemical derivatives that reflect how electrostatic potentials of the system change upon adsorption, one can intuit a BE by approximating the alchemical energy changes for a hypothetical material.[26, 27] A key advantage of APDFT is that no additional information is needed besides a fixed number of DFT computations, and once these are obtained, large numbers of adsorbate BEs for hypothetical surface structures can be obtained with effectively no computational cost.

In this dissertation, we report our findings on applications of APDFT in heterogeneous catalysis, predicting descriptors like adsorbate BE and reaction barrier heights on diverse classes of materials, and we assess how effective and reliable it will be for computational chemists and materials scientists to use APDFT coupled with other research methodologies. We first provide a brief background on conventional computational catalysis, and how current challenges call for robust methods that allow rapid property predictions using physical descriptors of materials without a prerequisite to obtain large numbers of DFT-calculated results. Then we provide a detailed introduction on APDFT that includes an outline on theoretical foundations, a summary of various formulations of quantum alchemy that align with APDFT, and a practical tutorial that the average computational scientist of heterogeneous catalysis can follow to produce APDFT predictions from the results of common periodic DFT software packages.

In the following chapters, we begin to summarize our work on APDFT applications. First, by employing our APDFT schemes on BEs of oxygen reduction reaction intermediates on catalysts beyond transition metal alloys, we report BE predictions that are in close agreement with DFT results on many variations of rocksalt TiC(111), TiN(100), and TiO(100) materials, and we identify that APDFT performs less reliably for band-gap materials like semiconductors. The next chapter covers our work benchmarking applications in kinetic modelling, where we predicted barrier heights for methane dehydrogenation on 32 alloy vari-

ations of Pt(111) with APDFT, relying on a single nudged elastic band (NEB) calculation as a reference case. Shifting focus on methods to improve APDFT and better understand why accuracy drops for certain systems, we discuss machine learning (ML) coupled APDFT in the following chapter. We trained support vector regression machine learning (ML) models on data sets of APDFT predicted BEs of carbon, nitrogen, and oxygen based adsorbates at different coverages on hypothetical alloys based on a Pt(111). These models capably predicted APDFT errors, which were observed to follow very systematic trends. The ML model predictions were then applied as corrections to the APDFT BE predictions on hypothetical alloys, decreasing the mean absolute error within our test sets by as much as an order of magnitude. We then identified variables about our target alloys and how they related to the large APDFT errors of BE on these alloys.

While machine learning was an insightful approach to find the correlations between prediction accuracy and reference and target materials, we cover second order APDFT approximations and alchemical energy pathways in the final chapter to highlight reasons why APDFT errors are correlated to types of applied systems. We re-evaluated select metallic and semiconducting systems using both first (APDFT-1) and second-order APDFT (APDFT-2), calculating values of the energy function at many points along the alchemical pathways, and comparing with DFT data at every point. Done with both VASP and CP2K data, we identified that systems that aren't well-described with APDFT-1 manifest alchemical functions that are high-order polynomials. APDFT-2 approximations become more accurate for some systems when done with CP2K, but VASP-based APDFT-2 predictions further decrease in accuracy. From these results, we discuss the advantages and disadvantages of using either code and how the computational catalysis community may continue to utilize APDFT to expedite their studies.

1.1 Conventional Computational Catalysis

Catalytic reaction mechanisms depend on the interplay of reaction thermodynamics and kinetics. These can often be modeled quite well with standard DFT calculations, but bringing this level of accuracy already brings substantial computational cost. Systems that are more complicated will be expected to require higher level quantum chemistry (QC) calculations that bring even greater costs.

The thermodynamics of elementary reaction steps can be assessed by computing reaction intermediate BE. Each BE determination normally requires three separate QC electronic energy calculations, each preferably modeled using geometrically relaxed structures, and suitably accurate levels of theory. Appropriate zero-point energy, thermal, and entropy corrections can be included as well, but these are sometimes neglected for simplicity when analyzing trends of thermodynamic descriptors. The required calculations include the model of the catalyst surface without the adsorbate (site), a model of the adsorbate not interacting with the catalyst (ads), and a model for the catalyst surface with the adsorbate bound (ads-site). The BE is then calculated with the following equation:

$$BE = E_{\text{site}} + E_{\text{ads}} - E_{\text{ads-site}} \quad (1.1)$$

In this convention, positive binding energies indicate thermodynamically favorable adsorptions. This equation clearly highlights a problem with high throughput screening studies that rely on QC calculations. Any BE calculation for a single adsorbate on any hypothetical surface site usually requires two separate QC calculations (for $E_{\text{ads-site}}$ and E_{site}), and those might take anywhere from minutes to weeks to complete on modern super computers. Predictive BEs also may require more considerations of multiple adsorbate configurations as well as adsorbate-induced surface reconstruction and segregation.

In a collaboration with Joshua Snyder’s group at Drexel University, we showed that it is usually necessary to investigate multiple thermodynamically accessible states to better interpret experimental observations.[28] In this work, nanoporous PdX skin alloys (X = Co, Ni, Cu, and Ag) were presented as electrocatalysts that produced formate from CO₂ with high selectivity and avoided deactivation from CO poisoning. Among these alloys,

Pd-skin/Pd₃Co was found to be most promising, which our group confirmed by evaluating stable configurations of CO and H binding on the surface with DFT. To assess the impact of alloying and hydrogenation on CO binding, we determined the low energy adsorption configurations for CO* and H* on Pd and the Pd-skinned alloys.

Reliably quantifying BEs for these systems under realistic electrochemical conditions can be very challenging, but qualitative trends in relative BEs for CO* and H* on a variety of hydrogenated configurations were able to provide physical insights into the experimental observations. We hypothesized that the most active electrocatalysts would exhibit weaker H binding to better facilitate hydrogenations of CO₂ into formate as well as weaker CO binding that would result in less CO poisoning. Results in Figure 1 show that adsorbed CO and H were both the most destabilized on Pd-skin/Pd₃Co, regardless of the co-adsorbed species and helped explain the exceptional operational stability of nanoporous Pd-skin/Pd₃Co. This served as an interesting example of the fact that untangling catalyst descriptors can be complex and may require numerous DFT computations.

Furthermore, relatively efficient DFT calculations for N electron systems typically scale as N^3 or greater, making these calculations quite expensive. This unfavorable scaling is especially significant when calculating E_a . The Nudged Elastic Band (NEB) algorithm,[29] which is now a standard approach for predicting E_a , works by computing a series of constrained geometric ‘images’ that are eventually optimized to a minimum energy pathway along a potential energy surface. Since standard NEB calculations generally require the optimization of 10–20 images per pathway, they bring significant computational cost that limits applications in high-throughput screening. A conventional approach to address this is with linear scaling relations,[30, 31, 32] but an open question is whether alternative approaches might be more accurate and useful.

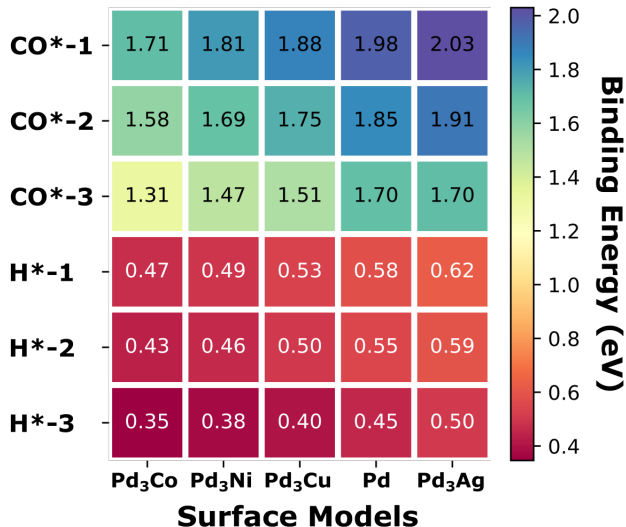


Figure 1: Tabulated above are binding energies (in eV) on Pd-skin Pd₃X alloys (X = Co, Ni, Cu, Ag) for CO binding to a clean surface (CO*-1), CO binding to a surface with H* (CO*-2), CO binding to a surface with 2 H* (CO*-3), H binding to a clean surface (H*-1), H binding to a surface with H* (H*-2), and H binding to a surface with CO* (H*-3).

Some approaches that have been developed include approximate models such as cluster expansions that model adsorbate interactions using an expanded series of polynomials that describe the energy of Ising-type interactions,[8, 9, 10] models that describe adsorptive scalability based on the coordination number of the active site,[11, 12, 13] and applications of machine learning models to catalyst data sets, learning statistical trends based on various features/fingerprints of catalyst models.[14, 15, 16, 17, 18, 6, 19] However, a downside of most of these machine learning approaches is that they need a lot of data for training, they can unreliably extrapolate to new systems, and they are often harder to intuitively or physically interpret.

One could also expedite catalyst screening using a more physically grounded approach like the d-band model,[33] which relates an adsorbate’s binding energy (BE) with the position of the catalyst surface’s d-band center. If a shift in the d-band center from a reference state

to a hypothetical one is known, that shift would correlate with the difference between an adsorbate’s BE in the two cases. While elegant and useful, the d-band model does not satisfactorily predict trends in calculated BEs for electronegative adsorbates such as OH, F, and Cl on materials having mostly filled d-states.[34] Furthermore, the d-band model can only be used on systems that have significant adsorption energy contributions that arise from d-orbitals, e.g. transition metal systems. For other classes of materials, extensions to the d-band model have been developed.[35, 36, 37]

Compared to these methods listed above, APDFT has a promising advantage to accelerate computational catalysis screening studies because there is technically no need for model training and, analogous to the d-band model, APDFT correlates an adsorbate’s BE to the material’s electrostatic potentials,[20, 21, 22] which in effect reflects an amalgam of the material’s full band structure. When given a known adsorbate BE as well as the alchemical derivative that reflects how electrostatic potentials of the system change upon adsorption, one can work out a BE by approximating the alchemical derivatives for a hypothetical material.[26, 27] The key advantage to APDFT is that no additional information is needed besides a single DFT-quality adsorbate BE, and once this is obtained, large numbers of adsorbate BEs for hypothetical surface structures can be obtained with effectively no computational cost.

1.2 Alchemical Perturbation Density Functional Theory

In this section, we present background, derivations, and a simple example of APDFT. Advocated by von Lilienfeld and co-workers,[20, 21, 38, 23] APDFT uses gradient-supported methods based on perturbation theory to measure a change in a specific materials property in response to a change in the material’s composition without the explicit QC calculation of the property itself. When referring to a change in a material’s composition, we technically mean an alchemical perturbation or alchemical transmutation, which translates to a change in nuclear charge of an atom in a reference material (R) that results in a new target material (T). In APDFT, the Hamiltonians (\hat{H}) of the reference and target material are linearly coupled through a mixing parameter, λ , as the following:

$$\hat{H}(\lambda) = \lambda\hat{H}_T + (1 - \lambda)\hat{H}_R \quad (1.2)$$

where $0 \leq \lambda \leq 1$, describing the size of the perturbation. von Lilienfeld also showed that these Hamiltonians may be coupled with a quadratic interpolation function, yielding superior results for some applications.[39] Next, the desired quantity resulting from this coupled Hamiltonian is the energy of the target system at $\lambda = 1$ (E^T), which can be expanded in a Taylor series around the reference material at $\lambda = 0$:

$$E^T = \sum_{n=0}^{\infty} \frac{1}{n!} \frac{\partial^n E(\lambda)}{\partial \lambda^n} \Big|_{\lambda=0} = E^R + \sum_{n=1}^{\infty} \frac{1}{n!} \frac{\partial^n E(\lambda)}{\partial \lambda^n} \Big|_{\lambda=0} \quad (1.3)$$

The first derivative from above can be evaluated via the Hellmann-Feynman theorem:[40]

$$\frac{\partial E}{\partial \lambda} \Big|_{\lambda=0} = \langle \psi_R | \hat{H}_T - \hat{H}_R | \psi_R \rangle = \int_{\Omega} d\mathbf{r} (\nu_T(\mathbf{r}) - \nu_R(\mathbf{r})) \rho(\mathbf{r}, \lambda) \quad (1.4)$$

which include the external coulomb potentials ν_R and ν_T corresponding to reference and target systems, respectively. We can then write Eqn. 1.4 in terms of higher order derivatives of the electron density with respect to the perturbations:[23]

$$\frac{\partial^{n+1} E}{\partial \lambda^{n+1}} = \int_{\Omega} d\mathbf{r} (\nu_T(\mathbf{r}) - \nu_R(\mathbf{r})) \frac{\partial^n \rho(\mathbf{r}, \lambda)}{\partial \lambda^n} \quad (1.5)$$

Finally, by combining Eqns. 1.3 and 1.5, we obtain:

$$E^T = E^R + \int_{\Omega} d\mathbf{r}(\nu_T(\mathbf{r}) - \nu_R(\mathbf{r}))\tilde{\rho}(\mathbf{r}) \quad (1.6)$$

where the averaged density is given by:

$$\tilde{\rho}(\mathbf{r}) = \sum_{n=1}^{\infty} \frac{1}{n!} \frac{\partial^{n-1} \rho(\mathbf{r}, \lambda)}{\partial \lambda^{n-1}} \Big|_{\lambda=0} \quad (1.7)$$

In the work we report here, we follow this form of APDFT truncated up to the second order terms (Eqns. 1.6 and 1.7). The integral of the electron density is evaluated in VASP,[41] when we calculate E^R , which is equivalent to the average electrostatic potential centered around the nuclei.

While the derivation above is the most current and rigorous form of this perturbational approach, there have been several other diverse applications of alchemical perturbations for gradient-supported property calculations. In an early study on screening the stability of heteroatomic clusters with first order perturbation, Weigend and coworkers considered integer perturbations to the nuclear charge of atoms in homoatomic clusters of Pt or Ir using electrostatic potentials arising from the total electron density surrounding the perturbed atom.[42] Energy derivatives were applied to ionization potentials, excitation energies, and bond potentials by integrating the derivative over a path defined by λ , which in this case represented a parameter that provided a smooth transition between the initial and final energy states.[43] For instance, the occupation numbers of orbitals were used as λ to approximate HOMO ionization potentials and excitation energies. Sheppard and coworkers provided a detailed description of the Taylor expansion in chemical compound space centered around a reference compound at $\lambda = 0$, expression the potential energy of a target compound at $\lambda = 1$. [27] They demonstrated how the first energy derivative could be composed of contributions (written later in this section) that depend on changes in nuclear charge, atomic positions, and number of electrons that may occur in the transition between the reference and target compounds. Restricting themselves to isoelectronic alchemical transmutations with fixed atomic coordinates, they successfully approximated protonation reaction enthalpies, reaction barriers based on the umbrella flipping of ammonia, and oxygen binding to nanoparticles using electrostatic potential differences.

First order APDFT approximations have been used in many other notable applications including estimating the chemical potential of binary mixtures;[44] permitting broader screening through chemical space;[45] studies of mixed-metallic clusters;[46] calculating bond potentials;[47] screening of alkali halide crystals;[48] predicting material properties of bulk transition metals;[38] predicting molecular adsorption energies on metal surfaces;[26] probing non-locality of electron density;[49] predicting band structures in III-V semiconductors;[50] properties of BN-doped variants of benzene, coronene, fullerene, and graphene;[51, 52, 53][51, 52, 53] and deprotonation energies.[24]

We now provide a tutorial on how a simple form of APDFT can be employed for computational catalysis applications [51, 27, 26, 54]. To begin, we consider the BE calculation for an OH molecule on Pt(111). This system will be referred as our reference state and subsequently labeled with $\lambda = 0$. Since zero-point, thermal, and entropy corrections will be neglected out of simplicity, its BE will be labeled as $\Delta E^0|_{\lambda=0}$. We now will consider a new system where the Pt(111) surface has been doped with a new element that results in a new binding site. This new state will be labeled as $\lambda = 1$, and its BE will be labeled as $\Delta E^0|_{\lambda=1}$. APDFT can be used to predict $\Delta E^0|_{\lambda=1}$ by relating it to $\Delta E^0|_{\lambda=0}$ using a thermodynamic cycle shown in Figure 2.

According to the cycle, we have 1.8:

$$\Delta E^0|_{\lambda=0} + \Delta E_{\lambda \rightarrow 1}^a = \Delta E^0|_{\lambda=1} + \Delta E_{\lambda \rightarrow 1}^s \quad (1.8)$$

Where $\Delta E^0|_{\lambda=0}$ will be obtained using Eqn. 1.1, $\Delta E_{\lambda \rightarrow 1}^s$ is the energy change of the bare catalyst surface (s = site) when doped with a new element, and $\Delta E_{\lambda \rightarrow 1}^a$ is the same type of change for the system of the catalyst site with a bound adsorbate (a = ads-site). Upon rearrangement of Eqn. 1.8, the change in BE due to the transition from the $\lambda = 0$ to the $\lambda = 1$ states is the difference of these two unknown terms:

$$\Delta BE = \Delta E^0|_{\lambda=1} - \Delta E^0|_{\lambda=0} = \Delta E_{\lambda \rightarrow 1}^a - \Delta E_{\lambda \rightarrow 1}^s \quad (1.9)$$

We now will demonstrate how to calculate the right side of Eqn. 1.9 using APDFT to obtain ΔBE . APDFT relates $\Delta E^0|_{\lambda=1}$ to $\Delta E^0|_{\lambda=0}$ with perturbation theory, specifically by




Figure 2: Thermodynamic cycle depicting the binding energies (BEs) of an adsorbate on a surface (horizontal legs) and atomic transmutations (vertical legs).. $\Delta E^0|_{\lambda=0}$ and $\Delta E^0|_{\lambda=1}$ denote the BEs for the top and bottom horizontal legs, respectively. $\Delta E_{\lambda \rightarrow 1}^s$ and $\Delta E_{\lambda \rightarrow 1}^a$ denote the energy change associated with the atomic transmutation for the left (s = surface) and right (a = ads-site) vertical legs, respectively.

approximating the exact result as a Taylor series expansion with the thermodynamic state function (ΔE^0).[27] We have the following expression:

$$\Delta E^0|_{\lambda=1} = \Delta E^0|_{\lambda=0} + \partial_{\lambda} \Delta E^0 \Delta \lambda + \frac{1}{2} \partial_{\lambda}^2 \Delta E^0 \Delta \lambda^2 + \dots \quad (1.10)$$

Here, the predicted BE ($\Delta E^0|_{\lambda=1}$) is approximated as the reference BE ($\Delta E^0|_{\lambda=0}$) plus additional corrections from the Taylor series based on so-called “alchemical derivatives”, which are other terms resulting from the transformation from the $\lambda = 0$ state to the $\lambda = 1$ state. The n th alchemical derivative with respect to λ is denoted by $\partial_{\lambda}^n \Delta E^0$. For the easiest applications of APDFT, one can simply truncate the expansion in Eqn. 1.10 to first order and assign $\Delta \lambda = 1$, but higher order corrections for non-periodic computational schemes have been developed by von Lilienfeld and co-workers.[55] The first-order approximation for APDFT is shown with the following:

$$\partial_{\lambda} \Delta E^0 = \sum_I \Delta \mu_{nI} \partial_{\lambda} N_I - \sum_I \Delta F_I \partial_{\lambda} R_I + \Delta \mu_e \partial_{\lambda} N_e \quad (1.11)$$

Here, the first term accounts for the nuclear chemical potential gradient ($\Delta\mu_{nI}$) due to variation in nuclear charge (N_I) from the $\lambda = 0$ to the $\lambda = 1$ state. The second term accounts for energy gradients due to the forces on atoms (ΔF_I) resulting from changes in atomic positions (R_I) from the $\lambda = 0$ to the $\lambda = 1$ state. The third term accounts for the electronic chemical potential gradient ($\Delta\mu_e$) due to variation in total number of electrons (N_e) from the $\lambda = 0$ to the $\lambda = 1$ state.

Two additional approximations are typically employed with simple applications of first order APDFT used for computational catalysis studies of extended surfaces. The first approximation is to assume that nuclear positions for the $\lambda = 0$ state and the $\lambda = 1$ state are the same. While changing a material’s composition will certainly impact interatomic forces and result in relaxations to achieve a minimum energy state, it will be assumed for simplicity that the relaxation energy contributions due to the doping atom in the $\lambda = 1$ states for the “ads-site” and “site” calculations are similar and thus largely cancel in the thermodynamic cycle represented in Figure 2. Mathematically, the result is that $\delta_\lambda R_I$ in the second term of Eqn. 1.11 becomes zero, and thus the second term can be neglected. The second approximation is to ensure that the $\lambda = 0$ state and the $\lambda = 1$ state have the same number of electrons so that $\delta_\lambda N_e$ becomes zero so that the third term becomes zero. (Note that this is not a general requirement for APDFT, but it is a practical way to ensure that a total surface charge density under periodic boundary conditions remains reasonably physical in these specific applications.) To apply this constraint, one must use “isoelectronic transmutations” so that the $\lambda = 1$ system has the same number of electrons as the $\lambda = 0$ system. For example, in a Pt(111) slab, if one Pt atom is transmuted into an Au atom ($\Delta Z = +1$), a second Pt atom must be transmuted into an Ir atom ($\Delta Z = -1$). This results in a zero net change in atomic number for the new slab ($\sum_{n=1}^{\text{no. of atoms}} \Delta Z = 0$). While the second transmutation is a departure from the system of interest, it can be made to occur far from the site of interest so that it has no significant effect on the BE of that site.

To summarize, the simplest application of first order APDFT for predicting a BE for a hypothetical catalyst surface will consider only fixed atomic coordinates and only net isoelectronic transmutations, and this simplifies Eqn. 1.11 into the following:

$$\partial_\lambda \Delta E^0 = \sum_I \Delta \mu_{nI} \partial_\lambda N_I \quad (1.12)$$

The nuclear chemical potential, $\Delta \mu_{nI}$ (also called the alchemical potential), is defined as the change in atomic electrostatic potential between the “site” and the “ads-site” reference states. This is represented as an array of differences in electrostatic potentials, and so the elements of this array have units of energy/charge. Note that this procedure accounts for changes in electrostatic potentials of atoms on the catalyst surface after relaxations due to molecular adsorption have taken place. This term mathematically refers to the difference shown by the following expression:

$$\Delta \mu_{nI} = \mu_n^{\text{ads-site}}(R_I^{\text{ads-site}}) - \mu_n^{\text{site}}(R_I^{\text{site}}) \quad (1.13)$$

These atom-centered electrostatic potentials for the “ads-site” and “site” reference states can be obtained with VASP,[41] a widely-used DFT code. Figure 3a depicts how alchemical potentials are constructed as an array of differences in electrostatic potentials for atoms from both the “site” and “ads-site” states. The nuclear charge variation quantity, $\partial_\lambda N_I$, in Eqn. 1.12 now should be accounted for. When one transmutes atoms from a $\lambda = 0$ state to form a hypothetical $\lambda = 1$ state, one notes changes in nuclear charges. Specifically, one uses an array of differences in nuclear charges due to the formation of the $\lambda = 1$ state, and so the elements in this array have units of charge. Figure 3b depicts how nuclear charge variations are constructed as an array of differences in atomic numbers of the atoms in the $\lambda = 0$ state and the $\lambda = 1$ state. Note that this array will be the same whether it is constructed for a “site” or “ads-site” state since they both will correspond to the same $\lambda = 1$ state. The overall calculation of the alchemical derivative used to obtain ΔBE between the $\lambda = 0$ state and the $\lambda = 1$ state is simply a dot product of the two arrays given above, and this gives a scalar that has units of energy (Figure 3c).




Figure 3: Illustrations of steps to compute corrections to binding energies (BEs) based on APDFT. **a)** Constructing the alchemical potential array with differences between μ_i^a and μ_i^s , the electrostatic potentials for the i th atom in the ads-site (a) and site (s) systems, respectively. The adsorbate atoms in the ads-site state are conventionally not included in the array because they are not subject to transmutations. **b)** Constructing the nuclear charge variation array with differences in atomic number for each atom in the system before and after the transmutation. **c)** Illustrations of Eqs 1.12 and 1.17 as a dot product of the two arrays constructed in **a)** and **b)** (Note that the illustrations in **a)** and **b)** do not show an isoelectronic change and the result of the equality shown in **c)** is provided to show a simple visual representation.)

For additional clarity, we can return to the two unknown terms ($\Delta E_{\lambda \rightarrow 1}^a$ and $\Delta E_{\lambda \rightarrow 1}^s$) in Eqn. 1.9 that compute ΔBE according to the cycle in Figure 2. We now show the direct relation to the results of Eqns. 1.12 and 1.13 and the steps in Figure 3. The left leg of Figure 2 and the expression in Figure 3b both illustrate a reference Pt(111) surface “site” undergoing an alchemical transmutation of one Pt atom into an Au atom. Using Eqn. 1.12, a first-order approximation of the energy change is:

$$\Delta E_{\lambda \rightarrow 1}^s = \sum_I \mu_{nI}^{\text{site}}(R_I^{\text{site}}) \partial_\lambda N_I^{\text{site}} \quad (1.14)$$

Note that the illustrations in Figure 2 and Figure 3b do not show an isoelectronic change to provide a simple visual representation, but the arrays used in Eqn. 1.14 should account for them. The energy change for the same alchemical transmutation done to the “ads-site” species is:

$$\Delta E_{\lambda \rightarrow 1}^a = \sum_I \mu_{nI}^{\text{ads-site}}(R_I^{\text{ads-site}}) \partial_\lambda N_I^{\text{ads-site}} \quad (1.15)$$

For the cycle in Figure 2 to be true, the transmutations done to the “slab” and “ads-site” states must be exact, and this results in $\partial_\lambda N_I^{\text{site}} = \partial_\lambda N_I^{\text{ads-site}} = \partial_\lambda N_I$. Returning to Eqn. 1.9, and using Eqns. 1.12 through 1.15:

$$\Delta E_{\lambda \rightarrow 1}^a - \Delta E_{\lambda \rightarrow 1}^s = \sum_I [\mu_{nI}^{\text{ads-site}}(R_I^{\text{ads-site}}) - \mu_{nI}^{\text{site}}(R_I^{\text{site}})] \partial_\lambda N_I = \sum_I \Delta \mu_{nI} \partial_\lambda N_I \quad (1.16)$$

Finally, we combine Eqns. 1.9 and 1.16 to rewrite our simple approximation, illustrated in Figure 3c, for the predicted BE change using APDFT:

$$\Delta BE_{alc} = \sum_I \Delta \mu_{nI} \partial_\lambda N_I \quad (1.17)$$

Next, we present a benchmark case of APDFT predicted BE for OH adsorption on a four-layer, 2×2 Pt(111) surface model that contains four Pt atoms in each layer. For the binding configuration of OH on the surface, the fourth layer has two unique atoms that are transmutable. Thus, we created hypothetical materials by transmuting one of the eight atoms in the top two layers by $\Delta Z = \pm 1$ to convert a Pt near a binding site into Au or Ir and one of the two atoms in the fourth layer by $\Delta Z = \mp 1$. This generated 32 unique cases that computational alchemy can be used for predictions based on a single DFT binding energy calculation.




Figure 4: ΔBE predictions for OH adsorption on 32 hypothetical alloys of Pt made by $\Delta Z = \pm 1$ transmutations to the reference slab. Energy differences are relative to the reference calculation and reported in eV. The size of the data points corresponds to the distance of the transmutation site from the adsorbate (large points are transmutations far from the adsorption site).

Figure 4 shows a benchmarking comparison of first order APDFT calculations compared to corresponding single point DFT energy calculations. Note that these data were previously reported elsewhere,[26] but the data shown now depict the relative location of transmuted

atoms in the slab model to the OH adsorbate. Larger data points depict cases where transmuted atoms are further from the adsorbate while smaller data points depict cases where they are closer. These predictions have a mean absolute error of 0.045 eV. The first order APDFT is notably less accurate for transmutations made directly at the binding site, where the alchemical derivatives are the greatest.[26]

In summary, Figure 4 reiterates that first order APDFT can quite accurately predict how BEs changes on hypothetical alloys with simple algebraic computations based on a single set of QC calculations and corresponding electrostatic potentials. This provides a computational lever to increase the utility of a single BE calculation by enabling the estimation of (in this case) about 30 additional BEs. Even more BEs can be predicted by employing more and/or larger (e.g. $\Delta Z = \pm 2$) transmutations, but we have found these will result in larger errors that would need to be remedied for accuracy. Not only does this tool give users BE predictions, it also gives users physical insight into which nearby sites are the most important for adsorbate binding. In principle, this screening allows us to eliminate cases that result in an insignificant change in ΔBE , and thus allow more careful attention to cases that result in more significant changes in BEs using more refined QC methods and/or experimental studies.

2.0 Benchmarking Computational Alchemy for Carbide, Nitride, and Oxide Catalysts

The content of this section is taken from Charles Griego, Karthikeyan Saravanan, and John A. Keith, “Benchmarking Computational Alchemy for Carbide, Nitride, and Oxide Catalysts.” *Adv. Theory Simul.*, 2019, 2: 1800142.

Kohn-Sham density functional theory (DFT)-based searches for hypothetical catalysts are too computationally demanding for wide searches through diverse materials space. Here, we assess the accuracy of computational alchemy schemes on carbides, nitrides, and oxides. With a single set of reference DFT calculations, computational alchemy approximates adsorbate binding energies (BEs) on a large numbers of hypothetical catalysts surfaces with negligible computational cost. Analogous to previous studies on metal alloys, computational alchemy predicts adsorbate BEs on rocksalt TiC(111), TiN(100), and TiO(100) materials, which have no bandgap, in close agreement with DFT results (with mean unsigned errors up to 0.33 eV). In contrast, we find that semiconducting systems such as rutile TiO₂(110), rutile SnO₂(110), and rocksalt ZnO(100) can present more significant challenges. This work identifies these challenges being linked to the density of states at the Fermi level and by adding Pt dopants in the surface layer of TiO₂, we demonstrate that computational alchemy can become more reliable with semiconductor-based systems. This remedy provides insight that promotes computational alchemy for broad searches for catalyst active sites through materials space beyond transition metal alloys.

2.1 Introduction

We previously showed that computational alchemy schemes can predict BEs of oxygen reduction reaction intermediates on alloys of Pt, Pd, and Ni within 0.1 eV of DFT results.[26] Computational alchemy performed reasonably well for OH intermediates as well as for skin

alloys that are known to be problematic for the d-band model.[34] Alchemy errors are highest when alchemical derivatives are high. These are highest when transmutations occur nearest to the adsorption site, or when net isoelectronic transmutations span multiple groups of the periodic table, or when more and more atoms are transmuted.

It remains unclear if computational alchemy can be trusted for predicting binding energies on other material surface besides transition metals. Rocksalt TiC (111), TiN (100), TiO (100), as well as rutile TiO₂ (110) were considered due to these surfaces being of interest as catalysts and/or supports.[56, 57, 58, 59, 60, 61, 62, 63, 64] We considered pristine surfaces as well as those containing a main group element vacancy. The carbon vacancy was at the three-fold site in TiC, nitrogen and oxygen vacancies were under the four-fold site in TiN and TiO, respectively, and the oxygen vacancy in TiO₂ was a bridging oxygen. Following our previous work, we chose to model charge-neutral H atoms and OH molecules as adsorbed species (H* and OH*) since they are relevant oxygen reduction reaction intermediates. Both intermediates are bound on the three-fold site with a metal atom directly underneath in TiC and the four-fold site of TiN. For TiO, H* is also bound to the four-fold site, but OH* binds on top of a metal atom. For TiO₂, H* binds on top of a bridging oxygen atom, and OH* binds on top of a metal surface atom. These binding sites are favorable for both the pristine slabs and the slabs with a vacancy.

2.2 Computational Methodology

Kohn-Sham density functional theory (DFT) calculations were conducted using PBE[65] exchange correlation functional with projector augmented wave pseudopotential representations of core electrons as implemented in VASP[66, 41]. A plane wave basis set was used to represent valence electrons where 350 eV energy cutoff was chosen for rocksalt TiC, TiN, and TiO, 400 eV for rutile SnO₂, and 450 eV for rocksalt ZnO and rutile TiO₂. A $4 \times 4 \times 1$ Monkhorst-Pack sampling of the k-point grid was employed for all materials except SnO₂, which used $5 \times 5 \times 1$. The minimum energy configuration of each reference surface was determined iteratively with conjugate gradient algorithm until the difference in steps was

less than 0.1 meV. Each surface is modelled as a 2×2 slab extended to 4 layers, and bottom two layers are fixed, while the top layers and adsorbates are relaxed. We note that TiC, TiN, TiO, and ZnO unit cells contain four Ti atoms in each layer whereas TiO₂ and SnO₂ contain eight.

2.3 Results and Discussion

Figure 5 shows the performance of alchemy with low-energy facets of three different rock salt materials: TiC(111), TiN(100), and TiO(100) without and with a main group element vacancies. Consistent with our previous study on transition metal alloys, alchemy can reasonably calculate adsorbate BEs for H* and OH* intermediates when two, four, and six atoms have been transmuted. The reference binding energies from DFT in Figure 5 range from -1.18 eV to +6.11 eV for all systems. Calculated Δ BE values range from -0.53 eV to +0.75 eV for all systems. Deviations from the parity line correlate with systems having higher alchemical derivatives and/or systems where adsorption results in more significant polarization near the site where transmutations occur. In some cases, including when H* is found near a vacancy, the presence of a vacancy results in smaller alchemical derivatives and thus lower errors. Overall, alchemy can capably model H* and OH* binding energies on these three materials with errors within a range of 0.01-0.33 eV depending on the material, even when a vacancy is present. In all of these cases, there is a significant density of states at the Fermi level of all of these systems as would be expected for any transition metal alloy system.

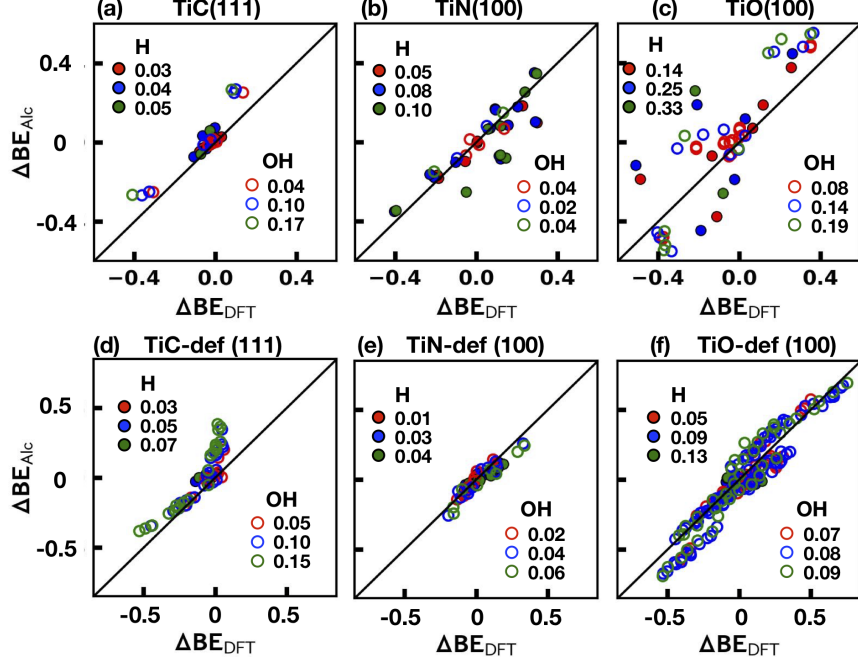


Figure 5: Parity plots comparing the accuracy of alchemy for predicting binding energy (BE) against DFT (in eV) for H* and OH* on pristine slabs (a-c) and slabs with the presence of a vacancy defect (“def”) (d-f): TiC, TiN, and TiO. ΔBE_{alc} is the alchemical derivative of BE between the hypothetical and reference slab. ΔBE_{DFT} is equal to $BE_{alloy} - BE_{ref}$ calculated from DFT. The hypothetical materials were obtained by alchemical transmutations to one (red), two (blue), and three (green) pair(s) of Ti atoms. Mean absolute errors (MAEs) are listed in each figure and given in eV.

We then evaluated computational alchemy when used on non-conducting materials (Figure 6). Interestingly, the performance of alchemy for approximating adsorbate BEs was rather poor even for just H*, and so other calculations involving OH* and vacancies were not pursued. For ZnO(100), there is a linear trend between alchemy and DFT, but the slope does not show parity. For SnO₂(110), there correlation is parallel to the parity line but consistently offset. The MAE for ZnO(100) is as high as the upper limit of errors seen in Figure 5, and MAEs were even greater for SnO₂ and TiO₂. As expected, the density of states

for these materials all show a bandgap. Alchemy exhibits very poor agreement with DFT in TiO_2 , but there is somewhat better agreement with DFT in SnO_2 and ZnO . We presently attribute this issue to the bandgaps in the materials. A significant density of states crossing the Fermi level signifies a conductive material that would also have a relatively high degree of shielding. High shielding makes atom-centered electrostatic potentials less sensitive to other nearby atoms, and this corresponds to the situations where first order approximations with alchemy would be most valid. This hypothesis will be validated in future work by testing second order alchemy approximations.

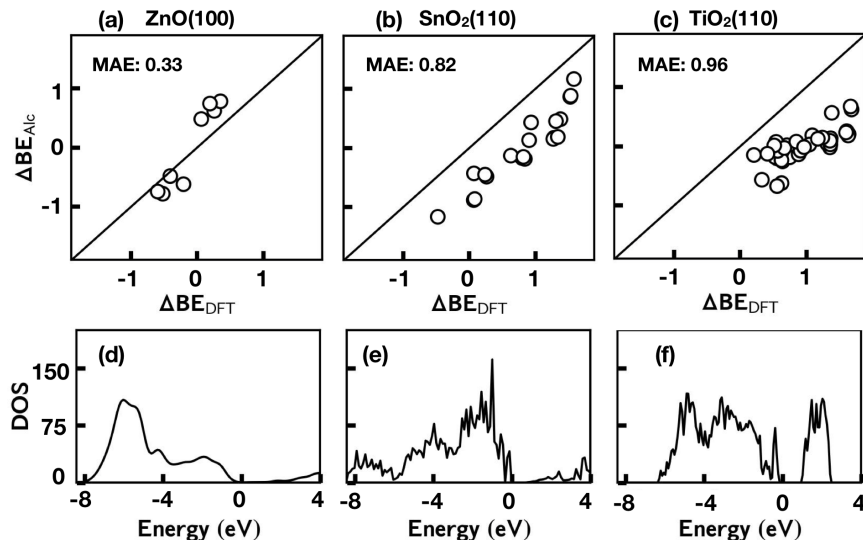


Figure 6: (a) Parity plots comparing the accuracy of alchemy for predicting binding energy (BE) against DFT (in eV) for H^* on pristine slabs of semiconducting materials: (a-c) ZnO , SnO_2 , and TiO_2 . $\Delta\text{BE}_{\text{alc}}$ is the alchemical derivative of BE between the hypothetical and reference slab. $\Delta\text{BE}_{\text{DFT}}$ is equal to $\text{BE}_{\text{alloy}} - \text{BE}_{\text{ref}}$ calculated from DFT. The hypothetical materials were obtained by alchemical transmutations to one pair of Zn, Sn, or Ti atoms. Mean absolute errors (MAEs) in each plot above are given in eV. (b) Density of states (DOS) for (d-f) ZnO , SnO_2 , and TiO_2 . The Fermi level is zero for each plot.

To further test that problems with alchemy approximations relate to bandgaps, we considered modified systems where 50% of the Ti atoms in the top layer or the second layer were replaced with Pt atoms. Placing Pt atoms in the TiC(111), TiN(100), or TiO(100) structures was expected not to significantly impact these systems since each was already electronically conducting. Indeed, alchemy in these systems with Pt atoms included led to similar errors as cases shown in Figure 5.

However, by placing Pt atoms in the second layer (Pt@2L) or top layer (Pt@1L) of TiO₂(110), we expected to tune the density of states near the Fermi level. We hypothesized that Pt atoms in the second layer do not directly interact with the vacuum layer, and so fewer states were expected to arise at the Fermi level, while Pt atoms in the first layer interact directly with vacuum, and so more states were expected to be found at the Fermi level. Figure 7 confirms this hypothesis and shows a dramatic improvement in alchemy predictions when there is a greater density of states at the Fermi level. Alchemy becomes significantly more reliable when Pt is in the first layer and predicts H* and OH* binding energies within the range of 0.16-0.29 eV.

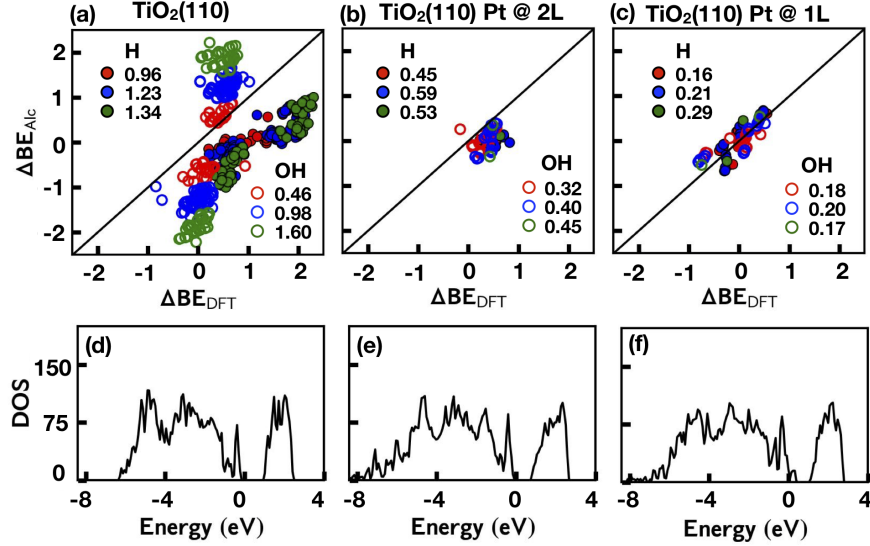


Figure 7: Parity plots comparing the accuracy of alchemy for predicting binding energy (BE) against DFT (in eV) for H* and OH* on (a-c) TiO₂ pristine and with 50% Pt-substituted in the second layer (Pt@2L) and first layer (Pt@1L). ΔBE_{alc} is the alchemical derivative of BE between the hypothetical and reference slab. ΔBE_{DFT} is equal to $BE_{alloy} - BE_{ref}$ as calculated from DFT. The hypothetical materials were obtained by alchemical transmutations of one (red), two (blue), and three (green) pair(s) of Ti atoms (pristine) or Pt and Ti atoms (Pt-substituted). Mean absolute errors (MAEs) in each plot above are given in eV. Density of states (DOS) for (d-f) TiO₂ pristine, Pt@2L, and Pt@1L. The Fermi level is zero for each plot.

2.4 Conclusion

In conclusion, we have demonstrated the utility of computational alchemy for approximating thermodynamic descriptors for catalysis for non-transition metal alloy systems. Computational alchemy performs reasonably well when predicting binding energies for H* and OH* on carbide, nitride, and oxide systems that have no band gap. Significant errors in this

computational scheme arise when used on systems that have a band gap. The present explanation for this is that first order approximations using alchemy are most valid for systems having high electronic shielding, which have been used to test alchemy in cases of transition metal alloys[26] and BN-doped graphene[51]. Results from alchemy can become more reliable by ensuring that the materials of interest have a substantial density of states at the Fermi level (which we did via in silico alloying with Pt in the top layer of the system). This shows a practical path forward to leverage computational alchemy for high-throughput searches of catalyst sites though broad regions of materials space.

3.0 Assessing APDFT Predicted Reaction Pathways and Barriers

The content of this section is taken from Charles Griego, John R. Kitchin, and John A. Keith, “Acceleration of catalyst discovery with easy, fast, and reproducible computational alchemy.” *International Journal of Quantum Chemistry.*, 2020, 121:e26380.

The expense of quantum chemistry calculations significantly hinders the search for novel catalysts. Here, we provide a tutorial for using an easy and highly cost-efficient calculation scheme called alchemical perturbation density functional theory (APDFT) for rapid predictions of binding energies of reaction intermediates and reaction barrier heights based on Kohn-Sham density functional theory reference data. Using a single nudged elastic band calculation for CH₄ dehydrogenation on Pt(111) as a reference case, we generate 32 new pathways with barrier heights having mean unsigned errors of less than 0.3 eV relative to single-point DFT calculations. Notably, this easy APDFT scheme brings no appreciable computational cost once reference calculations are done, and this shows that simple applications of computational alchemy can significantly impact DFT-driven explorations for catalysts.

3.1 Introduction

Limitations from costly QC methods become vastly present in computational studies of the kinetics of reactions on surfaces. The energy barrier (E_a) between two reaction steps, which largely dictates reaction kinetics, is a challenging quantity to measure with QC calculations since the transition state exists on a peak of the potential energy surface. Only the most relevant barriers ideally would be studied with costly QC methods, but determining whether a barrier is relevant requires iterative kinetic analyses that themselves require numerous barrier height predictions to be useful. A common approach for predicting E_a is to employ the nudged elastic band (NEB) algorithm[29] to a collection of interpolated

images between the initial and final states of a reaction step. Simultaneously, these images are optimized towards a minimum pathway on the potential energy surface, allowing one to identify the transition state. However, optimizing 10-20 images at once is a computationally expensive approach that hinders attempts to calculate E_a for all steps in a reaction network on multiple catalysts. In this study, we address this issue by uncovering many hypothetical reaction pathways and barrier heights from a single NEB with APDFT.

3.2 Results and Discussion

To benchmark first order APDFT predictions of reaction barriers, we use CH_4^* dehydrogenation on Pt(111) as a reference system, where each species bound to a surface site is denoted by $*$:



Figure 8 shows snapshots of the NEB reaction pathway for this process. The reactant state and first image in the NEB (a) contains CH_4^* geometrically relaxed to a distance of 3.7 Å from an ontop site. The reaction proceeds as one C-H bond breaks and a hydrogen is adsorbed to an adjacent ontop site. In the transition state, the detaching hydrogen is drawn toward a bridge site. The product state contains both a CH_3^* and an H^* bound at 2.1 and 1.5 Å at ontop sites, respectively. Other work reports that H^* also binds favorably to an fcc site in the product state.[67] However, our system is among many reasonable choices of reference points to predict changing trends of various possible reaction pathways on many hypothetical catalysts.

To compute barrier heights using first order APDFT, we apply the BE procedure from above on each of the images from the NEB calculation. From this, we can generate a variety of different pathways for up to 32 different hypothetical alloy configurations as shown in Figure 9. Figure 9a shows predicted reaction energy profiles following Eqn. 3.1. The reference energy profile calculated using DFT is denoted by red asterisks, while the solid blue line denotes the most affected Au alloys and the solid green line denotes the most




Figure 8: Snapshots of the (a) reactant, (b) transition, and (c) product states of the dehydrogenation of CH_4 adsorbed at $1/4$ coverage of ontop sites on Pt.

affected Ir alloys. The other lines pertain to all other cases where the alloy systems result in a negligible change relative to the reference system (due to very small alchemical derivatives present in these cases). Note that energy profiles for similar systems overlap and may not be distinguishable. Figure 9b shows a parity plot for ΔE_a values relative to the reference calculation. There is generally an increase in barrier when the transmuted site becomes Au and a decrease when the site becomes Ir. Notably, the largest errors in barrier heights are 0.3 eV relative to DFT benchmarks (for alloys that exhibited the highest alchemical derivatives), but the vast majority of other data are more accurate.

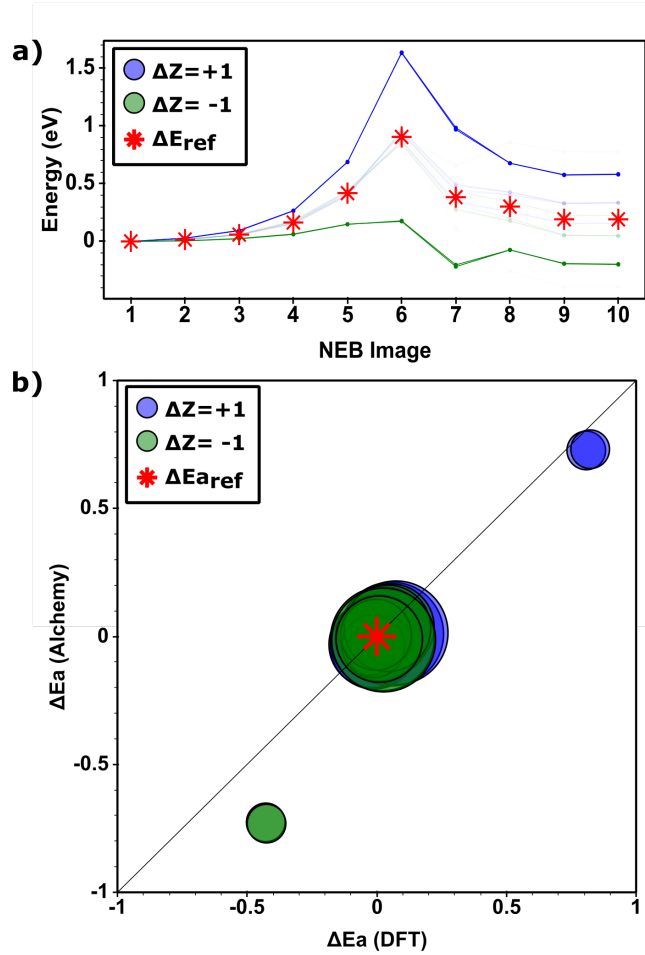


Figure 9: a) Energy profiles for the CH_4 dehydrogenation mechanism on hypothetical alloys of Pt. The reference pathway occurs on pure Pt and is denoted with red asterisks. The most significant effect from a transmutation with $\Delta Z = +1$ is shown in blue, the most significant effect from a transmutation with $\Delta Z = -1$ is shown in green, and other reaction pathways computed with alchemy are shown in light blue/green. b) First order APDFT benchmarking of the change in ΔE_a for CH_4 dehydrogenation on 32 hypothetical alloys of Pt made by $\Delta Z = \pm 1$ relative to the reference barrier (energies in eV). The size of the data points correspond to the distance of the transmutation site from the adsorbate in the reactant state.

3.3 Conclusion

APDFT is a promising gradient-supported method for leveraging QC calculations in large scale screening. This method relies on alchemical derivatives which are straight-forward to compute and which can be used for accurate predictions of binding energies and activation energies on electronically conductive systems relevant for computational catalysis. Using a single nudged elastic band calculation for CH_4 dehydrogenation on Pt(111) as a reference case, we generate 32 new pathways with barrier heights having mean unsigned errors of less than 0.3 eV relative to single-point DFT calculations. Even though first order APDFT can exhibit errors as large as 0.3 eV in barrier heights for reference system doped with just a single atom, it is promising that such simple approximations can be useful to calculate a computationally expensive descriptor that guides screening studies. Furthermore, as the source for errors in different APDFT approximations become better understood, it becomes reasonable to imagine that more accurate approaches might be developed based on APDFT that would have a predictive power comparable (or even indistinguishable) to standard DFT.

4.0 Machine Learning Corrected Alchemical Perturbation Density Functional Theory for Catalysis Applications

The content of this section is taken from Charles Griego, Lingyan Zhao, Karthikeyan Saravanan, and John A. Keith, “Machine learning corrected alchemical perturbation density functional theory for catalysis applications.” *AIChE J.* 2020; 66:e17041.

Alchemical perturbation density functional theory (APDFT) has promise for enabling computational screening of hypothetical catalyst sites. Here, we analyze errors in first order APDFT calculation schemes for binding energies of CH_x , NH_x , OH_x , and OOH adsorbates over a range of different coverages on hypothetical alloys based on a Pt(111) reference system. We then train three different support vector regression machine learning models that correct systematic APDFT prediction errors for each of the three classes of carbon, nitrogen, and oxygen based adsorbates. While uncorrected first order APDFT alone approximates accurate adsorbate binding energies on up to 36 hypothetical alloys based on a single Kohn-Sham DFT calculation on a 3×3 unit cell for Pt(111), the machine learning-corrected APDFT extends this number to more than 20,000 and provides a recipe for developing other machine learning-based APDFT models.

4.1 Introduction

Here, we investigated the feasibility of correcting errors in first order APDFT using machine learning (ML) approaches. We hypothesized that errors in adsorbate binding energies in these systems reflect a change from a linear energy relationship with alchemical derivatives to an unknown non-linear relationship. Thus, ML models that effectively use statistical techniques to learn non-linear relationships from training sets of data should be a productive way to correct these errors.

Indeed, ML has wide applications in biology, chemistry and materials science.[68, 69, 70, 71, 72] Notable examples range from small molecule chemistry to protein interactions to materials screening.[73, 74, 75, 76, 12, 77, 78, 79] ML models have also been applied to catalyst screening studies.[80] Combined ML and descriptor-based kinetic analysis framework was found to rapidly screen bimetallic catalysts for CO₂ reduction.[16, 15] Fully automated screening methods that combined ML and optimization techniques facilitated screening through a massive chemical space of inter-metallic catalysts.[6] Catalyst activity predictions were also accelerated coupling ML with active site configurations,[81] scaling relations,[14] and volcano plots.[82] We assessed the relationship between accuracy (relative to DFT) and the properties of systems evaluated with first order APDFT, we chose features that most correlate to APDFT’s accuracy, and we trained and tested a machine learning model on a large data set of adsorbate BEs on hypothetical alloys.

4.2 Computational Methodology

In this work, we benchmark BEs for different systems on a four-layer Pt(111) surface with the lower two layers fixed and the upper two layers relaxed. The adsorbates we chose are CH_x and NH_x where $x = 0 - 3$, as well as OH_x where $x = 0 - 2$ and OOH. Surface coverage (θ) is calculated as the number of adsorbates per atom in a surface layer. Here we model $\theta = 1, 1/4, 1/9$ by placing one adsorbate molecule on $1 \times 1, 2 \times 2$ and 3×3 slab models. We transmute atoms in the top two layers by $\pm\Delta Z$ and maintain isoelectronicity by transmuting an atom in the bottom layer by $\mp\Delta Z$.

All DFT calculations in this work were carried out using PBE[65] as the exchange-correlation functional and projector augmented wave pseudopotentials as implemented in VASP.[41] An energy cutoff of 350 eV and an $8 \times 8 \times 1$ Monkhorst-Pack grid sampling of k-point were used for all models. Structures were relaxed using conjugate gradient iterative minimization until the difference in energies between consecutive geometry steps was less than 1 meV. Adsorbate binding energies were calculated using the equation $BE = E_{\text{site}} + E_{\text{ads}} - E_{\text{ads-site}}$. Where E_{site} is the energy of the bare catalyst surface, E_{ads} the energy of the catalyst with an adsorbate bound to the surface site, and $E_{\text{ads-site}}$ is the energy of the lone molecular adsorbate.

Once a reference BE was obtained from DFT calculations, we constructed a list of atom-centered electrostatic potential differences. Then, we constructed a second list that describes the ΔZ change to each atom in the slab model after we made transmutations. We approximated ΔBE_{APDFT} by taking the dot product of the two lists and benchmarking it to ΔBE_{DFT} values obtained from single point DFT calculations (again, relaxed with respect to adsorption of the reference system but not with respect to alchemical transmutations). Detailed descriptions of this method are described in our previous work.[83] For each reference system, we evaluated alloys made from $N_{\text{T}} = 1 - 4$ and with $\Delta Z = \pm 1, 2, 3$. All possible N_{T} changes were made among all sites in the first layer and the second layer for the 1×1 and 2×2 surface models. For the 3×3 model, we chose a random subset of about 150 configurations to reduce the number of DFT calculations needed to benchmark since 24,282 configurations are combinatorially possible. The data production was facilitated using Phystone, an APDFT Python library developed by our team found at https://github.com/KeithLabPitt/comp_alchemy.

We first assessed the errors with APDFT predicted BE with respect to different characteristics of the alloys and the adsorbate binding. To analyze these variables individually, we have bucketed the APDFT BE errors from our entire dataset into categories. In Figure 10, we show bar graphs separated by coverage. Individual bars show the mean absolute error (MAE) with respect to both ΔZ and N_{T} . For $\theta = 1/9$, the MAE for all cases is roughly below 0.6 eV and generally increases when ΔZ or N_{T} increase. For $\theta = 1/4$, the MAE for each case is below 1.6 eV and systematically increases with ΔZ and N_{T} . For $\theta = 1, N_{\text{T}} = 1$

or 2 because there are only two sites available in the first and second layers. We found that the MAE is below 2 eV and increases with both variables. Overall, MAE increased when the coverage increased. This is because with higher coverage, the concentration of sites with large alchemical derivatives, that overestimate BE, were greater.




Figure 10: Mean absolute error (MAE) with respect to change in nuclear charge at each transmuted site (ΔZ) and the number of transmuted atoms (N_T) for binding energy calculations of all adsorbates at coverages ($\theta = 1/9, 1/4, 1$).

In Figure 11, we compare the MAE with respect to the type of adsorbate in the reference system. Here, the adsorbate type is described by the amount of central atom hydrogenation (ie. the value of x in CH_x , NH_x , or OH_x (or OOH)). There are two notable observations from this analysis. The APDFT predicted MAE decreases with the degree of increasing hydrogenation and when the central atom in the adsorbate is farther to the right of the

periodic table. In Figures 10 and 11, we see strong dependence of MAE with ΔZ , N_T , θ , and the type of adsorbate, suggesting that these variables are viable features for making predictions.




Figure 11: MAE with respect to hydrogenation of the central atom for CH_x, NH_x, and OH_x adsorbate BE.

4.3 ML Model

Here we developed ML models that can predict the magnitude of errors of APDFT relative to DFT accuracy. Figure 12 shows the workflow we followed to obtain ML corrected APDFT BE predictions. With a hypothetical alloy surface as the input, we fingerprinted the hypothetical dopant configuration that resulted from transmutations, and we constructed a feature vector from these fingerprints and other descriptors unique to the adsorbate-alloy system. The feature vector was fed to the model and outputted a predicted value of the APDFT’s deviation from DFT. The error was then added to the original APDFT predictions as a correction.

In order for a machine learning model to correctly learn the underlying patterns in the training data, the most relevant features must be identified and processed. These features are variables that describe particular properties among all groups in a large dataset. We included in our model the variables that strongly correlate with BE predictions errors as identified

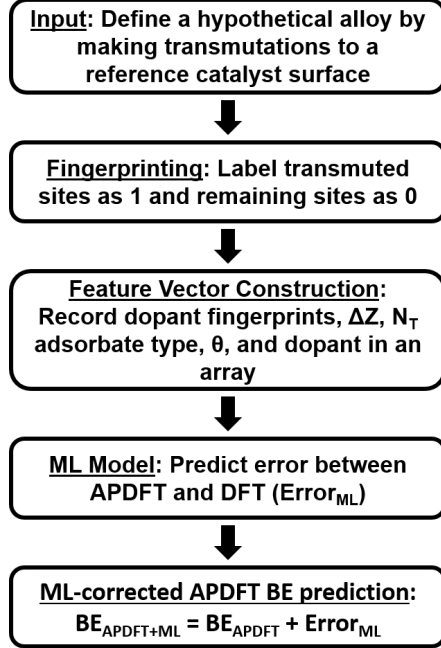


Figure 12: Machine learning workflow for obtaining corrected alchemy BE predictions.

in Figures 10 and 11: ΔZ , N_T , θ , and the degree of hydrogenation of the central atom in the adsorbate. While these variables give a simple description of the kind of adsorption and the type of changes made in the reference system, a feature representation of the dopant location in the alloys would provide a much more unique description of the alloy to link to APDFT prediction accuracy. This was done by considering a feature vector with length equal to the number of atoms in the top two layers of our slab models. Each element of the vector represents one of the sites in the top two layers where a binary indication is used to describe the site being transmuted or not. For more details, our feature vector construction is explained in the Appendix. Once we construct the feature vector for the dopant configuration, we add additional elements with binary descriptions of the selected features.

The dataset we use to train our model includes all surface alloys outlined in the error analysis above. For each of these alloys, a feature vector was constructed as outlined above, and the error in BE between APDFT and DFT was calculated. To most accurately assess ML models, we partition the data into a training set and a test set and make validations using k-fold cross validation (CV) (more information in the Appendix).

We considered several machine learning algorithms to predict APDFT error corrections. We evaluated support vector regression (SVR) models with radial basis function (RBF) and polynomial kernels, random forest, linear regression, ridge regression, and elastic net regression by measuring the performance with a 10-fold CV. The CV MAE for each model trained using the CH_x , NH_x , and OH_x subsets is listed in Table 1. The SVR model with the RBF kernel gave the best prediction accuracy, and the hyperparameters were optimized also with a 10-fold CV. We found that the default parameters provided through the scikit-learn library[84] gave the best performing version of this model. With the optimized hyperparameters, the model is retrained for each subset. When building the final models, we split the data of each subset into 80% training and 20% testing partitions. The assessment of this split, the learning performance, and how well the model is generalized is show in Figure 13. We have plotted the MAE of model predictions of the error between APDFT and DFT calculated BEs ($\Delta\text{BE}_{\text{APDFT}} - \Delta\text{BE}_{\text{DFT}}$) using both the training (red circles, solid line) and test (black circle, solid line) set against the training sample size. With increasing training samples (BEs), the training score curve converges and becomes stable at a minimum MAE. The test score curve also exhibits this behavior and quickly approaches the training score curve with increasing training samples.

Table 1: Summary of machine learning model evaluation and selection based on the CV MAEs for each dataset.

Model Type	MAE CH _x (eV)	MAE NH _x (eV)	MAE OH _x (eV)
SVR (RBF)	0.07±0.006	0.08±0.006	0.07±0.004
SVR (Polynomial)	0.14±0.008	0.21±0.006	0.21±0.01
Random Forest	0.11±0.008	0.13±0.01	0.10±0.006
Linear	0.31±0.01	0.37±0.02	0.35±0.01
Ridge	0.31±0.01	0.37±0.02	0.35±0.01
Elastic Net	0.43±0.02	0.44±0.02	0.36±0.01

Figure 13 also shows learning curves of ML model predictions of DFT BEs ($\Delta\text{BE}_{\text{DFT}}$) using only DFT training data, where MAE of predictions on both the training (red stars, dashed line) and test (black stars, dashed line) set are plotted against the training sample size. For the OH_x and CH_x datasets, the test score curves for the ML model trained to APDFT data are consistently lower than those for the ML model using strictly DFT data and show that more accurate ML models can be trained with APDFT data. For the NH_x dataset, the test score curves are more similar, and with some training set sizes there is not an apparent benefit to using APDFT data. Developing a physically grounded explanation for this as well as how learning curves change when using higher order corrections of APDFT will be addressed in future work.

4.4 Results and Discussion

As shown in the final steps of Figure 12, our models output predicted errors that are added to the APDFT calculated BE as a correction. Three separate SVR models were trained using datasets containing CH_x, NH_x, OH_x and OOH species. Figure 14 shows pure APDFT

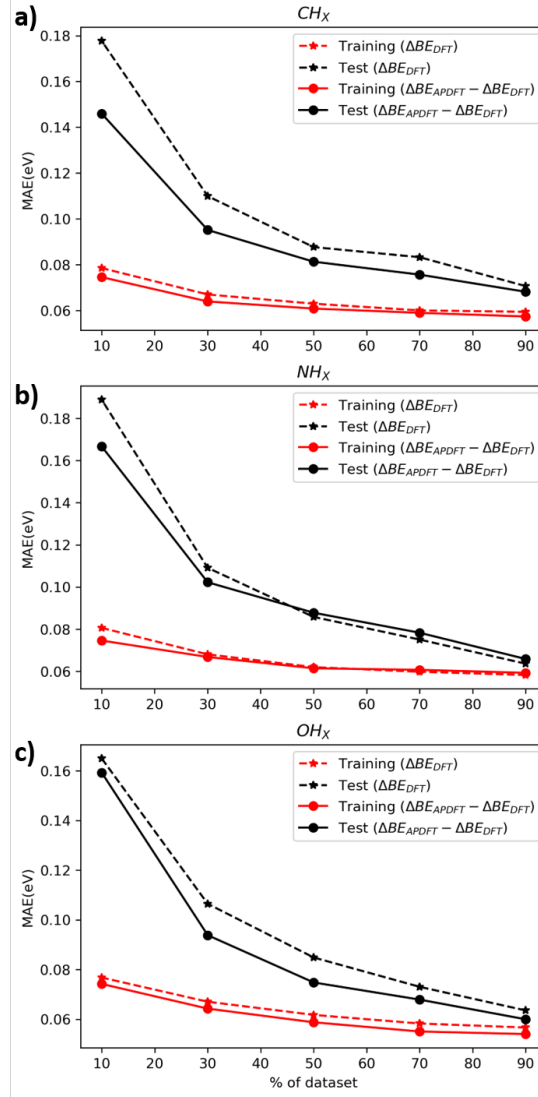


Figure 13: Comparison of learning curves for ML models that predict DFT BEs using a training set of DFT calculated BEs (labelled ΔBE_{DFT} in the legends) or a training set that learns deviations between APDFT and DFT (labelled $\Delta BE_{APDFT} - \Delta BE_{DFT}$ in the legends).

(a) and ML-corrected APDFT (b) predictions for alloys in the CH_x test set against respective DFT predictions of . Similar plots for NH_x and OH_x adsorbate test sets are included in the Appendix. Upon ML corrections, the prediction accuracy notably increases for each case,

and with many improved cases where APDFT systematically under and over-predicted the BE. For CH_x , MAE decreased by more than an order of magnitude from 0.78 to 0.07 eV. MAE decreases from 0.65 to 0.07 eV for NH_x , and for OH_x and OOH , the MAE decreased from 0.39 to 0.06 eV.




Figure 14: Parity plot of the BE change (ΔBE) of CH_x adsorbates on alloys of transmuted Pt predicted with APDFT (a) and ML-corrected APDFT (b) compared to DFT. ΔBE are in units of eV. The MAE decreases from 0.78 (a) to 0.07 eV (b) upon ML-correction.

We then investigated which features were most important in the model training. Each feature was scored with a normalized negative log of the p-value. These values are recorded in the bar chart in Figure 15. A feature is most important if the value is close to 1 and least important when it is close to 0. We initially found it curious that the transmutation site is the least important feature because our previous work reports that APDFT's largest prediction errors belong to systems where the transmutation is at the binding location. However, this result is sensible because there are up to 18 possible transmutations sites in our models, and

only a relatively small number of these sites are close to the adsorbate. This means that it is less common to see in the dataset cases where large errors come from changing sites near the adsorbate. Conversely, this argument explains why ΔZ is the most important feature. It was also reported in our previous work[26, 54] (and shown in Figure 10) that for a larger nuclear charge perturbation, the errors are larger, regardless of the location of the change. Thus, for the part of the dataset where the change is large (ie. $\Delta Z = 3$), almost all cases have large errors, and this feature plays a greater role in the model training.




Figure 15: Feature importance plot showing the significance of different parameters for the machine learning model. Features are scored based on a normalized negative log of the p -value, where the most important feature (ΔZ) ranks at a score of 1.0.

Since ΔZ and N_T were the top two most important features, we outlined how our models were especially effective in improving predictions for alloys made from large ΔZ and N_T transmutations, where first order APDFT fails. Figure 16 illustrates the distribution and level of improvement of predictions with alloys in the test data set relative to ΔZ and N_T . The radius of each circle corresponds to the percentage of alloys present in the test data set made from that particular pair of ΔZ and N_T . Each circle is colored based on the difference in MAEs for the uncorrected APDFT and ML-corrected APDFT in that set of alloys. For example, dark blue represents that ML-corrections effectively resulted in zero improvement while dark red represents substantial improvement by more than 1 eV. Note that alloys with larger N_T had more possible combinations than smaller N_T and thus comprised greater

percentages of the data set. With increasing ΔZ , the relative percentages of different alloys that were tested does not change, but the magnitude of the corrected MAE significantly increases. This shows where ML corrections are most impactful while also showing which subsets of alloys that were tested are most at risk for being underfit. For advanced ML models that interface with APDFT, it is recommended to use of datasets that account for intrinsic symmetries of each adsorbate configuration as well as to focus on allocating even more training data on hypothetical alloys using the larger $|\Delta Z|$ and N_T values.




Figure 16: Visualization of the distribution of hypothetical alloy types in the test data set and how well ML improves APDFT predictions for those alloy types. N_T is the quantity of transmutations used to generate a hypothetical alloy and is dimensionless. ΔZ is the size of the transmutation step in generating an alloy and has units of charge. Circle radii reflect the percentage of alloys from the test data set that are made up of that particular combination of N_T and $|\Delta Z|$. The filled color represents the change in MAE (in units of eV) to that subset of data after ML correction have been applied.

4.5 Conclusion

We have demonstrated how an ML-based treatment can significantly improve the accuracy of first order APDFT BE predictions of a range of adsorbates on hypothetical alloys of Pt(111). First, we identify the source of errors related to the system variables. Generally, the

MAE among systems in our dataset increases with increasing ΔZ , N_T , and coverage. There is also a correlation to the type of adsorbate in the reference system as the MAE is found to decrease with increasing levels of hydrogenation at the central atom. We used this insight to fingerprint our alloy systems and construct feature vectors that include these variables. We found through a k-fold CV analysis that an SVR model with a radial basis function kernel gave the most accurate prediction of the APDFT errors. Training separate SVR models on CH_x , NH_x , and OH_x datasets, BE predictions for systems in our test set had accuracy improved by 0.71, 0.58, and 0.33 eV, respectively. Each model used at least 3,600 BE data points for training and were validated using at least 900 data points. While an appreciable amount of DFT calculations were needed to develop the machine learning models, their necessity extends beyond the numbers of BE predictions in our test sets. Our models could provide reliable BE predictions on alloy configurations that weren't included in this study. For instance, in the section of our dataset where $\theta = 1/9$ (using a 3×3 surface model) we included a rather small subset of all possible combinations involving $N_T = 1, 2, 3, 4$ and $\Delta Z = \pm 1, \pm 2, \pm 3$ transmutations among the 18 atoms in the top two layers. In this case, there are $2 \cdot (3 \cdot \binom{18}{1} + 3 \cdot \binom{18}{2} + 3 \cdot \binom{18}{3} + 3 \cdot \binom{18}{4}) = 24,282$ hypothetical alloy configurations that could be modeled reasonably accurately given a single reference BE calculation and a trained ML model. However, only about 1,600 systems with $\theta = 1/9$ were used for training all models, and about 400 were included in validating all models. Upon investigating the feature importance, we found that the ΔZ and N_T variables were most important, while the choice in the transmutation site was least important in the model training. This is apparent because choices in ΔZ and N_T that are higher in magnitude lead to more systems in the dataset with low prediction accuracy. These findings suggest that ML models provide a straightforward and practical remedy for first order APDFT predictions of BEs, and that ML models can in principle achieve the same accuracy as DFT calculations. Future work will still focus on developing schemes to more correct these errors more physically with higher order terms as well as ensure that ML training sets are not underfit. This is expected to allow more accurate predictions of catalysis through materials space.

5.0 Second Order APDFT: Methods and Applications for Heterogeneous Catalysis

Expansive searches of hypothetical catalysts based in Kohn-Sham density functional theory (DFT) can become more feasible by extrapolating properties from existing data. Alchemical perturbation density functional theory (APDFT) is a physically-grounded approach towards extrapolating catalyst properties via a Taylor expansion of alchemical energy derivatives. Truncating the Taylor series to first derivatives (APDFT-1) has provided adequate adsorbate binding energy predictions, but several shortcomings have been identified that are correlated to characteristics of the catalyst system under scrutiny. Here, we assess second derivatives in APDFT (APDFT-2) to make improved binding energy predictions using a collection of reference DFT data obtained from VASP and CP2K. Hydrogen adsorption was evaluated with APDFT-2 on hypothetical variations of Pt and TiO_2 surfaces that are formed by alchemical transmutations. Results are compared to DFT data by illustrating the alchemical energy function over theoretical transmutation pathways that consist of states with non-integer nuclear charges. We find that with CP2K, APDFT-2 is more feasible, and more accurate than APDFT-1, by conducting finite differences of derivatives evaluated at states with effective, non-integer nuclear charge perturbations.

5.1 Introduction

The field of computational catalysis has made great efforts to produce efficient and accurate screening methodologies that help identify stable and abundant materials that exhibit promising activity. Quantum mechanics calculations (QM) typically provide the foundation for these methods, with Kohn-Sham Density Functional Theory (DFT) schemes that allow an adequate balance of accuracy and computational expense. A straight-forward approach to screen materials is to use brute force, and evaluate the properties for all species in a defined search space. However, while the expense of DFT is reliable for many systems individually, this cost will certainly add up for large search spaces, and it simply would not be feasible, or desirable, to search all of materials space in this way.[1]

Binding energy linear scaling relations[85] and Bronsted-Evans-Polanyi[86] relations were a very useful discovery for catalyst design, where one could use linear relations between binding energies of similar adsorbates or the activation energy of a surface reaction versus its reaction energy. Additionally, the *d*-band model that has become widely accepted to approximate changes in adsorbate binding energy between shifts in catalyst *d*-band centers.[33] All these methods have helped alleviate the conceptual effort to measure how an entire reaction series may change with catalyst species. As this became an advantage, it also revealed yet another challenge in catalyst design.

Scaling relations create a limit to the catalyst activity, because this phenomenon causes binding energies of all intermediates in a reaction to shift equally for different species of active sites, and the activation energy barriers remain unchanged. This requires ways to circumvent scaling relations, by tuning binding energies of intermediates independently of one another. This can be done by alloying a surface with multiple species to create an environment where intermediates favorably bind at respective sites, but the search space opens immensely with countless combinations of atom species one could investigate given a single template of a surface facet. Many alternative approaches have been developed to expedite the process of evaluating reaction energies on alloy materials such as coordination number models,[11, 12, 13] cluster expansions,[87, 88, 89, 90, 91] and machine learning models.[92, 93, 94] However, some of these approaches require a large number of initial DFT calculations to train and validate models, and while there are strengths predicting properties interpolated within the training space, extrapolating quantities is still an issue.

Gradient-supported approaches may have more potential to extrapolate properties along farther lengths of materials space. Such approaches rely on derivatives of properties, such as energy, respective to changes in an independent variable. Nuclear charge is a practical variable in chemical/material search spaces, as a change in nuclear charge directly relates to a compositional change in a compound. Quantum alchemy methods employ this relationship between properties and nuclear change with perturbation theory, where a transmutation of an atom is considered as a perturbation to the nuclear charge.

There are a handful of quantum alchemy methods that consider energy or electron density derivatives with respect to varying nuclear charge. These include using a Taylor series expansion of the energy,[39] where derivatives can be determined with finite differences of QM energies,[95] derivatives being computed from conceptual density functional theory,[96, 97, 98, 99] density derivatives based in Alchemical Perturbation Density Functional Theory,[23] and more recently, derivatives obtained via automatic differentiation.[100, 101] Quantum alchemy methods should be considered separate from classical alchemy approaches, which are thermodynamic integration schemes in molecular dynamics, where free energy differences are obtained by smoothly varying force field parameters between two potential energy surfaces which may result by removing/placing nuclear charge gradually at fractional increments.[25]

von Lilienfeld, von Rudorff, and coworkers have disseminated a thorough understanding of quantum alchemy methods for atoms and molecules. Recent work from this group has covered atomic contributions in molecules,[55] deprotonation energies in small organic molecules,[24] effects from basis set sizes on alchemical predictions,[102] alchemical chirality in chemical compound space,[103] and convergence of Taylor series expansions.[104] Our group has adopted these techniques of using a Taylor series expansion of the energy into the realm of materials screening with applications in catalysis,[26, 54, 83, 105] where we use the APDFT formulation with density derivatives, as well as energy derivatives obtained from finite differences of QM energies.[95]

In catalysis studies we identified that first-order APDFT extrapolations from VASP data were quite promising on properties evaluated in thermodynamic cycles. In the thermodynamic cycle scheme, the change in energy of a process, such as an adsorbate binding to a surface, is evaluated by calculating the energy change from an alchemical transmutation in both the reactant and product states. The resulting quantity is an energy difference between two states, which induces error cancellations and allows accurate APDFT predictions.[83, 95] Although the thermodynamic cycle scheme proved reliable in predicting catalyst properties, there are still several notable shortcomings for certain systems. First, we observed that accuracy decreases when a site is transmuted near an adsorbate, which is especially problematic in catalyst design, as we are usually more interested in the influence the species of the active site has on adsorption. And although we found that APDFT was reliable for metallic systems such as Pt, Pd, Ni, Pt/Pd-Ni skin alloys,[26] accuracy significantly decreases for semiconductors like ZnO, SnO₂, and TiO₂. [54]

Before now, we had only identified that these correlations between APDFT accuracy and the choices in reference/target systems existed, and we had demonstrated that these correlations were strong enough to reliably train machine learning models.[105] However, we have yet to identify the underlying reasons behind these correlations nor have we explored applications of higher ordered APDFT approximations, which require accurate evaluations of non-integer nuclear charge states with DFT. The aim of this work was to gain a deeper understanding of first and second order APDFT applications in catalysis by investigating alchemical potential energy functions for systems previously studied. By doing so, we could

accomplish two things: (1) assess ways to evaluate non-integer nuclear charge states, that allow higher ordered APDFT predictions, with common DFT software packages, and (2) visualize the alchemical energy functions that APDFT approximates. We hypothesized that the alchemical potential energy surfaces for cases that exhibited reliable first-order APDFT predictions would be linear functions, and for cases with low accuracy first-order APDFT predictions, the alchemical curves would resemble parabolic or higher ordered polynomials that second order APDFT may effectively treat. This work also serves as a guide for computational materials scientists that wish to pursue APDFT screening studies with data from popular planewave DFT codes.

5.2 Taylor Series Approximation with Alchemical Derivatives

In APDFT, the Hamiltonians (\hat{H}) of the reference and target material are linearly coupled through a mixing parameter, λ , as the following:

$$\hat{H}(\lambda) = \lambda\hat{H}_T + (1 - \lambda)\hat{H}_R \quad (5.1)$$

where $0 \leq \lambda \leq 1$ and describes the size of the perturbation. The desired quantity resulting from this coupled Hamiltonian is the energy of the target system at $\lambda = 1$ (E^T), which can be expanded in a Taylor series around the energy of the reference system at $\lambda = 0$:

$$E^T = \sum_{n=0}^{\infty} \frac{1}{n!} \frac{\partial^n E(\lambda)}{\partial \lambda^n} \Big|_{\lambda=0} = E^R + \sum_{n=1}^{\infty} \frac{1}{n!} \frac{\partial^n E(\lambda)}{\partial \lambda^n} \Big|_{\lambda=0} \quad (5.2)$$

Here, the Taylor series approximates the curvature of the alchemical potential energy function using local derivatives around a reference point on the curve. For a linear alchemical function, a first-order Taylor series approximation that only includes one term in the sum in Equation 5.2 (APDFT-1) may sufficiently replicate the curve. However, the alchemical function may likely be parabolic or a higher-ordered polynomial, and more terms will need to be included in the sum to approximate the curvature of this function (APDFT-N).

Assuming that the Taylor series converges, there will eventually be a limit to the number of terms that achieve maximum accuracy of the approximation, as the magnitudes of higher-ordered derivatives will become insignificant. von Rudorff identified the convergence radius of the Taylor series expansion for dimers, recovering self-consistent energies of Hartree Fock calculations.[104] Achieving accurate derivatives at higher orders has presented challenges due to quantities in QM codes being evaluated at double precision (15 significant digits). The Taylor series requires larger precision for higher orders due to the errors of the derivatives becoming amplified as they are divided by the $n!$ factor that grows quickly with subsequent inclusion of terms. Due to this issue, we experienced gross over-estimates of quantities when including third and fourth order terms, and thus this body of work only includes approximations up to second order.

Next we will elaborate on how to calculate first and second derivatives for the Taylor's series from both VASP and CP2K calculations. The first derivative from Eqn. 5.2 above can be evaluated via the Hellmann-Feynman theorem:[40]

$$\left. \frac{\partial E}{\partial \lambda} \right|_{\lambda=0} = \langle \psi_R | \hat{H}_T - \hat{H}_R | \psi_R \rangle = \int_{\Omega} d\mathbf{r} (\nu_T(\mathbf{r}) - \nu_R(\mathbf{r})) \rho(\mathbf{r}, \lambda) \quad (5.3)$$

which include the external coulomb potentials ν_R and ν_T corresponding to reference and target systems, respectively. This term is calculated differently between VASP and CP2K, but the same general assumptions apply for evaluating systems with APDFT: (1) the atomic coordinates remain fixed after all transmutions and (2) the transmutions are done in pairs isoelectronically, where one atom experiences a $+\Delta Z$ nuclear charge change, and the other experiences a $-\Delta Z$ change.

In other papers,[26, 83] we have covered extensively first alchemical derivatives from VASP calculations and presented the formula:

$$\left. \frac{\partial E}{\partial \lambda} \right|_{\lambda=0} = \sum_I \mu_{nI} \left. \frac{\partial N_I}{\partial \lambda} \right|_{\lambda=0} = \sum_I \mu_{nI} \Delta Z_I \quad (5.4)$$

which results in a dot product between atom-centered electrostatic potentials (μ_{nI}) and the integer nuclear charge difference over λ (ΔZ_I) at atom I .

Unfortunately, there is currently not a straight-forward method to obtain atom-centered electrostatic potentials in CP2K. However, we can conduct a procedure that follows the Hellmann-Feynman Theorem, which states that the first derivative of energy with respect to the parameter λ is equal to the difference between energies obtained by evaluating both \hat{H}_T and \hat{H}_R with the reference wavefunction, ψ_R . ψ_R is easily obtainable by collecting the restart files from a converged calculation that optimized the electronic structure of the reference system. With this file, we can evaluate the energies from \hat{H}_T and \hat{H}_R by “restarting” calculations for both systems, where each have the tag *SCF_GUESS* set to *RESTART* and the restart file is copied to the respective directories. Since we are only interested in the energies that result from evaluating each Hamiltonian with ψ_R and aren’t intending to re-optimize the electronic structure, the maximum number of self-consistent field iterations (*MAX_SCF*) should be set to 1. Finally, the difference between both resulting energy values equals the first derivative needed for APDFT-1.

To conduct second order APDFT approximations (APDFT-2), we evaluate second order alchemical energy derivatives with central finite differences. By evaluating first order (N order) derivatives along small displacements from a reference point on the alchemical potential energy curve, we can use finite differences to compute second order ($N + 1$ order) derivatives with an approximation of the formal definition of a derivative:

$$\frac{\partial f}{\partial \lambda} = \lim_{\lambda \rightarrow 0} \frac{f(\lambda + h) - f(\lambda - h)}{2h} \approx \frac{f(\lambda + h) - f(\lambda - h)}{2h} \quad (5.5)$$

where h represents the width of the finite displacements from the reference point ($\lambda = 0$). In order to obtain most accurate results, we used $h = 0.01$ for this work.

5.3 Modelling Non-integer Nuclear Charges

In order to perform APDFT-N predictions for periodic systems in VASP and CP2K, energy values and first derivatives need to be evaluated at fictitious states that result from fractional perturbations to the nuclear charge of atoms in a reference system. This section serves as a tutorial on how to evaluate such types of systems in VASP and CP2K.

5.3.1 VASP

For calculations in VASP, fictitious states are modelled using the virtual crystal approximation method (*VCA* tag in *INCAR* file), where weights are assigned to two overlapping atoms. The two overlapping atoms include an atom that already exists in the reference calculation and a desired target atom that results from the transmutation. The sum of the weights must equal 1, and the weight assigned to the target atom equals λ , the mixing parameter in the perturbation. For example, for $\lambda = 0.1$, 0.9 must be assigned to the reference atom (90%) and 0.1 to the target atom (10%). These weights are used to scale values in the *POTCAR* (pseudopotential) files supplied for each atom, which is an adaption of the scheme from Bellaiche and Vanderbilt.[106] We must note that the VASP manual recommends to use this method cautiously, as the radial cutoffs in the PAW potentials need to be adjusted to accurately model the psuedo-plane waves of the virtual species. This issue is further discussed later in this writing, regarding the influence of this method on APDFT predictions.

5.3.2 CP2K

For CP2K calculations, we can model fictitious states by assigning a non-integer effective nuclear charge using the *CORE CORRECTION* tag in the *KIND* section of the input file, where the details on each atom are supplied. This approach isn't straightforward, however, because this tag will apply the core correction to every atom of the same kind in the structure. An easy work around is to supply separate, distinct *KIND* sections for each transmuted atom. This can be done by giving the atom a new symbol in the list of *XYZ* coordinates supplied in the *xyz* file or *COORD* input section, with the requirement that the new symbol must match the label for the new *KIND* section. For an example, if a reference system is a surface model of fcc Pt (111) with one surface atom being transmuted, there should be a regular *KIND* section for Pt and an additional *KIND* section for the transmuted Pt atom. This section could be labelled, for example, "*PtXX*," where *XX* represents the index of the atom, and for the corresponding atom in the list of *XYZ* coordinates, the Pt label should be replaced with *PtXX*. The core correction is equal to the size of the perturbation, or the

mixing parameter λ , so for a perturbation of $\lambda = 0.1$, the core correction is set to 0.1, and the resulting nuclear charge of the atom is $Z + 0.1$, where Z is the original nuclear charge of the atom. Finally, the main caveat of this method is that the basis set of the reference atom must be supplied for the transmuted atom. We observed that the errors from doing this can become increasingly apparent for larger core corrections and rely on the reference system where the core correction is applied, however, these errors do also cancel in thermodynamic cycles.[95]

5.4 Modelled Materials

In previous work, we studied APDFT applications predicting binding energies for carbide, nitride, and oxide catalysts and identified shortcomings for predictions on semiconductors.[54] We assessed APDFT BE predictions of hydrogen on surfaces of ZnO(100), SnO₂(110), and TiO₂(110), where we considered unique combinations of sites that experienced isoelectronic transmutations of $|\Delta Z| = 1$ nuclear charge differences. Among all predictions, the mean absolute errors (MAE) for ZnO(100), SnO₂(110), and TiO₂(110) were 0.33, 0.82, and 0.96 eV, respectively, and we identified a qualitative trend where MAE is largest when the reference material has the most distinct band gap (TiO₂). We then observed that by doping Ti sites in TiO₂ with Pt, the band gap decreased, and we obtained more accurate APDFT BE predictions of H on surfaces where $|\Delta Z| = 1$ transmutations were made to the Pt sites. For transmutations to Pt dopants in the second layer of TiO₂ (Pt@2L), the MAE for H BE predictions decreased to 0.45 eV, and for Pt dopants in the first layer (Pt@1L), the MAE decreased further to 0.16 eV. In this report, we revisit APDFT predictions on these systems to address multiple points: the functional forms of the alchemical potential energy curves of adsorbate binding on semiconducting surfaces compared to metal surfaces, the changes in accuracy for higher ordered APDFT predictions in VASP, and the overall prediction of APDFT predictions on all these systems based on CP2K data.

We modelled hydrogen adsorption on multiple slabs with DFT, using each as a reference system to create hypothetical materials via individual transmutations to surface and subsurface sites. First, we modelled fcc Pt (111) with H adsorbed directly on top of a site, which serves as a baseline for a metallic system that APDFT accurately treats.[83] Next we modelled slabs of rutile TiO_2 (110) in pristine form, and with 50% of Ti atoms in the unit cell (three sites) replaced with Pt in the first layer (Pt@1L), and second layer (Pt@2L). H adsorption on a bridging oxygen site was modelled for all three slabs. We created hypothetical systems by making isoelectronic pairs of transmutations, where one transmutation of $\Delta Z = \pm 1$ occurs at one of several unique sites in the first or second layer of the surface, and the other transmutation of $\Delta Z = \mp 1$ occurs at one site chosen arbitrarily in the bottom layer. In the Pt slab, we chose five unique sites to transmute, with three in the first layer, and two in the second layer. In the pristine TiO_2 slab, six unique sites were transmuted, with three in the first layer (same chosen for Pt@1L), and three in the second layer (same for Pt@2L). One of the three Pt dopants in the Pt@1L and Pt@2L TiO_2 slabs was transmuted in the respective layer. Alchemical potential energy curves were then evaluated with DFT to illustrate energy pathways for the formation of each hypothetical system. Non-integer nuclear charge states were evaluated at 0.1 spaced increments along the $\lambda = 0 \rightarrow \lambda = \pm 1$ pathway for each transmutation. Points at $\lambda = \pm 0.01$ were also evaluated to compute higher-ordered derivatives with finite differences.

5.5 Computational Details

All calculations in VASP employed the PBE[65] exchange-correlation functional and projector augmented wave (PAW) pseudopotentials.[41] A planewave kinetic energy cutoff of 350 eV was used for the Pt slab models and 450 eV was used for TiO_2 slab models (including those with Pt dopants). For k-point sampling, Monkhorst-Pack grids of $8 \times 8 \times 1$ and $4 \times 4 \times 1$ were used for Pt and TiO_2 slab models, respectively. Structural geometries for all models were relaxed using iterative conjugate gradient minimization of forces between atoms until the difference in energies between consecutive iterations was less than 1 meV.

CP2K calculations [107] used a combined Gaussian and plane wave (GPW) formalism to optimize electronic structure,[108] where we chose double- ζ plus polarization (DZVP) basis sets optimized according to the Mol-Opt method.[109] We employed the pseudopotentials by Goedecker-Teter-Hutter to describe the interactions between valence electrons and the frozen atomic cores,[110] and the PBE[65] exchange-correlation functional was used again. A plane wave expansion cutoff energy of 400 Ry for Pt slab models and 800 Ry for TiO_2 slabs were used, and a k-point sampling of $4 \times 4 \times 1$ on Monkhorst-Pack grids was employed for all systems. Geometries were relaxed using the iterative minimization algorithm of Broyden-Fletcher-Goldfarb-Shanno (BFGS) until the forces converged to $4.5 \times 10^{-4} E_h$ per bohr. All energy cutoffs and k-point grid densities were determined through SCF energy convergence tests.

In this work, the target property approximated with APDFT is the binding energy (BE) of H on the hypothetical surfaces resulting from transmutations. The BE is evaluated with $BE = E_{\text{site}} + E_{\text{ads}} - E_{\text{ads-site}}$, which requires the energy of the catalyst surface without the adsorbate (E_{site}), the energy of the adsorbate not interacting with the catalyst (E_{ads}), and energy of the catalyst surface with the adsorbate bound ($E_{\text{ads-site}}$). BE predictions with APDFT rely on a thermodynamic cycle, which involves the energy of transmuting the same site(s) in both the slab model with (ads-site) and without (site) the adsorbate. For more details of APDFT employed in thermodynamic cycles, we refer the reader to our previous work.[83]. Here, alchemical potential energy curves depict BE as a function of alchemical changes, and will be referred as alchemical binding energy curves in the remainder of the report.

5.6 Results and Discussion

5.6.1 Overall Performance of APDFT

First, we will address the overall performance of BE predictions with APDFT-1 and APDFT-2 by analyzing the mean absolute error (MAE) among predictions on all hypothetical materials created by a reference system. Figure 17 shows the MAE for both APDFT-1 and APDFT-2 predictions on all reference systems using VASP (top row) and CP2K data (bottom row). MAE was evaluated separately for hypothetical systems resulting from $\Delta Z = -1$ (left column) and $\Delta Z = 1$ (right column) transmutations, and data were partitioned by groups of reference systems: TiO₂-1L, Pt-TiO₂-1L, TiO₂-2L, Pt-TiO₂-2L, and Pt-Slab. TiO₂-1L contains APDFT BE predictions on hypothetical materials formed from transmuting one of the three Ti sites selected in the first layer of TiO₂, and Pt-TiO₂-1L contains predictions on materials formed from transmuting Pt sites that replaced the original three Ti sites in the same layer. For TiO₂-2L and Pt-TiO₂-2L, these sets contained data for the hypothetical materials from transmutations of the three Ti sites and Pt sites in the second layer, respectively.

First we discuss the results from VASP (Figure 17 a-b). For $\Delta Z = -1$ (a), the MAE (eV) for APDFT-1 (dark blue) predictions on each data set compare following: TiO₂-2L (1.34) > TiO₂-1L (1.10) > Pt-TiO₂-2L (0.41) > Pt-TiO₂-1L (0.20) > Pt-Slab (0.08), and APDFT-2 (orange) predictions follow: TiO₂-2L (1.38) > TiO₂-1L (1.16) > Pt-TiO₂-2L (1.0) > Pt-Slab (0.29) > Pt-TiO₂-1L (0.18). For $\Delta Z = 1$ (b), the MAE (eV) compare following: TiO₂-1L (1.09) > TiO₂-2L (0.89) > Pt-TiO₂-2L (0.58) > Pt-TiO₂-1L (0.09) > Pt-Slab (0.07). APDFT-2 predictions follow: TiO₂-2L (1.82) > TiO₂-1L (1.32) > Pt-TiO₂-2L (0.83) > Pt-Slab (0.30) > Pt-TiO₂-1L (0.29). The first points to discuss is that though these data sets are smaller, we do see similar trends to previous work,[54] where the largest MAE values belong to pure TiO₂ systems (TiO₂-1L and TiO₂-2L), then MAE decreases first for Pt-doped TiO₂ in the second layer (Pt-TiO₂-2L) followed by further decreases for Pt-doped TiO₂ in the first layer (Pt-TiO₂-1L). Furthermore, the Pt slab system, exhibits the most accurate predictions.

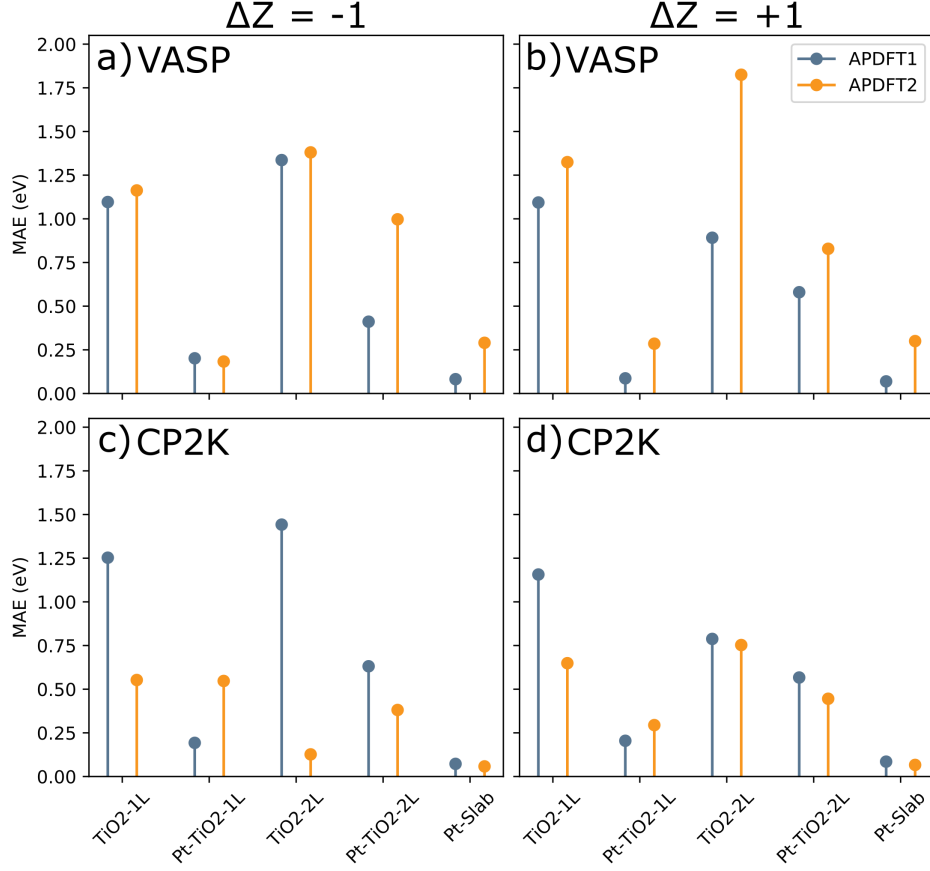


Figure 17: Plots showing mean absolute error (MAE) among APDFT-1 (dark blue) and APDFT-2 (orange) predictions on all reference systems using VASP (top row) and CP2K data (bottom row). MAE was evaluated separately for hypothetical systems resulting from $\Delta Z = -1$ (left column) and $\Delta Z = 1$ (right column) transmutations. Data was partitioned by groups of reference systems that include transmutations of three Ti sites in the first layer of TiO₂ (TiO₂-1L), three Pt sites in the same locations (Pt-TiO₂-1L), three Ti sites in the second layer (TiO₂-2L), Pt sites in the same locations (Pt-TiO₂-2L), and five sites Pt (Pt-Slab).

Moving to APDFT-2 results on VASP data, we would normally expect a decrease in MAE, as an additional term is included in the Taylor series to further approximate the alchemical binding energy curve. However, we observe an increase in MAE in Figure 17 (a-

b) for all data sets excluding Pt-TiO₂-1L in (a). The largest MAE increase among $\Delta Z = -1$ systems was +0.59 eV for Pt-TiO₂-2L references and +0.93 eV for TiO₂-2L references among the $\Delta Z = 1$ systems. The MAE decrease for transmutations of Pt-TiO₂-1L in (a) was only -0.02 eV.

We see a different patterns for APDFT predictions with CP2K data (Figure 17 c-d). For $\Delta Z = -1$ (c), the MAE (eV) for APDFT-1 (dark blue) predictions on each data set compare following: TiO₂-2L (1.44) > TiO₂-1L (1.25) > Pt-TiO₂-2L (0.63) > Pt-TiO₂-1L (0.19) > Pt-Slab (0.07), and APDFT-2 (orange) predictions follow: TiO₂-1L (0.55) = Pt-TiO₂-1L (0.55) > Pt-TiO₂-2L (0.38) > TiO₂-2L (0.13) > Pt-Slab (0.06). For $\Delta Z = 1$ (d), the MAE (eV) compare following: TiO₂-1L (1.16) > TiO₂-2L (0.79) > Pt-TiO₂-2L (0.57) > Pt-TiO₂-1L (0.20) > Pt-Slab (0.09), and APDFT-2 predictions follow: TiO₂-2L (0.75) > TiO₂-1L (0.65) > Pt-TiO₂-2L (0.45) > Pt-TiO₂-1L (0.29) > Pt-Slab (0.07). Again, we see similar trends to our previous work,[54] but some APDFT-1 errors are higher from CP2K data than the respective VASP data, such as the cases from pure TiO₂ references in (c). Notably, there are decreases in MAE for APDFT-2 predictions in all cases except Pt-TiO₂-1L. The largest MAE decrease among $\Delta Z = -1$ systems was -1.32 eV for TiO₂-2L references and -0.51 eV for TiO₂-1L references among the $\Delta Z = 1$ systems. The MAE increase for $\Delta Z = \pm 1$ transmutations of Pt-TiO₂-1L was +0.09 and +0.35 eV, respectively. We can further investigate the origin of these errors by constructing the binding energy curves along the alchemical pathway and comparing the resulting curve from the Taylor series approximation.

5.6.2 Alchemical Binding Energy Curves

In this section, we discuss alchemical binding energy curves from VASP and CP2K data for selected transmutations of each reference system (Figures 18 - 20). Alchemical curves for every transmutation studied in this work can be found in the Appendix. Overall, we see that APDFT-1 poorly replicates parabolic alchemical curves, and APDFT-2 remedies some of these systems, as the second order derivatives help capture the parabolic curvature, however APDFT-2 with VASP over-estimates the curvature for many cases.

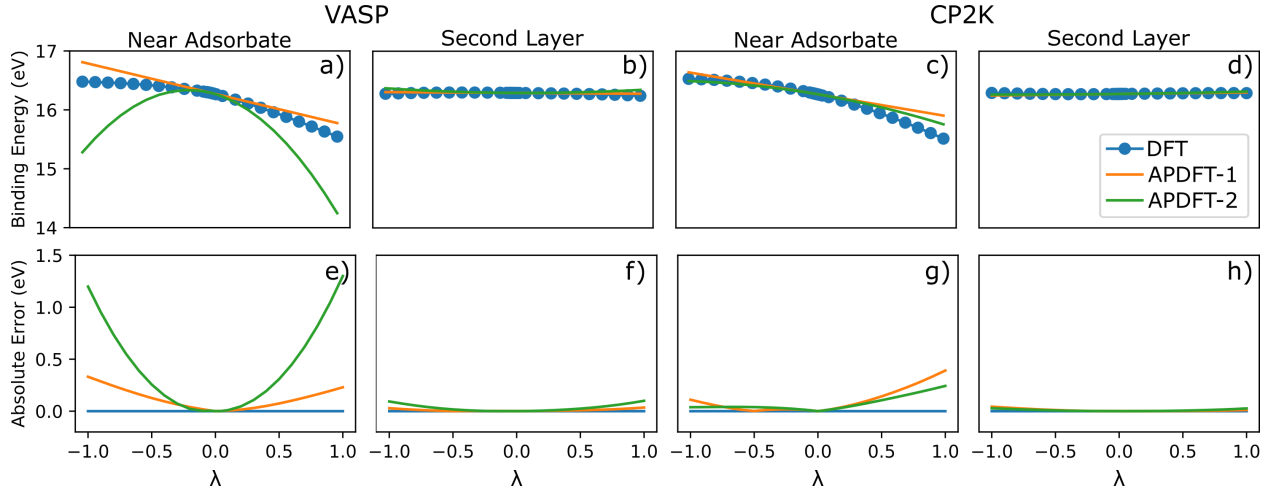


Figure 18: Alchemical binding energy curves of H on selected transmuted variations of Pt (111). The top row (a-d) depicts DFT data (blue circles) from VASP (a-b) and CP2K (c-d) with functions resulting from APDFT-1 (orange) and APDFT-2 (green) approximations. These curves depict BE along the pathway (λ) of $\Delta Z = \pm 1$ transmutations to the H binding site in the top layer (a,c) and a site in the second layer under the binding site (b,d). The bottom row (e-h) depicts absolute errors between the APDFT predicted functions and DFT, with horizontal lines (blue) drawn at zero to serve as a reference to guide the eye.

We first discuss Figure 18 to highlight the nature of alchemical binding energy curves for metallic systems that APDFT-1 normally describes accurately. The top row of figure 18 depicts DFT data (blue circles) for H binding on a Pt surface from VASP (a-b) and CP2K (c-d) with functions resulting from APDFT-1 (orange) and APDFT-2 (green) approximations. These curves result from $\Delta Z = \pm 1$ transmutations to the H binding site in the top layer

(a,c) and a site in the second layer under the binding site (b,d). The horizontal axes label the size of the perturbation with the mixing parameter (λ), where $\lambda = \pm 1$ marks a complete perturbation that results in a nuclear charge difference of $\Delta Z = \pm 1$. The bottom row (e-h) depicts absolute errors between the APDFT predicted functions and DFT data, with horizontal lines (blue) drawn at zero to serve as a reference to guide the eye.

For VASP (a-b), we see a distinct difference in the shape of the alchemical curve between the two transmutation sites. The alchemical curve is a wide parabola with lower curvature for the transmutation at the binding site, and at the site in the second layer, the alchemical curve is nearly flat. The shapes of the curves for all other sites in Pt are similar to the latter, where the curves are linear with a very low slope or flat. This helps explain why APDFT-1 predictions are typically less accurate for a transmutation near the adsorbate. Since it only utilizes first derivatives, APDFT-1 approximates a sloped line, which will diverge from a parabolic alchemical curve that has increasing derivatives. Second order derivatives in the Taylor series can accurately replicate the parabola, however, in Figure 18 (a), the resulting APDFT-2 function is more steep. This means that the second order derivatives evaluated with VASP data are too large, and we can link this to the MAE increase for APDFT-2 predictions for Pt systems in Figure 17 (a-b).

For the CP2K (c-d), we also see a parabolic function resulting from a transmutation to the binding site and a flat line for the transmutation in the second layer. APDFT-1 is also less accurate for the transmutation near the adsorbate as the linear function diverges, however, the parabolic function from APDFT-2 matches the DFT data more closely. We can again link these observations to the decreases in MAE for APDFT-2 predictions for Pt systems in Figure 17 (c-d). Although these predictions are more accurate compared to VASP, the absolute errors are greater for the $\lambda = +1$ target. This may be due to asymmetries of the second derivatives in the negative and positive directions in λ .

Next we will discuss the alchemical binding energy curves for TiO_2 -based systems from VASP data in Figure 19 with plots constructed the same as Figure 18. Figure 19 (a,c) contain alchemical binding energy curves for transmutations to sites in the pure TiO_2 reference system, and (b,d) show transmutations to Pt sites in Pt-doped TiO_2 reference systems. Figure 19 (a) depicts the alchemical binding energy curve for $\Delta Z = \pm 1$ transmutations to

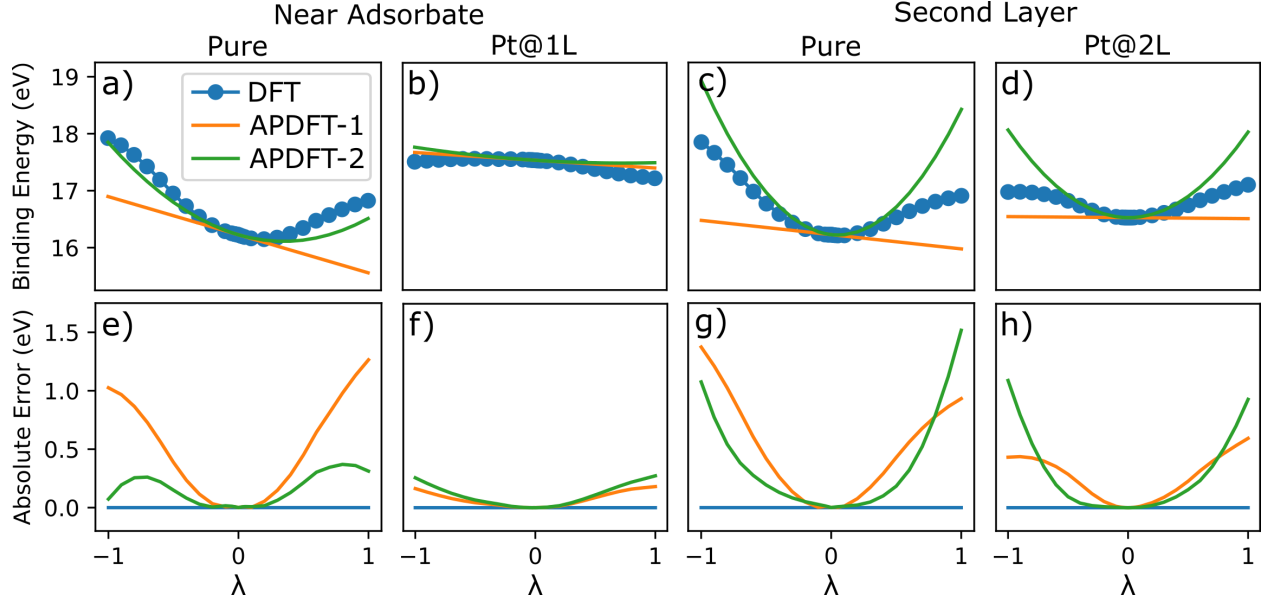


Figure 19: Alchemical binding energy curves of H on selected transmuted variations of TiO_2 -based systems from VASP data. The top row (a-d) depicts DFT data (blue circles) with functions resulting from APDFT-1 (orange) and APDFT-2 (green) approximations. These curves depict BE along the pathway (λ) of $\Delta Z = \pm 1$ transmutations to sites in the pure TiO_2 reference system (a,c), and transmutations to Pt sites in Pt-doped TiO_2 reference systems (b,d). The bottom row (e-h) depicts absolute errors between the APDFT predicted functions and DFT, with horizontal lines (blue) drawn at zero to serve as a reference to guide the eye.

a Ti site near the bridging oxygen site where H is bound, and (c) depicts the curve for transmutations to a Ti site in the second layer. Figure 19 (b) and (d) show transmutations to the Pt site located in the same position as the Ti site for the system in (a) and (c), respectively.

We see that for pure TiO_2 , the alchemical binding curves exhibit the shape of a high ordered polynomial, and similar to the case with Pt slabs, the linear approximations from APDFT-1 diverge from the target values at $\lambda = \pm 1$. For parabolic approximations with APDFT-2, the absolute errors are much lower for values near $\lambda = 0$, but near the target values, the results depend on the site of the transmutation. For the transmutation near

the adsorbate, the APDFT-2 errors are much lower than APDFT-1 at the target values at $\lambda = \pm 1$, and for transmutations in the second layer, the APDFT-2 errors are comparable to APDFT-1. The curves for the remaining sites (Appendix) show similar behavior to (c). This suggests that for transmutations near the adsorption site, the second derivatives are lower in magnitude, and APDFT-2 agrees more closely to the DFT. However, if extrapolated, the function from APDFT-2 would likely diverge from the actual alchemical curve, as it increases at a greater rate. Overall, these observations explain the increase in MAE for TiO_2 systems in Figure 17 (a-b).

Pt-doped TiO_2 systems show different behavior based on the layer. Transmuting the Pt dopant is in the top layer (b), the alchemical curve is flat and APDFT errors are very low, comparable to pure Pt systems, and transmuting Pt in the second layer (d), the shape of the curve and the accuracy of APDFT are more similar to TiO_2 systems in (c). These results highlight again the increases in MAE in Figure 17 (a-b) as well as the trends observed in previous work,[54] where APDFT-1 errors were significantly more accurate for Pt-doped TiO_2 in the first layer than pure TiO_2 .

Figure 20 depicts alchemical binding energy curves from CP2K data for the same transmutations sites in reference systems as Figure 19. Again, the alchemical curves for transmutations of pure TiO_2 (a,c) exhibit the shape of a high-order polynomial, and the linear approximations from APDFT-1 diverge from the target values. For APDFT-2, the errors decrease significantly compared to VASP, and the resulting parabolic functions are similar between both transmutation locations. However, the errors are asymmetric for the transmutation in the second layer, with the prediction for $\Delta Z = -1$ more accurate than $\Delta Z = 1$, and extrapolating the resulting APDFT-2 functions in (a) and (c) would again lead to diverging results from the actual alchemical curve.

We see different shapes of alchemical curves for Pt-doped TiO_2 systems (b,d). For transmutations to Pt dopant in the top layer (b), the curve closely resembles a shallow cubic function, and the APDFT predictions are thus more accurate for $\lambda > 1$. The APDFT-2 function also appears very linear because the second derivatives, evaluated at the inflection point of the cubic function, are very low in magnitude. This explains why Pt- TiO_2 -1L for $\Delta Z = -1$ transmutations is the only subset of data that sees an increase in MAE for

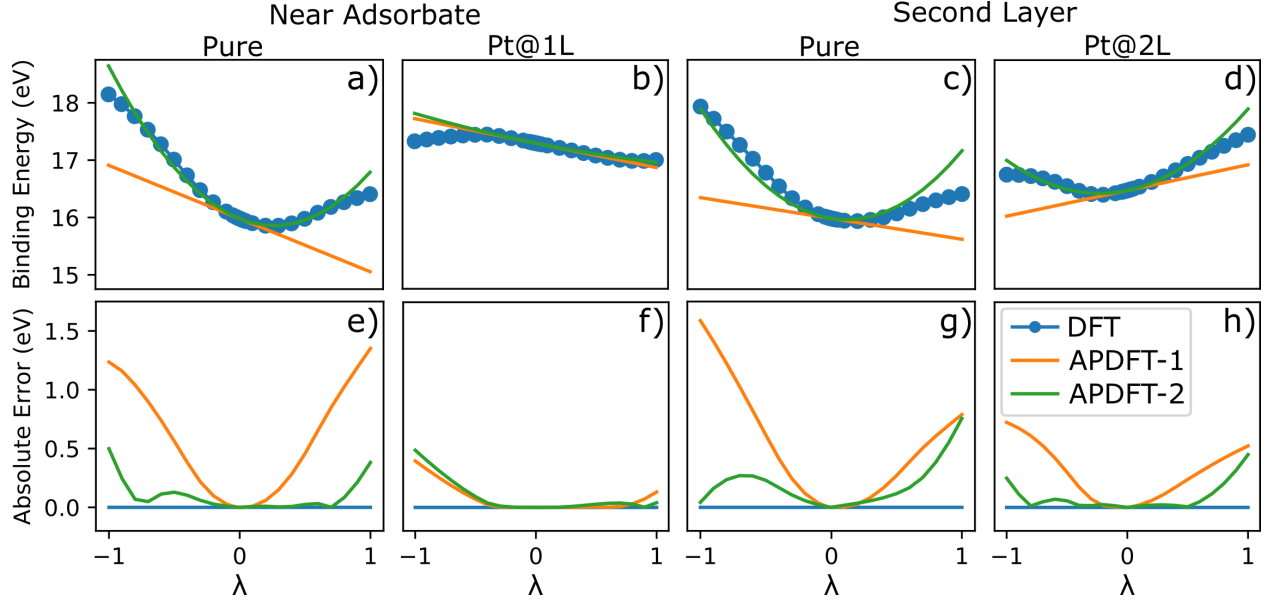


Figure 20: Alchemical binding energy curves of H on selected transmuted variations of TiO_2 -based systems from CP2K data. The top row (a-d) depicts DFT data (blue circles) with functions resulting from APDFT-1 (orange) and APDFT-2 (green) approximations. These curves depict BE along the pathway (λ) of $\Delta Z = \pm 1$ transmutations to sites in the pure TiO_2 reference system (a,c), and transmutations to Pt sites in Pt-doped TiO_2 reference systems (b,d). The bottom row (e-h) depicts absolute errors between the APDFT predicted functions and DFT, with horizontal lines (blue) drawn at zero to serve as a reference to guide the eye.

APDFT-2 predictions in Figure 17 (c). For transmutations to the Pt dopant in the second layer (d), the resulting functions are similar to the pure TiO_2 cases, where APDFT-1 is inaccurate for $\lambda = \pm 1$, and while APDFT-2 improves the accuracy, the parabolic function likely diverges from the alchemical curve upon extrapolation.

Overall, we see that the curvature for alchemical binding energy curves varies based on the system. For transmutations that are far from the adsorption site on a metal surface, the derivatives of the alchemical function are very low, and APDFT-1 reliably treats these systems. Transmuting the binding site in a metal leads to a more significant binding energy change, and the function matches a second order polynomial. APDFT-1 is unable to

approximate this function accurately because the Taylor series relies on the first derivative of the function near a local minimum, and since the magnitude of this derivative is so low, the resulting APDFT-1 function has a lower slope and the function diverges from curve that increases in slope. For the semiconducting system, we observed that most alchemical curves were high-order polynomials, and by adding metal dopants, the derivatives of these curves lessen. While APDFT-1 poorly approximates the high-order polynomials, and APDFT-2 provides higher quality treatments, with high accuracy for regions on the alchemical curve near the reference energy. The performance of APDFT-2 in the regions near the target energies varies based on the DFT code as the large magnitudes of second derivatives from VASP over-predict the energies, while the magnitudes of derivatives from CP2K are more reasonable and the errors are lower.

This shortcoming of VASP is likely due to the virtual crystal approximation approach to modeling the non-integer nuclear charged states. To be fully compatible with the VCA implementations the PAW pseudopotentials need to be reconstructed so that they are similar between interpolating atoms.[106] We found that this method drastically breaks down for atoms with large cores like the transition metals (see Appendix). We constructed alchemical curves for $\Delta Z = \pm 1, 2$ isoelectronic transmutations of simple dimers containing atoms with small cores (N_2) and large cores (Pt_2). For N_2 , the alchemical curves showed smooth transitions for the two sizes of transmutation, but for Pt_2 , the curves did not align in the region where identical non-integer nuclear charge states are being modelled with different mixtures of pseudopotentials. Similar curves were constructed by applying core corrections to N_2 and Pt_2 in CP2k, and both alchemical curves showed smooth transitions from both transmutations.

5.7 Conclusion

In summary, we have presented a guide and analysis of catalysis applications with APDFT approximations up to second order. First we outlined two approaches to conduct APDFT with quantities evaluated with VASP and CP2K and compared the results directly for identical systems. VASP produces atom-centered electrostatic potentials (density derivatives), that produce alchemical derivatives via a dot product of said values with nuclear charge differences from transmutations. with CP2K, alchemical derivatives are evaluated via Hellmann-Feynman Theorem, which requires a difference of two energy values evaluated with the same wavefunction parameters from restart files. Second order derivatives were obtained with central finite differences of first derivatives evaluated at small, non-integer perturbations the reference state. The functions that result from the first (APDFT-1) and second (APDFT-2) order Taylor series approximations were compared to alchemical curves of the binding energy changes constructed with DFT-evaluated energies of non-integer nuclear charge states between reference and target values. Energies of non-integer states were obtained with the virtual crystal approximation in VASP, which superimposes pseudopotentials of atoms in the transmutation based on user-supplied weights, and with core corrections in CP2K, which defines an effective nuclear point charge for an atom.

Overall, we identified that APDFT-1 poorly approximates alchemical binding energy curves that exhibit shapes of second order, or higher, polynomials, and APDFT-1 accurately describes linear alchemical binding energy curves. Transitioning to APDFT-2, approximations with VASP data saw further decrease in accuracy from APDFT-1, while accuracy improved for most cases with CP2K data. APDFT-2 approximations with VASP over-predicted the curvature of the alchemical functions due to large second order derivatives computed from data at non-integer states that come from implementations of the virtual crystal approximation that are unreliable for atoms with large cores. For future uses of APDFT with periodic DFT calculations, we recommend CP2K as a more reliable code to extrapolate catalyst properties with up to second order APDFT. The procedures outlined here capably produce second order derivatives that allow reliable screening for a variety of reference systems. Finally, this work also presents opportunities for further developments

with periodic DFT software packages, such as developments for PAW potentials to be more compatible with tools that evaluate small perturbations of systems like the VCA in VASP, and optimizing codes to evaluate quantities at greater numerical precision, so that higher orders of APDFT may be feasible.

6.0 Conclusions

In summary, this dissertation presents promising possibilities for computational screening studies of heterogeneous catalysts with APDFT. APDFT approximates energy values through alchemical energy derivatives in a Taylor series truncated up to a number of terms, and APDFT offers the advantage of producing many hypothetical catalyst descriptors from a limited number of calculations for one reference catalyst system. In Chapters 2 and 3 we outlined two examples of first order APDFT reliably predicting catalyst descriptors like adsorbate BE and reaction energy barriers for numerous hypothetical states. For BE predictions of H and OH on TiC, TiN, and TiO, we observed accuracy within 0.33 eV, which is comparable to similar predictions for transition metal catalysts. Conversely, first order APDFT poorly approximated these quantities with semiconducting TiO₂, but by replacing sites with Pt dopants, adding states at the Fermi level, the accuracy improves. We then found that BE predictions along a reaction pathway of CH₄ dehydrogenation on Pt, provided by a set of NEB images, could additionally provide accurate energy barrier measurements, within 0.3 eV, for 32 variations of the reference images.

Chapter 4 showed that there are even advantages coupling machine learning with APDFT. By training support vector regression models on a large data set of APDFT-predicted BEs for different classes of adsorbates, we could correct the original errors, leading to more accurate BE predictions with ML that requires less training data than using pure DFT-based training sets. This procedure presented a route to obtain thousands of additional BE predictions on materials formed from multiple and/or larger transmutations, and we learned how these types of transmutations influence APDFT errors. Finally, we scrutinized second order APDFT approximations in Chapter 5, and illustrated how both first and second order APDFT predictions agree with DFT by constructing alchemical curves that depict the energy as a function of the transmutation pathway. Systems that are typically described well with first order APDFT reveal linear alchemical energy functions, and systems not accurately described reveal functions that are high order polynomials. While second order APDFT includes a higher quality description of the curvature, accurate second derivatives could not be

obtained with current implementations of VASP. However, DFT results from CP2K, which allow core corrections to atomic nuclear charges, offered sufficient second derivatives, and second order APDFT predictions improved in accuracy.

6.1 Future Work

Our findings suggest many routes for future development and applications of APDFT. First, there are ambitious developments for future versions of DFT software that could elevate APDFT. In VASP, there are limitations to the kinds of species that can be evaluated reliably with the virtual crystal approximation. This could potentially be addressed by adjusting pseudopotentials so that linear combinations of two potentials are more compatible for large core elements, such as transition metals in later rows of the d-block. Furthermore, it is evident that alchemical binding energy curves can exist as higher order polynomials. APDFT-2 is currently not sufficient at replicating the curvature of these functions in regions far from the reference state ($\lambda = 0$), and for further extrapolations, approximations will likely diverge as parabolic approximations continuously increase. An obvious next step is to incorporate further approximations with APDFT-3, and APDFT-4, which include third and fourth order derivatives that are obtainable from finite differences of additional states evaluated at small values of λ . However, these quantities must be evaluated at higher numerical precision so that rounding errors aren't accentuated with the increasing $n!^{-1}$ term. By implementing arbitrary precision in future versions of DFT software packages, achieving higher ordered APDFT approximations becomes more feasible and opens opportunities to better treat more complicated alchemical pathways.

Another interesting approach to assess would be piece-wise APDFT approximations. It may be possible to achieve higher quality alchemical functions that reach further points in materials space by evaluating derivatives at higher λ values. For example, derivatives at $\lambda = \pm 0.5$ states should provide more accurate descriptions of the curvature in the regions at $\lambda = \pm 1$ and beyond. This approach could be especially useful for cases like TiO_2 , where derivatives quickly deviate from parabolic behavior in regions closer to $\lambda = \pm 1$.

While many of these approaches call for further DFT calculations, these requirements become less concerning for high-throughput screening applications of APDFT. So far, we have only applied APDFT-2 to single transmutations to sites in catalyst surfaces that undergo $\Delta Z = \pm 1$ nuclear charge changes. An open question still exists regarding the limits of target systems that can be evaluated with second order APDFT, ie. forming hypothetical materials with multiple and/or larger transmutations. We can simultaneously study the performance of APDFT-2 on these systems and observe how the alchemical curves vary. Going forward, with the ability to replicate alchemical curves for systems resulting from multiple transmutations at variable ΔZ , the cost of additional DFT calculations required for increasing orders of APDFT quickly becomes insignificant. For one site, five states must be evaluated along λ to make APDFT-4 approximations with finite differences. Thus, $5N$ DFT calculations are required to evaluate any combination of transmutations to N sites in a catalyst slab model. For a slab with 8 atoms, 40 initial DFT calculations are required, but the number of ways to combine 1, 2, 3, or 4 simultaneous transmutations results in 162 systems that APDFT-4 can evaluate instantly. This only considers a fixed magnitude of ΔZ for each transmutation, and thus the number of combinations could continue to increase with this variable.

This offers much promise for ambitious screening studies of hypothetical catalysts across materials space. With higher orders of APDFT, the search space extends with countless combinations of materials configurations, and we can simultaneously measure multiple descriptors that measure catalyst activity and stability. Our work shows that we can use APDFT to measure energies of reaction intermediates bound to surfaces and subsequently construct an entire reaction energy profile. This can be extended to multiple pathways, allowing us to evaluate entire reaction networks for thousands of candidates. Additionally, quantities such as surface energies and adsorbate-induced segregation energies are currently untouched by these schemes. These metrics are crucial for knowing stable surface states under operating conditions, as they are used to construct high-quality phase diagrams. However, a single phase diagram requires surface energies for numerous configurations. APDFT may leverage the effort to construct supplemental phase diagrams of thousands of hypothetical surfaces that stem from a single reference material. Overall, the findings in this dissertation serve as a collection of approaches with APDFT that give future researchers opportunities to im-

prove or develop schemes that enhance the capabilities of APDFT or simultaneously screen properties listed here using large, existing data sets to explore vast reaches of hypothetical materials space.

Appendix A Supporting Information for Machine Learning Corrected APDFT

A.1 Numbers of Systems Evaluated

In this section, we explain how many unique binding energies (BEs) were calculated. We considered twelve adsorbates in total (CH_x , NH_x , with $x = 0 - 3$, OH_x , with $x = 0 - 2$, and OOH) and modelled multiple adsorbate coverages ($\theta = 1, 1/4, 1/9$) using 1×1 , 2×2 , and 3×3 surface slab dimensions. Hypothetical alloys were formed from each slab by making 1-4 numbers of transmutations (N_T) with nuclear charge change of $\Delta Z = -3, -2, -1, 1, 2, 3$ (six possibilities). For the 1×1 and 2×2 slabs, the total number of hypothetical alloys we made was equal to the number of ways you can choose N_T changes of ΔZ from all sites in the top two layers (two and eight, for 1×1 and 2×2 respectively). For 1×1 , we could only do $N_T = 1$ or 2 transmutations, with six possible choices in ΔZ , and twelve adsorbates. The total number of systems evaluated was $6 \cdot 12 \cdot \left(\binom{2}{1} + \binom{2}{2}\right) = 216$, where the binomial notation is used for showing the resulting combinations of N_T transmuted atom in the two atomic sites in the top two layers. For the 2×2 slab, the total number of systems was $6 \cdot 12 \cdot \left(\binom{8}{1} + \binom{8}{2} + \binom{8}{3} + \binom{8}{4}\right) = 11,664$. For the 3×3 slab, we could not evaluate all possible systems because this would require as many as 291,384 DFT calculations for benchmarking. Instead, we chose approximately 2,000 systems at random. These details are outlined explicitly in Table 2 below, where we list the number of systems evaluated in each adsorbate class broken down by the adsorbate coverage and the data set partitions. In the model construction phase, we split each data set into 80% belonging to the training set and 20% belong to the test set.

Table 2: Detailed breakdown of the number of systems evaluated for each adsorbate class with respect to the adsorbate coverage and partition of the data sets (total, train, or test). We also include the overall sum of the data set partitions per group of adsorbates (ads.) and per coverage (θ)

Data Set	Partition	$\theta = 1$	$\theta = 1/4$	$\theta = 1/9$	Overall (ads.)
CH _x	Total	72	3888	584	4544
	Train	55	3109	471	3635
	Test	17	779	113	909
NH _x	Total	72	3888	622	4582
	Train	58	3114	493	3665
	Test	14	774	129	917
OH _x	Total	72	3888	797	4757
	Train	60	3109	636	3805
	Test	12	779	161	952
Overall (cov.)	Train	173	9332	1600	
	Test	43	2332	403	
	Total	216	11664	2003	

A.2 Feature Vector Construction

We represented the dopant location in the alloys by considering a feature vector with length equal to the number of atoms in each surface layers of our slab models, and each element of the vector represents an atomic site. This presents a problem where the length of the vector varies depending on the size of the surface unit cell used to model varying adsorbate coverage. We remediated this by normalizing the vector, making its length equal to the number of atoms in a layer of the largest unit cell included in the dataset. This is illustrated in Figure 21, where we show top views of the first and second layers of three varying sized unit cells of Pt with Au dopants. Since a 3x3 unit cell of fcc Pt(111) was the

largest unit cell included, and we considered all atoms in the top two layers, the length of the feature vector 18 elements. Additionally, the positions of each atom in the vector are set according to their positions in the 3x3 unit cell. In Figure 21, the atoms in the first layer are indexed from 0 to 8, following the atom ordering from the unit cell origin at the bottom left to the top right. Since atoms indexed 0, 4, and 8 are Au dopants, we encode the vector with 1's at these positions. If we follow this constraint when encoding changes to the first layer of a 2x2 unit cell, we would only allow the elements whose positions match those in the 3x3 unit cell to be changed to 1's, and the remaining elements would be a 0. According to Figure 21, we are only able to encode atoms indexed by 0, 1, 3, 4, 9, 10, 12, and 13 in a 2x2 unit cell. Since atoms 0 and 4 are Au, those two positions are encoded as 1's. The padded zeroes, for positions that can't be encoded for a 2x2 unit cell, are colored in red. For a 1x1 unit cell, only atoms indexed by 0 and 9 can be encoded, and atom 0 is an Au dopant, so the zeroth position of the vector is changed to 1. The remaining positions that can't be changed are the padded zeroes colored in red. A second group of examples are illustrated in Figure 22 that include Ir dopants arranged in various configurations in Pt slab models with adsorbates.




Figure 21: Fingerprinting scheme to construct feature vectors that describe dopant locations.

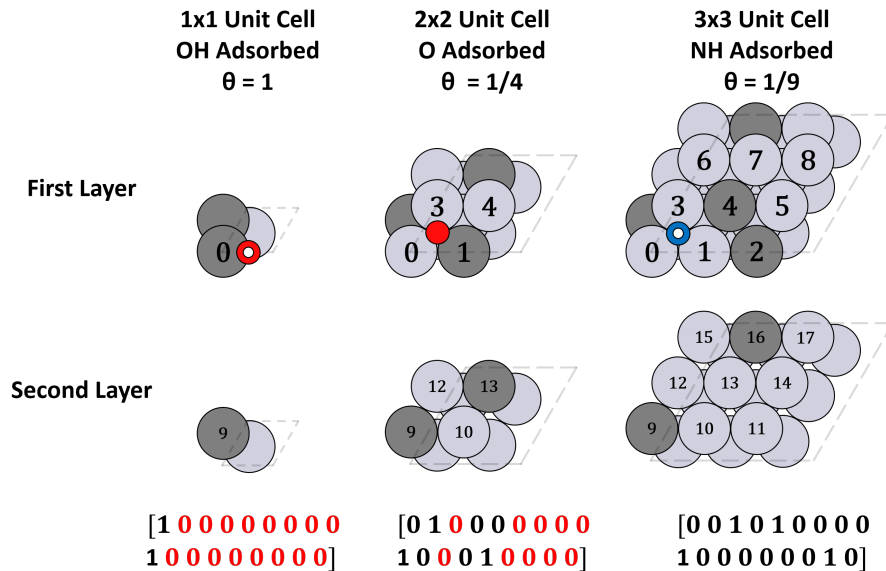


Figure 22: Fingerprinting scheme that describes dopant locations in Pt slab models with adsorbates.

For features like type of adsorbate, coverage, and resulting atom, we add an element to our vector for each possibility in those groups. For example, we add three elements that represent possibilities of θ , and the first, second, and third pertain to coverages of $1/9$, $1/4$, and 1 , respectively. If our input alloy is a 3×3 unit cell, the first element would equal 1 and the second and third would be zero.

A.3 Model Validation

Our models were validated with a test set partitioned from the entire dataset. A test set is necessary because if you make predictions on data using a model that was trained with the same data, the conclusions are trivial. Typically, the data is split at a ratio where a majority of the data is used to train our model, and with the amount of remaining test data, a

reasonable number of predictions are made to validate the model’s performance. This tactic is useful when addressing overfitting, a common problem that is encountered when trying to make reliable and transferable machine learning models. A model is overfit when the fit is optimized so close to the training data that the performance is poor when making predictions from new data. Overfitting can be avoided with resampling techniques that employ the train-test split approach like k -fold cross validation (CV). In k -fold CV, the training set is split into k numbers of smaller datasets. For k iterations, a model is trained using $k - 1$ folds of training data from the k -fold split dataset, and the resulting model is validated on the remaining fold using the mean absolute error (MAE). For our model selection, we performed 10-fold CV on each, and the average MAE of all k iterations measured the performance of the model.

A.4 NH_x Parity Plots

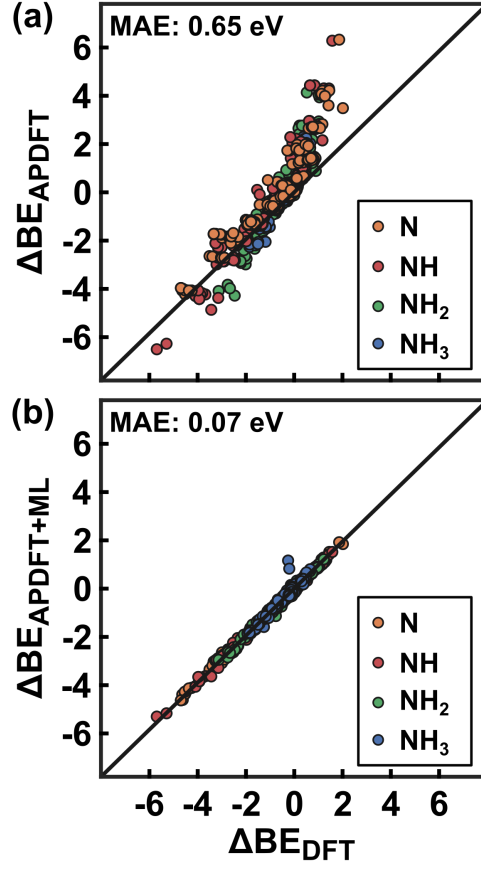


Figure 23: Parity plot of the BE change (ΔBE) of NH_x adsorbates on alloys of transmuted Pt predicted with APDFT (a) and ML-corrected APDFT (b) compared to DFT. ΔBE are in units of eV. The MAE decreases from 0.65 (a) to 0.07 eV (b) upon ML-correction.

A.5 OH_x and OOH Parity Plots

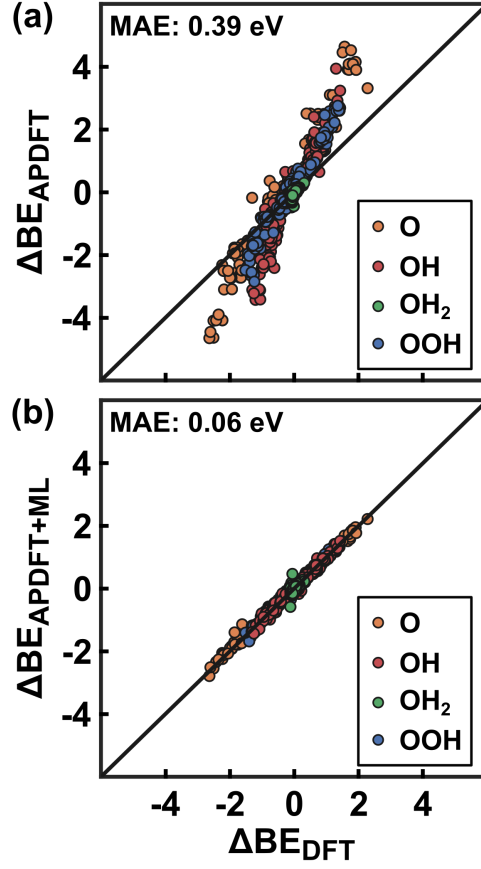


Figure 24: Parity plot of the BE change (ΔBE) of OH_x adsorbates on alloys of transmuted Pt predicted with APDFT (a) and ML-corrected APDFT (b) compared to DFT. ΔBE are in units of eV. The MAE decreases from 0.39 (a) to 0.06 eV (b) upon ML-correction.

A.6 Error Analysis by Alloy Type

Figure 7 in the main text provides a breakdown of errors across the full training set, where we specify the percentage of data and reduction of errors with respect to ΔZ and N_T . The merits represented in that figure are tabulated below:

Table 3: Distribution of the types of hypothetical alloys found in the test data set and how well ML improves APDFT predictions for those types of alloys.

ΔZ	N_T	% of Test Data Set	Δ MAE (eV)
1	1	3.1	0.00
1	2	7.3	0.07
1	3	9.8	0.16
1	4	12.9	0.25
2	1	3.7	0.13
2	2	7.1	0.37
2	3	10.5	0.58
2	4	12.5	0.76
3	1	2.5	0.12
3	2	7.3	0.59
3	3	10.7	0.87
3	4	12.9	1.23

Appendix B Supporting Information for Second Order APDFT Studies

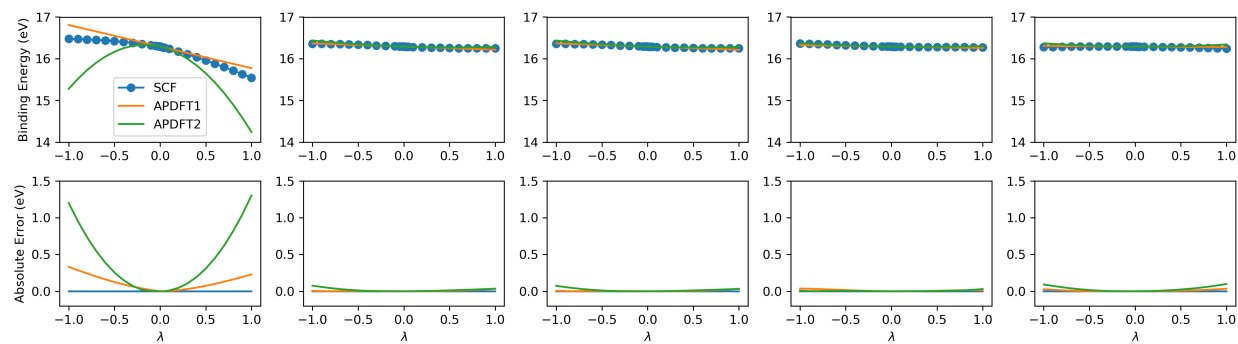


Figure 25: Alchemical binding energy curves of H on all transmuted variations of Pt (111) (VASP)

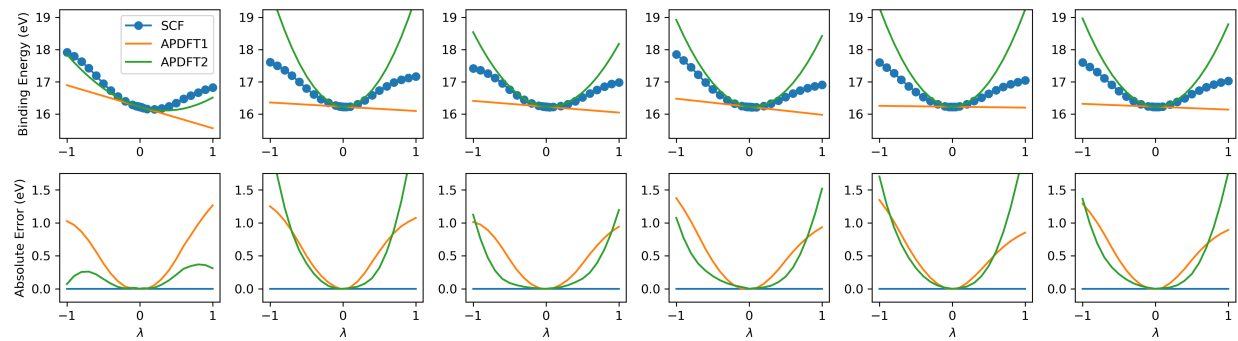


Figure 26: Alchemical binding energy curves of H on all transmuted variations of TiO_2 (VASP)




Figure 27: Alchemical binding energy curves of H on all transmuted variations of (Pt@1L)-TiO₂ (VASP)




Figure 28: Alchemical binding energy curves of H on all transmuted variations of (Pt@2L)-TiO₂ (VASP)

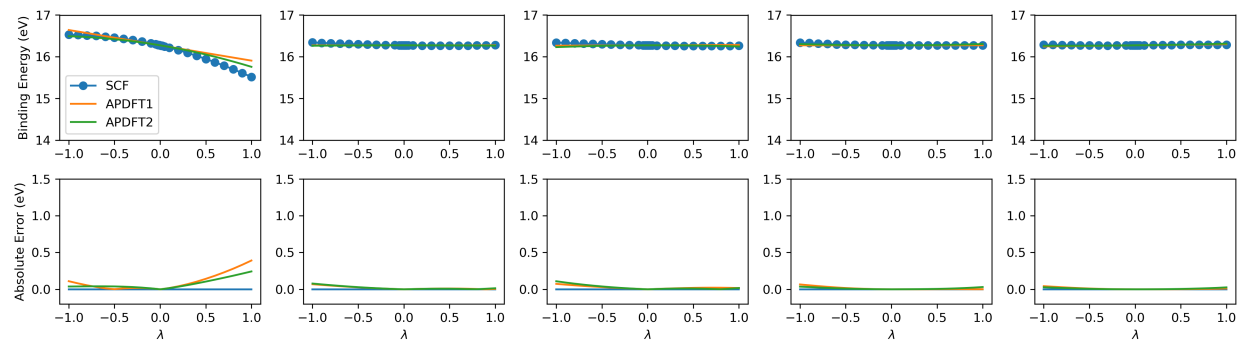


Figure 29: Alchemical binding energy curves of H on all transmuted variations of Pt (111) (CP2K)

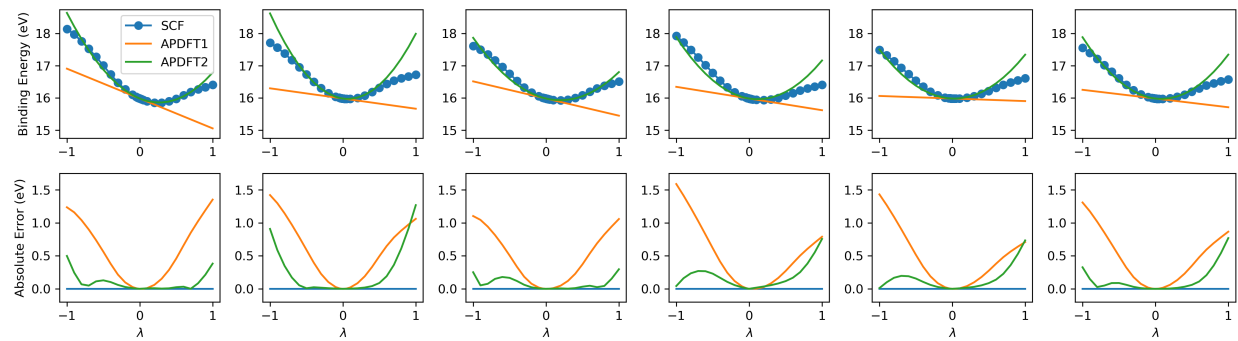


Figure 30: Alchemical binding energy curves of H on all transmutated variations of TiO_2 (CP2K)




Figure 31: Alchemical binding energy curves of H on all transmuted variations of (Pt@1L)-TiO₂ (CP2K)




Figure 32: Alchemical binding energy curves of H on all transmuted variations of (Pt@2L)-TiO₂ (CP2K)

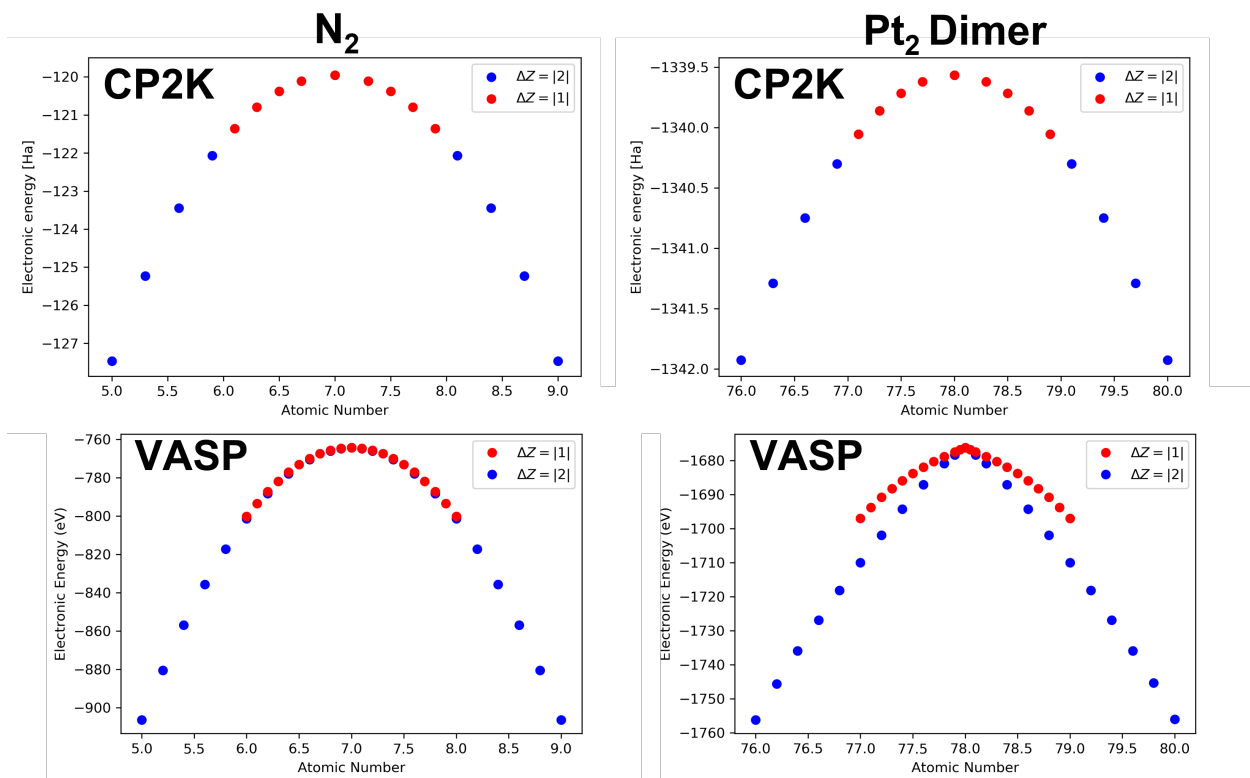


Figure 33: Alchemical energy curves of $\Delta Z = \pm 1, 2$ transmutations to N_2 and Pt_2 dimers with VASP and CP2K data.

Bibliography

- [1] B. W. J. Chen, L. Xu, and M. Mavrikakis, “Computational Methods in Heterogeneous Catalysis,” *Chemical Reviews*, vol. 121, no. 2, pp. 1007–1048, 2021.
- [2] B. R. Goldsmith, J. Esterhuizen, J. X. Liu, C. J. Bartel, and C. Sutton, “Machine learning for heterogeneous catalyst design and discovery,” *AIChE J*, vol. 64, no. 7, pp. 2311–2323, 2018.
- [3] Z. W. Seh, J. Kibsgaard, C. F. Dickens, I. Chorkendorff, J. K. Nørskov, and T. F. Jaramillo, “Combining theory and experiment in electrocatalysis: Insights into materials design,” *Science*, vol. 355, no. 6321, p. eaad4998, 2017.
- [4] J. Greeley, “Theoretical Heterogeneous Catalysis: Scaling Relationships and Computational Catalyst Design,” *Annual Review of Chemical and Biomolecular Engineering*, vol. 7, no. 1, pp. 605–635, 2016.
- [5] A. Kulkarni, S. Siahrostami, A. Patel, and J. K. Nørskov, “Understanding Catalytic Activity Trends in the Oxygen Reduction Reaction,” *Chemical Reviews*, vol. 118, no. 5, pp. 2302–2312, 2018.
- [6] K. Tran and Z. W. Ulissi, “Active learning across intermetallics to guide discovery of electrocatalysts for CO₂ reduction and H₂ evolution,” *Nat. Catal.*, vol. 1, no. 9, pp. 696–703, 2018.
- [7] J. Greeley and M. Mavrikakis, “Alloy catalysts designed from first principles,” *Nat. Mater.*, vol. 3, no. 11, pp. 810–815, 2004.
- [8] V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler, “Ab initio molecular simulations with numeric atom-centered orbitals,” *Computer Physics Communications*, vol. 180, no. 11, pp. 2175–2196, 2009.
- [9] S. D. Miller and J. R. Kitchin, “Uncertainty and figure selection for DFT based cluster expansions for oxygen adsorption on Au and Pt (111) surfaces,” *Mol Simul*, vol. 35, no. 10-11, pp. 920–927, 2009.
- [10] L. M. Herder, J. M. Bray, and W. F. Schneider, “Comparison of cluster expansion fitting algorithms for interactions at surfaces,” *Surf Sci*, vol. 640, pp. 104–111, 2015.

- [11] F. Calle-Vallejo, D. Loffreda, M. T. M. Koper, and P. Sautet, "Introducing structural sensitivity into adsorption–energy scaling relations by means of coordination numbers," *Nature Chemistry*, vol. 7, no. 5, pp. 403–410, 2015.
- [12] G. Yang, J. Wu, S. Chen, W. Zhou, J. Sun, and G. Chen, "Size-independent neural networks based first-principles method for accurate prediction of heat of formation of fuels," *J Chem Phys*, vol. 148, no. 24, p. 241738, 2018.
- [13] L. T. Roling, T. S. Choksi, and F. Abild-Pedersen, "A coordination-based model for transition metal alloy nanoparticles," *Nanoscale*, vol. 11, no. 10, pp. 4438–4452, 2019.
- [14] Z. W. Ulissi, A. J. Medford, T. Bligaard, and J. K. Nørskov, "To address surface reaction network complexity using scaling relations machine learning and DFT calculations," *Nat Commun*, vol. 8, p. 14621, 2017.
- [15] Z. Li, X. Ma, and H. Xin, "Feature engineering of machine-learning chemisorption models for catalyst design," *Catalysis Today*, vol. 280, pp. 232–238, 2017.
- [16] Z. Li, S. Wang, W. S. Chin, L. E. Achenie, and H. Xin, "High-throughput screening of bimetallic catalysts enabled by machine learning," *Journal of Materials Chemistry A*, vol. 5, no. 46, pp. 24131–24138, 2017.
- [17] J. R. Boes and J. R. Kitchin, "Neural network predictions of oxygen interactions on a dynamic Pd surface," *Molecular Simulation*, vol. 43, no. 5-6, pp. 346–354, 2017.
- [18] J. P. Janet and H. J. Kulik, "Resolving Transition Metal Chemical Space: Feature Selection for Machine Learning and Structure-Property Relationships," *J Phys Chem A*, vol. 121, no. 46, pp. 8939–8954, 2017.
- [19] J. P. Janet, L. Chan, and H. J. Kulik, "Accelerating Chemical Discovery with Machine Learning: Simulated Evolution of Spin Crossover Complexes with an Artificial Neural Network," *J Phys Chem Lett*, vol. 9, no. 5, pp. 1064–1071, 2018.
- [20] O. A. von Lilienfeld, R. D. Lins, and U. Rothlisberger, "Variational Particle Number Approach for Rational Compound Design," *Physical Review Letters*, vol. 95, no. 15, p. 153002, 2005.
- [21] O. A. von Lilienfeld and M. E. Tuckerman, "Molecular grand-canonical ensemble density functional theory and exploration of chemical space," *The Journal of Chemical Physics*, vol. 125, no. 15, p. 154104, 2006.

- [22] M. Muñoz and C. Cárdenas, “How predictive could alchemical derivatives be?,” *Physical Chemistry Chemical Physics*, vol. 19, no. 24, pp. 16003–16012, 2017.
- [23] G. F. von Rudorff and O. A. von Lilienfeld, “Alchemical perturbation density functional theory,” *Physical Review Research*, vol. 2, no. 2, p. 023220, 2020.
- [24] G. F. von Rudorff and O. A. von Lilienfeld, “Rapid and accurate molecular deprotonation energies from quantum alchemy,” *Physical Chemistry Chemical Physics*, vol. 22, no. 19, pp. 10519–10525, 2020.
- [25] T. P. Straatsma and J. A. McCammon, “Computational Alchemy,” *Annual Review of Physical Chemistry*, vol. 43, no. 1, pp. 407–435, 1992.
- [26] K. Saravanan, J. R. Kitchin, O. A. von Lilienfeld, and J. A. Keith, “Alchemical Predictions for Computational Catalysis: Potential and Limitations,” *The Journal of Physical Chemistry Letters*, vol. 8, no. 20, pp. 5002–5007, 2017.
- [27] D. Sheppard, G. Henkelman, and O. A. von Lilienfeld, “Alchemical derivatives of reaction energetics,” *The Journal of Chemical Physics*, vol. 133, no. 8, p. 084104, 2010.
- [28] S. Chatterjee, C. Griego, J. L. Hart, Y. Li, M. L. Taheri, J. Keith, and J. D. Snyder, “Free standing nanoporous palladium alloys as co poisoning tolerant electrocatalysts for the electrochemical reduction of co₂ to formate,” *ACS Catalysis*, vol. 9, no. 6, pp. 5290–5301, 2019.
- [29] G. Henkelman, B. P. Uberuaga, and H. Jónsson, “A climbing image nudged elastic band method for finding saddle points and minimum energy paths,” *The Journal of Chemical Physics*, vol. 113, no. 22, pp. 9901–9904, 2000.
- [30] F. Abild-Pedersen, J. Greeley, F. Studt, J. Rossmeisl, T. R. Munter, P. G. Moses, E. Skúlason, T. Bligaard, and J. K. Nørskov, “Scaling Properties of Adsorption Energies for Hydrogen-Containing Molecules on Transition-Metal Surfaces,” *Physical Review Letters*, vol. 99, no. 1, p. 016105, 2007.
- [31] A. Vojvodic, F. Calle-Vallejo, W. Guo, S. Wang, A. Toftelund, F. Studt, J. I. Martínez, J. Shen, I. C. Man, J. Rossmeisl, T. Bligaard, J. K. Nørskov, and F. Abild-Pedersen, “On the behavior of Brønsted-Evans-Polanyi relations for transition metal oxides,” *The Journal of Chemical Physics*, vol. 134, no. 24, p. 244509, 2011.

- [32] S. Wang, B. Temel, J. Shen, G. Jones, L. C. Grabow, F. Studt, T. Bligaard, F. Abild-Pedersen, C. H. Christensen, and J. K. Nørskov, “Universal Brønsted-Evans-Polanyi Relations for C–C, C–O, C–N, N–O, N–N, and O–O Dissociation Reactions,” *Catalysis Letters*, vol. 141, no. 3, pp. 370–373, 2011.
- [33] B. Hammer and J. K. Nørskov, “Why gold is the noblest of all the metals,” *Nature*, vol. 376, no. 6537, pp. 238–240, 1995.
- [34] H. Xin and S. Linic, “Communications: Exceptions to the d-band model of chemisorption on metal surfaces: The dominant role of repulsion between adsorbate states and metal d-states,” *J. Chem. Phys.*, vol. 132, no. 22, p. 221101, 2010.
- [35] A. Vojvodic, A. Hellman, C. Ruberto, and B. I. Lundqvist, “From electronic structure to catalytic activity: a single descriptor for adsorption and reactivity on transition-metal carbides,” *Phys. Rev. Lett.*, vol. 103, no. 14, p. 146103, 2009.
- [36] S. Bhattacharjee, U. V. Waghmare, and S. C. Lee, “An improved d-band model of the catalytic activity of magnetic transition metal surfaces,” *Sci. Rep.*, vol. 6, p. 35916, 2016.
- [37] L. Yu, L. Vilella, and F. Abild-Pedersen, “Generic approach to access barriers in dehydrogenation reactions,” *Comm. Chem.*, vol. 1, no. 1, 2018.
- [38] M. to Baben, J. O. Achenbach, and O. A. von Lilienfeld, “Guiding ab initio calculations by alchemical derivatives,” *The Journal of Chemical Physics*, vol. 144, no. 10, p. 104103, 2016.
- [39] O. Anatole von Lilienfeld, “Accurate ab initio energy gradients in chemical compound space,” *The Journal of Chemical Physics*, vol. 131, no. 16, p. 164102, 2009.
- [40] R. P. Feynman, “Forces in Molecules,” *Phys. Rev.*, vol. 56, no. 4, pp. 340–343, 1939.
- [41] G. Kresse and J. Furthmüller, “Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set,” *Physical Review B*, vol. 54, no. 16, pp. 11169–11186, 1996.
- [42] F. Weigend, C. Schrod, and R. Ahlrichs, “Atom distributions in binary atom clusters: A perturbational approach and its validation in a case study,” *The Journal of Chemical Physics*, vol. 121, no. 21, pp. 10380–10384, 2004.

- [43] A. Beste, R. J. Harrison, and T. Yanai, “Direct computation of general chemical energy differences: Application to ionization potentials, excitation, and bond energies,” *The Journal of Chemical Physics*, vol. 125, no. 7, p. 074101, 2006.
- [44] D. Alfè, M. J. Gillan, and G. D. Price, “Constraints on the composition of the Earth’s core from ab initio calculations,” *Nature*, vol. 405, no. 6783, pp. 172–175, 2000.
- [45] R. Balawender, M. A. Welearegay, M. Lesiuk, F. De Proft, and P. Geerlings, “Exploring Chemical Space with the Alchemical Derivatives,” *Journal of Chemical Theory and Computation*, vol. 9, no. 12, pp. 5327–5340, 2013.
- [46] F. Weigend, “Extending DFT-based genetic algorithms by atom-to-place re-assignment via perturbation theory: A systematic and unbiased approach to structures of mixed-metallic clusters,” *The Journal of Chemical Physics*, vol. 141, no. 13, p. 134103, 2014.
- [47] K. Y. S. Chang, S. Fias, R. Ramakrishnan, and O. A. von Lilienfeld, “Fast and accurate predictions of covalent bonds in chemical space,” *The Journal of Chemical Physics*, vol. 144, no. 17, p. 174110, 2016.
- [48] A. Solovyeva and O. A. von Lilienfeld, “Alchemical screening of ionic crystals,” *Physical Chemistry Chemical Physics*, vol. 18, no. 45, pp. 31078–31091, 2016.
- [49] S. Fias, F. Heidar-Zadeh, P. Geerlings, and P. W. Ayers, “Chemical transferability of functional groups follows from the nearsightedness of electronic matter,” *Proceedings of the National Academy of Sciences*, vol. 114, no. 44, pp. 11633–11638, 2017.
- [50] K. Y. S. Chang and O. A. von Lilienfeld, “Al_xGa_{1-x}As crystals with direct 2 eV band gaps from computational alchemy,” *Phys Rev Mater*, vol. 2, no. 7, 2018.
- [51] Y. S. Al-Hamdani, A. Michaelides, and O. A. von Lilienfeld, “Exploring dissociative water adsorption on isoelectronically BN doped graphene using alchemical derivatives,” *The Journal of Chemical Physics*, vol. 147, no. 16, p. 164113, 2017.
- [52] R. Balawender, M. Lesiuk, F. De Proft, and P. Geerlings, “Exploring Chemical Space with Alchemical Derivatives: BN -Simultaneous Substitution Patterns in C 60,” *Journal of Chemical Theory and Computation*, vol. 14, no. 2, pp. 1154–1168, 2018.
- [53] S. Fias, K. Y. S. Chang, and O. A. von Lilienfeld, “Alchemical Normal Modes Unify Chemical Space,” *J Phys Chem Lett*, vol. 10, no. 1, pp. 30–39, 2019.

- [54] C. D. Griego, K. Saravanan, and J. A. Keith, “Benchmarking Computational Alchemy for Carbide, Nitride, and Oxide Catalysts,” *Advanced Theory and Simulations*, vol. 2, no. 4, p. 1800142, 2019.
- [55] G. F. von Rudorff and O. A. von Lilienfeld, “Atoms in Molecules from Alchemical Perturbation Density Functional Theory,” *The Journal of Physical Chemistry B*, vol. 123, no. 47, pp. 10073–10082, 2019.
- [56] J. R. Kitchin, J. K. Nørskov, M. A. Barteau, and J. G. Chen, “Trends in the chemical properties of early transition metal carbide surfaces: A density functional study,” *Catalysis Today*, vol. 105, no. 1, pp. 66–73, 2005.
- [57] R. Michalsky, Y.-J. Zhang, A. J. Medford, and A. A. Peterson, “Departures from the Adsorption Energy Scaling Relations for Metal Carbide Catalysts,” *The Journal of Physical Chemistry C*, vol. 118, no. 24, pp. 13026–13034, 2014.
- [58] R. Michalsky, Y.-J. Zhang, and A. A. Peterson, “Trends in the Hydrogen Evolution Activity of Metal Carbide Catalysts,” *ACS Catalysis*, vol. 4, no. 5, pp. 1274–1278, 2014.
- [59] H. Abroshan, P. Bothra, S. Back, A. Kulkarni, J. K. Nørskov, and S. Siahrostami, “Ultrathin Cobalt Oxide Overlay Promotes Catalytic Activity of Cobalt Nitride for the Oxygen Reduction Reaction,” *The Journal of Physical Chemistry C*, vol. 122, no. 9, pp. 4783–4791, 2018.
- [60] Y. Abghoui and E. Skúlason, “Electrochemical synthesis of ammonia via Mars-van Krevelen mechanism on the (111) facets of group III–VII transition metal mononitrides,” *Catalysis Today*, vol. 286, pp. 78–84, 2017.
- [61] Y. Abghoui, A. L. Garden, V. F. Hlynsson, S. Bjorgvinsdottir, H. Olafsdottir, and E. Skúlason, “Enabling electrochemical reduction of nitrogen to ammonia at ambient conditions through rational catalyst design,” *Phys. Chem. Chem. Phys.*, vol. 17, no. 7, pp. 4909–4918, 2015.
- [62] Y. Abghoui, A. L. Garden, J. G. Howalt, T. Vegge, and E. Skúlason, “Electroreduction of N₂ to Ammonia at Ambient Conditions on Mononitrides of Zr, Nb, Cr, and V: A DFT Guide for Experiments,” *ACS Catalysis*, vol. 6, no. 2, pp. 635–646, 2016.

- [63] Y. Abghoui and E. Skúlason, “Hydrogen Evolution Reaction Catalyzed by Transition-Metal Nitrides,” *The Journal of Physical Chemistry C*, vol. 121, no. 43, pp. 24036–24045, 2017.
- [64] B. M. Pabon, J. I. Beltran, G. Sanchez-Santolino, I. Palacio, J. Lopez-Sanchez, J. Rubio-Zuazo, J. M. Rojo, P. Ferrer, A. Mascaraque, M. C. Munoz, M. Varela, G. R. Castro, and O. Rodriguez de la Fuente, “Formation of titanium monoxide (001) single-crystalline thin film induced by ion bombardment of titanium dioxide (110),” *Nat. Commun.*, vol. 6, p. 6147, 2015.
- [65] J. P. Perdew, K. Burke, and M. Ernzerhof, “Generalized Gradient Approximation Made Simple,” *Phys. Rev. Lett.*, vol. 77, no. 18, pp. 3865–3868, 1996.
- [66] J. J. Mortensen, L. B. Hansen, and K. W. Jacobsen, “Real-space grid implementation of the projector augmented wave method,” *Phys. Rev. B*, vol. 71, no. 3, p. 35109, 2005.
- [67] R. Zhang, L. Song, and Y. Wang, “Insight into the adsorption and dissociation of CH₄ on Pt(hkl) surfaces: A theoretical study,” *Applied Surface Science*, vol. 258, no. 18, pp. 7154–7160, 2012.
- [68] Y. Liu, T. Zhao, W. Ju, and S. Shi, “Materials discovery and design using machine learning,” *J Materiomics*, vol. 3, no. 3, pp. 159–177, 2017.
- [69] H. Chen, O. Engkvist, Y. Wang, M. Olivecrona, and T. Blaschke, “The rise of deep learning in drug discovery,” *Drug Discov Today*, vol. 23, no. 6, pp. 1241–1250, 2018.
- [70] O. A. von Lilienfeld, “Quantum Machine Learning in Chemical Compound Space,” *Angew Chem Int Ed*, vol. 57, no. 16, pp. 4164–4169, 2018.
- [71] J. G. Freeze, H. R. Kelly, and V. S. Batista, “Search for Catalysts by Inverse Design: Artificial Intelligence, Mountain Climbers, and Alchemists,” *Chem Rev*, vol. 119, no. 11, pp. 6595–6612, 2019.
- [72] J. Schmidt, M. R. G. Marques, S. Botti, and M. A. L. Marques, “Recent advances and applications of machine learning in solid-state materials science,” *Comput Mater*, vol. 5, no. 1, 2019.

- [73] G. Montavon, M. Rupp, V. Gobre, A. Vazquez-Mayagoitia, K. Hansen, A. Tkatchenko, K. R. Müller, and O. A. von Lilienfeld, “Machine learning of molecular electronic properties in chemical compound space,” *New J Phys*, vol. 15, no. 9, 2013.
- [74] M. Rupp, R. Ramakrishnan, and O. A. von Lilienfeld, “Machine Learning for Quantum Mechanical Properties of Atoms in Molecules,” *J Phys Chem Lett*, vol. 6, no. 16, pp. 3309–3313, 2015.
- [75] F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman, K. Burke, and K. R. Muller, “Bypassing the Kohn-Sham equations with machine learning,” *Nat Commun*, vol. 8, no. 1, p. 872, 2017.
- [76] F. A. Faber, A. Lindmaa, O. A. von Lilienfeld, and R. Armiento, “Machine Learning Energies of 2 Million Elpasolite (ABC₂SD₆) Crystals,” *Phys Rev Lett*, vol. 117, no. 13, p. 135502, 2016.
- [77] J. S. Smith, O. Isayev, and A. E. Roitberg, “ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost,” *Chem Sci*, vol. 8, no. 4, pp. 3192–3203, 2017.
- [78] S. Hashemifar, B. Neyshabur, A. A. Khan, and J. Xu, “Predicting protein-protein interactions through sequence-based deep learning,” *Bioinformatics*, vol. 34, no. 17, pp. i802–i810, 2018.
- [79] B. Meredig, A. Agrawal, S. Kirklin, J. E. Saal, J. W. Doak, A. Thompson, K. Zhang, A. Choudhary, and C. Wolverton, “Combinatorial screening for new materials in unconstrained composition space with machine learning,” *Phys Rev B*, vol. 89, no. 9, 2014.
- [80] J. R. Kitchin, “Machine learning in catalysis,” *Nat Catal*, vol. 1, no. 4, pp. 230–232, 2018.
- [81] R. Jinnouchi and R. Asahi, “Predicting Catalytic Activity of Nanoparticles by a DFT-Aided Machine-Learning Algorithm,” *J Phys Chem Lett*, vol. 8, no. 17, pp. 4279–4283, 2017.
- [82] B. Meyer, B. Sawatlon, S. Heinen, O. A. von Lilienfeld, and C. Corminboeuf, “Machine learning meets volcano plots: computational discovery of cross-coupling catalysts,” *Chem Sci*, vol. 9, no. 35, pp. 7069–7077, 2018.

- [83] C. D. Griego, J. R. Kitchin, and J. A. Keith, “Acceleration of catalyst discovery with easy, fast, and reproducible computational alchemy,” *International Journal of Quantum Chemistry*, vol. 121, no. 1, p. e26380, 2021.
- [84] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine Learning in {P}ython,” *J Mach Learn Res*, vol. 12, pp. 2825–2830, 2011.
- [85] F. Abild-Pedersen, J. Greeley, F. Studt, J. Rossmeisl, T. R. Munter, P. G. Moses, E. Skúlason, T. Bligaard, and J. K. Nørskov, “Scaling properties of adsorption energies for hydrogen-containing molecules on transition-metal surfaces,” *Phys. Rev. Lett.*, vol. 99, p. 016105, 2007.
- [86] J. Nørskov, T. Bligaard, A. Logadottir, S. Bahn, L. Hansen, M. Bollinger, H. Bengaard, B. Hammer, Z. Sljivancanin, M. Mavrikakis, Y. Xu, S. Dahl, and C. Jacobsen, “Universality in heterogeneous catalysis,” *Journal of Catalysis*, vol. 209, no. 2, pp. 275–278, 2002.
- [87] D. J. Schmidt, W. Chen, C. Wolverton, and W. F. Schneider, “Performance of cluster expansions of coverage-dependent adsorption of atomic oxygen on Pt(111),” *J. Chem. Theory Comput.*, vol. 8, p. 264, 2012.
- [88] T. Rehman, M. Jaipal, and A. Chatterjee, “A cluster expansion model for predicting activation barrier of atomic processes,” *J. Comput. Phys.*, vol. 243, p. 244, 2013.
- [89] S. Verma, T. Rehman, and A. Chatterjee, “A cluster expansion model for rate constants of surface diffusion processes on Ag, Al, Cu, Ni, Pd and Pt(100) surfaces,” *Surf. Sci.*, vol. 613, p. 114, 2013.
- [90] L. M. Herder, J. M. Bray, and W. F. Schneider, “Comparison of cluster expansion fitting algorithms for interactions at surfaces,” *Surf. Sci.*, vol. 640, p. 104, 2015.
- [91] A. Bajpai, K. Frey, and W. F. Schneider, “Binary approach to ternary cluster expansions: NO–O–vacancy system on Pt(111),” *J. Phys. Chem. C*, vol. 121, p. 7344, 2017.
- [92] J. R. Kitchin, “Machine learning in catalysis,” *Nature Catalysis*, vol. 1, no. 4, pp. 230–232, 2018.

- [93] P. Schlexer, J. Lamoureux, K. T. Winther, J. A. Garrido, V. Streibel, M. Zhao, M. Bajdich, F. Abild-Pedersen, and T. Bligaard, "Machine learning for computational heterogeneous catalysis," *ChemCatChem*, vol. 11, no. 16, pp. 3581–3601, 2019.
- [94] M. E. Gázquez and R. Yáñez, "Recent advances in knowledge discovery for heterogeneous catalysis using machine learning," *Catalysis Reviews*, vol. 63, no. 1, pp. 120–164, 2021.
- [95] E. A. Eikey, A. M. Maldonado, C. D. Griego, G. F. von Rudorff, and J. A. Keith, "Evaluating quantum alchemy of atoms with thermodynamic cycles: Beyond ground electronic states," *The Journal of Chemical Physics*, vol. 156, no. 6, p. 064106, 2022.
- [96] R. G. Parr and W. Yang, "Density-functional theory of the electronic structure of molecules," *Annual Review of Physical Chemistry*, vol. 46, no. 1, pp. 701–728, 1995. PMID: 24341393.
- [97] R. Balawender, M. A. Welearegay, M. Lesiuk, F. De Proft, and P. Geerlings, "Exploring chemical space with the alchemical derivatives," *J. Chem. Theory Comput.*, vol. 9, no. 12, pp. 5327–5340, 2013.
- [98] R. Balawender, M. Lesiuk, F. De Proft, and P. Geerlings, "Exploring chemical space with alchemical derivatives: BN-simultaneous substitution patterns in C₆₀," *J. Chem. Theory Comput.*, vol. 14, no. 2, pp. 1154–1168, 2018.
- [99] R. Balawender, M. Lesiuk, F. D. Proft, C. V. Alsenoy, and P. Geerlings, "Exploring chemical space with alchemical derivatives: Alchemical transformations of H through Ar and their ions as a proof of concept," *Phys. Chem. Chem. Phys.*, vol. 21, no. 43, pp. 23865–23879, 2019.
- [100] T. Tamayo-Mendoza, C. Kreisbeck, R. Lindh, and A. Aspuru-Guzik, "Automatic differentiation in quantum chemistry with applications to fully variational hartree-fock," *ACS Central Science*, vol. 4, no. 5, pp. 559–566, 2018. PMID: 29806002.
- [101] A. S. Abbott, B. Z. Abbott, J. M. Turney, and H. F. Schaefer, "Arbitrary-order derivatives of quantum chemical methods via automatic differentiation," *The Journal of Physical Chemistry Letters*, vol. 12, no. 12, pp. 3232–3239, 2021. PMID: 33764068.

- [102] G. Domenichini, G. F. von Rudorff, and O. A. von Lilienfeld, “Effects of perturbation order and basis set on alchemical predictions,” *The Journal of Chemical Physics*, vol. 153, no. 14, p. 144118, 2020.
- [103] G. F. von Rudorff and O. A. von Lilienfeld, “Simplifying inverse materials design problems for fixed lattices with alchemical chirality,” *Science Advances*, vol. 7, no. 21, p. eabf1173, 2021.
- [104] G. F. von Rudorff, “Arbitrarily accurate quantum alchemy,” *The Journal of Chemical Physics*, vol. 155, no. 22, p. 224103, 2021.
- [105] C. D. Griego, L. Zhao, K. Saravanan, and J. A. Keith, “Machine learning corrected alchemical perturbation density functional theory for catalysis applications,” *AIChE Journal*, vol. 66, no. 12, p. e17041, 2020.
- [106] L. Bellaiche and D. Vanderbilt, “Virtual crystal approximation revisited: Application to dielectric and piezoelectric properties of perovskites,” *Phys. Rev. B*, vol. 61, pp. 7877–7882, 2000.
- [107] J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, and J. Hutter, “Quickstep: Fast and accurate density functional calculations using a mixed gaussian and plane waves approach,” *Computer Physics Communications*, vol. 167, no. 2, pp. 103–128, 2005.
- [108] G. Lippert, J. Hutter, and M. Parrinello, “The Gaussian and augmented-plane-wave density functional method for ab initio molecular dynamics simulations,” *Theoretical Chemistry Accounts*, vol. 103, no. 2, pp. 124–140, 1999.
- [109] J. VandeVondele and J. Hutter, “Gaussian basis sets for accurate calculations on molecular systems in gas and condensed phases,” *The Journal of Chemical Physics*, vol. 127, no. 11, p. 114105, 2007.
- [110] S. Goedecker, M. Teter, and J. Hutter, “Separable dual-space gaussian pseudopotentials,” *Phys. Rev. B*, vol. 54, pp. 1703–1710, 1996.