

Resolving the Coverage Dependence of Surface Reaction Kinetics with Machine Learning and Automated Quantum Chemistry Workflows

Matthew S. Johnson,[†] David H. Bross,[‡] and Judit Zádor^{*,†}

*[†]Combustion Research Facility, Sandia National Laboratories, California 94551-0969,
United States*

*[‡]Chemical Sciences and Engineering Division, Argonne National Laboratory, Lemont,
Illinois 60439, United States*

E-mail: jzador@sandia.gov

Abstract

Microkinetic models for catalytic systems require estimation of many thermodynamic and kinetic parameters that can be calculated for isolated species and transition states using ab initio methods. However, the presence of nearby co-adsorbates on the surface can dramatically alter these thermodynamic and kinetic parameters causing them to be dependent on species coverage fractions. As there are combinatorially many co-adsorbed configurations on the surface, computing the coverage dependence of these parameters is far less straightforward.

We present a framework for generating and applying machine learning models to predict coverage-dependent parameters for microkinetic models. Our toolkit enables automatic calculation and evaluation of co-adsorbed configurations allowing us to sample 2,000 co-adsorbed adsorbates and transition states (TSs) for a diverse set of 9 reactions on Cu(111), a challenging surface, with four possible co-adsorbates. This dataset was then used to train subgraph isomorphic decision trees (SIDTs) to predict the stability and association energy of configurations. We were able to achieve mean absolute errors (MAEs) of 0.106 eV on adsorbates, 0.172 eV on TSs, and due to natural error cancellation in SIDTs for relative properties, 0.130 eV on reaction energies and 0.180 eV on activation barriers. We describe how to use these models to predict coverage-dependent corrections for adsorbates and TSs, and demonstrate on H*, HO* and O* comparing the generated SIDT model with an iteratively refined version.

Introduction

Heterogeneous catalysis plays an incredibly important role in energy technologies and chemical manufacturing. Catalytic systems involve many elementary reactions and are sensitive to temperature, pressure, and the nature of the catalyst. In order to predict the behavior of these systems at a range of conditions, usually one would build and simulate a microkinetic model (MKM) or a kinetic Monte Carlo (KMC) model.

Under low-coverage conditions it is often adequate to treat adsorbates and reactions as if they are isolated on the surface, assuming an otherwise empty surface. However, in many systems and under many conditions the surface has many adsorbates packed closely together. In these co-adsorbed systems the lateral interactions between co-adsorbed species can very significantly alter the thermochemistry and kinetics.¹⁻⁹ However, brute force computation of the minimum association energy of an adsorbate and co-adsorbates at a given coverage requires one to compute the average binding energy of every conceivable configuration at that coverage. Even just computing the minimum association energy of an adsorbate in the presence of a single type of co-adsorbate is computationally very expensive. Consideration of several co-adsorbates is even less computationally feasible, therefore, the effects are often either ignored or approximated crudely. Furthermore, due to the additional challenges most researchers do not attempt to compute the coverage dependence of transition state properties, and instead use Brønsted-Evans-Polanyi (BEP) relations.^{10,11}

To tackle these challenges a variety of approaches have been developed. In one approach, the dimensions of the periodic system are adjusted to achieve the desired coverage with the smallest system possible and reduce the computational cost.¹² Additionally, researchers have built models for predicting the energies of co-adsorbed configurations like the cluster expansion (CE) model¹³ that decomposes the overall energy into contributions from groups of N specific adsorbates to a given N .

However, these approaches have a number of weaknesses. Size reduced periodic cells inherently assume that the minimum energy coverage pattern is periodic in the size of the unit cell. In many cases lowest energy configurations cannot be captured with reduced size unit cells.¹⁴ In fact, such reduced size cell configurations may not even be minima on the potential energy surface of a larger cell. Consider a 1×1 cell with one O^* . The overall barrier for the reaction sequence $2 \text{ O}^* \rightarrow \text{O}_2^* \rightarrow \text{O}_2 + *$ can be comparable to predicted lateral interactions between two O^* .^{12,15} Given the highest barrier in this reaction sequence is usually not much higher in energy than $\text{O}_2 + *$, it is very believable that at high coverage

$2 \text{ O}^* \rightarrow \text{O}_2^*$ or $\text{O}_2 + *$ could occur barrierlessly spontaneously, with the 2 O^* state not being stable. However, in the 1×1 periodic cell the distance between O^* 's is fixed making the stable configuration with O 's bonded pairwise with each other unachievable, causing an optimization to find a configuration that is unstable on any realistically sized slab.

CE parameters are particular to the adsorbates and interaction terms they were computed for.¹⁶ Thus these models are primarily only useful for set interactions of adsorbates and co-adsorbates they have been explicitly fit to, and cannot be used to predict fit interactions with new adsorbates or co-adsorbates. Additionally, due to the combinatorial increase in possible clusters as cluster size increases it is rare to fit them out to more than clusters of $N = 3$. Therefore, in practice this method is likely less predictive at higher coverages.

Machine learning techniques, in principle, do not have the limitations of cluster expansion methods. They can learn from diverse datasets and predict on interactions not included in the training data. However, machine-learning techniques based on popular deep neural networks (DNNs) require large amounts of training data that can be computationally very expensive to obtain. Additionally, DNNs are not interpretable, making them difficult to analyze and improve.

The subgraph isomorphic decision tree (SIDT) machine learning method provides an alternative that is free of the above weaknesses.¹⁷⁻¹⁹ SIDTs are made up of nodes associated with molecular substructures represented as molecular subgraphs. They are evaluated by descending a target graph structure down the tree to the nodes with subgraphs it matches until it reaches the most specific matching node and making a prediction based on either the nodes matched or the final node. SIDTs can be applied to datasets too small for DNNs, they are straightforward to extend and retrain, and the substructures in the tree are inherently visualizable making SIDTs easy to analyze and interpret, and thus modify the generation process and tree itself to achieve desired outcomes. Additionally, SIDTs have a property that enables unique inherent error cancellation on many important chemical problems involving differences of physical quantities (e.g., barrier heights, reactions energies), in contrast

to DNNs that do not have this advantage. In kinetics, rate coefficients are not dependent on the absolute energy of any given configuration, only on the energy differences between reactants and the transition state and reactants and products, which are used to compute the forward and reverse rate coefficients for a given reaction. When the SIDT predicts the association energy of a reactant configuration and that of a transition state or product configuration, most of the matched subgraphs associated with interactions between adsorbates are unchanged, because only interactions close to the reaction center are different in these structures. As a result, predictions on unchanged interactions will cancel exactly in the SIDT and our error is only associated with the subgraphs that are modified by the reaction, i.e., bond breaking and forming. This property makes SIDTs significantly more accurate on the relative properties that actually matter for kinetics than one would expect from a given level of absolute accuracy.

In this work we perform an initial exploration applying SIDT techniques combined with automated quantum chemistry to estimate coverage dependence across a wide chemical space. We present a framework for automatically computing co-adsorbed configurations and using machine learning to predict the coverage dependent energetics of adsorbates and transition states and thus the coverage dependence of rate coefficients and thermochemistry. We automatically generate a training dataset of co-adsorbed adsorbates and transition states on a $3 \times 3 \times 4$ Cu(111) slab, a challenging surface. The size of the slab is chosen to minimize interaction between periodic images of the species. We use the dataset to train a sequence of SIDTs to predict whether a configuration is stable or not and to predict the association energy of the configuration for both co-adsorbed adsorbates and transition states. We show the effectiveness of the SIDT predictors and demonstrate the error-cancellation property discussed above on single co-adsorbate coverage dependence of adsorbates and TSs considered in the training set. We also show how to go from association energy predictions to coverage-dependent rate coefficients and reaction energies.

Methods

Dataset Generation

We started our co-adsorbed calculations from a set of calculations for isolated reactions, i.e., in the absence of co-adsorbates. We took the lowest energy configurations for 12 adsorbates and 9 transition states on Cu(111) calculated by Johnson et al.²⁰ using our software, Pynta. This set, listed in Table 1, includes transition states for a range of different reaction classes, adsorbates consisting of H, C, O, and N atoms and one bidentate adsorbate. In this work we selected four adsorbates to be co-adsorbates for purposes of sampling: H, N, and O atoms, and HO.

Table 1: Reactions on Cu(111) considered in this work for lateral interactions.

Reaction
$H^* + O^* \longleftrightarrow HO^* + *$
$OCH^* + * \longleftrightarrow OC^* + H^*$
$H^* + * \longleftrightarrow H^* + *$
$OC^* + O^* \longleftrightarrow CO_2 + 2*$
$HO^* + H^* \longleftrightarrow H_2O + 2*$
$HOCH_2^* + * \longleftrightarrow CH_3O^* + *$
$OCHO^* + * \longleftrightarrow CO_2 + H^* + *$
$H^* + OCH_2O^{**} \longleftrightarrow HOCH_2O^* + 2*$
$N^* + CH_3O^* \longleftrightarrow CH_2O + NH^* + *$

All of our calculations were done using the same software, methods and parameters as in Johnson et al. We used the BEEF-vdW functional²¹ with PBE-KJPAW pseudopotentials and an energy cutoff of 40 Ry as implemented in Quantum Espresso^{22,23} for a $3 \times 3 \times 4$ Cu(111) slab with a $3 \times 3 \times 1$ k -point grid. All geometry optimizations targeting minima were done in two stages: first using the MDMin method implemented in the Atomic Simulation Environment (ASE) until $f_{\max} \leq 0.5$ eV/A and then using ASE's BFGSLineSearch algorithm until $f_{\max} \leq 0.02$ eV/A.²⁴ Saddle point optimizations and intrinsic reaction coordinate (IRC) calculations were done using Sella until $f_{\max} \leq 0.02$ eV/A and $f_{\max} \leq 0.1$ eV/A respectively.²⁵⁻²⁷ Vibrational calculations were run using ASE's vibrations module.²⁴

Unless otherwise specified, all placements for optimization on the surface were done by starting from the isolated configuration for the adsorbate or transition state and then placing co-adsorbates on the selected sites using Pynta’s placement algorithm.²⁰ 3D configurations were converted to 2D graphs automatically using a set of algorithms contained within Pynta. These 2D representations are key to employing SIDT. For adsorbates, this process is relatively straightforward. ASE’s analysis tool was used to identify covalent bonds within the adsorbates; possible surface bonds are identified by searching for the closest site within 2.5 Å of the relevant adsorbate atom, and only considering sites on which the associated adsorbate is stable under isolated conditions. We then completed the 2D description (the graph of the system) by incrementing bond multiplicities to satisfy the octet rule, with surface bonds having the lowest priority.

For TSs we considered a number of additional factors. Reaction bonds, i.e., bonds that break or form in the reaction, were identified based on the original reaction template associated with the TS. Since atoms that are part of the reaction center may be close to the surface without being properly associated with a site, we included both reactant and product sites in the set of valid sites to determine the 2D structure. Since the reaction bonds in the 2D representation do not have a well-defined order, we cannot always satisfy the octet rule when we increment bonds. As a result, we can end up with extra bonds to the surface that are artifacts of the process. To handle this we removed surface bonds if they were a single bond and the associated atom already had two reaction bonds. This covers all reactions considered in this work and likely most common cases as well.

TSs for diffusions and similar reactions, where one atom has two reaction bonds to different sites, pose a further challenge for generating 2D representations, because we cannot simply form the reaction bonds with the closest stable site. Instead, we seek a pair of sites that capture the origin and destination of the diffusion process. To identify the right pair of sites, we considered five criteria. First, we required that the normalized vector connecting the two sites $\mathbf{v}_{\text{sites}}$ and the normalized normal mode vector associated with the imaginary

frequency $\mathbf{v}_{\text{imagfreq}}$ are closely aligned

$$|\langle \mathbf{v}_{\text{sites}} | \mathbf{v}_{\text{imagfreq}} \rangle| \geq 0.95 \quad (1)$$

Second, we required that the distance between the sites is less than 3 Å. Third, we required that the distance between the binding atom of the adsorbate and any site, d_{site} be less than 2 Å. Fourth, we only considered sites where the reactants or products were stable in the isolated calculations. Finally, we also defined a measure, h , of how closely the atom is positioned to the halfway point between the sites

$$h = \frac{||\mathbf{u}_1 + \mathbf{u}_2||}{||\mathbf{u}_1|| + ||\mathbf{u}_2||} \quad (2)$$

where \mathbf{u}_i is the vector from site i to the atom. When more than one pair of sites fulfilled the previous criteria, we chose the pair that maximizes $1/(d_{\text{site}} \times h)$.

The techniques discussed above are sufficient to generate the 2D representations of TSs, including ones of diffusion reactions. However, for TSs we also need to validate that the saddle point optimized in the presence of the co-adsorbates still connects the reactants and products of the original reaction. We found that two criteria were sufficient to separate correct from incorrect TSs in a subset of our data consisting of 100 randomly selected TSs, which were checked manually. We first required that the normalized normal mode corresponding to the imaginary frequency in the co-adsorbed case, \mathbf{v}_{coad} , aligns with the normalized isolated mode $\mathbf{v}_{\text{isolated}}$

$$|\langle \mathbf{v}_{\text{isolated}} | \mathbf{v}_{\text{coad}} \rangle| \geq 0.7 \quad (3)$$

We additionally required that any co-adsorbate atoms be more than 1.1 times the covalent bond cutoff threshold (as defined in ASE) away from any reaction center atoms.

One major challenge of sampling the co-adsorbed space is that many configurations one might propose are not stable. Lateral interactions may prevent two adsorbates from being

placed next to each other, surface restructuring may affect the stability of old sites or create new ones, and reactions that normally have a barrier may occur spontaneously at higher coverages. We simultaneously mitigated these challenges and provided a useful base set of samples by first calculating every unique and valid placement pairing between every adsorbate and every selected co-adsorbate that put them within a maximum distance of 3 Å. We analyzed the results comparing the initial 2D graph and the 2D graph after optimization to determine which pairwise configurations are stable.

Finally, to generate a random, yet balanced set, we first chose either to sample an adsorbate or a TS with equal probability. We then chose a single random co-adsorbate at 90% probability and a random sample of mixed co-adsorbates at 10% probability. To determine the sampled coverage fraction we drew a uniformly distributed sample from [0,1], which was then rounded to an integer number of co-adsorbates to put on the slab. The co-adsorbates were distributed randomly over the stable sites. Samples that involved sub-configurations that matched a failed pairwise optimization were rejected. For all unique successful TS optimizations we ran an IRC to find the reactant and product configurations and optimized and computed frequencies for both. This enabled explicit sampling of activation barriers and reaction energies.

Drawing and running calculations for 2,000 samples gave us 477 unique and valid co-adsorbed adsorbate configurations and 207 unique and valid co-adsorbed TS configurations. Figure 1 shows the energy of these configurations as a function of the number of surface bonds, while Figure 2 shows the distribution of the elements bonded to the surface across the whole dataset, i.e., including the adsorbing atoms of the central adsorbate or TS and of the co-adsorbates. In general, H atoms tend to have weaker inter-adsorbate interactions owing to its smaller size, making co-adsorbed configurations with H more likely to be stable and optimize successfully, causing H to have a large number of occurrences in the dataset. The similar number of O occurrences is unsurprising, given O and HO are co-adsorbates and many of the adsorbate and TSs targeted involve bonds between O and the surface. Given

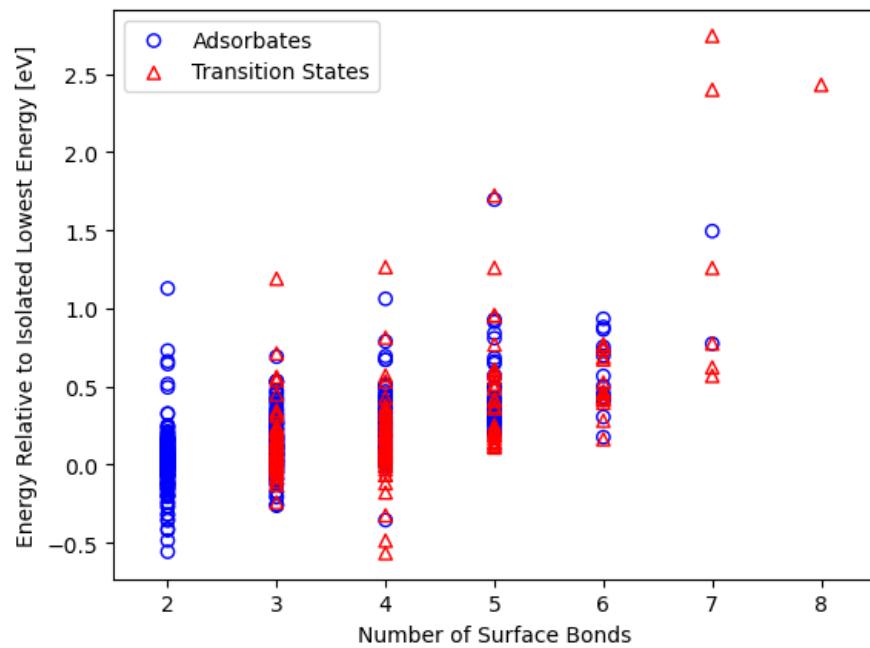


Figure 1: Relative energies of the successful configurations in the dataset arranged by the number of surface bonds.

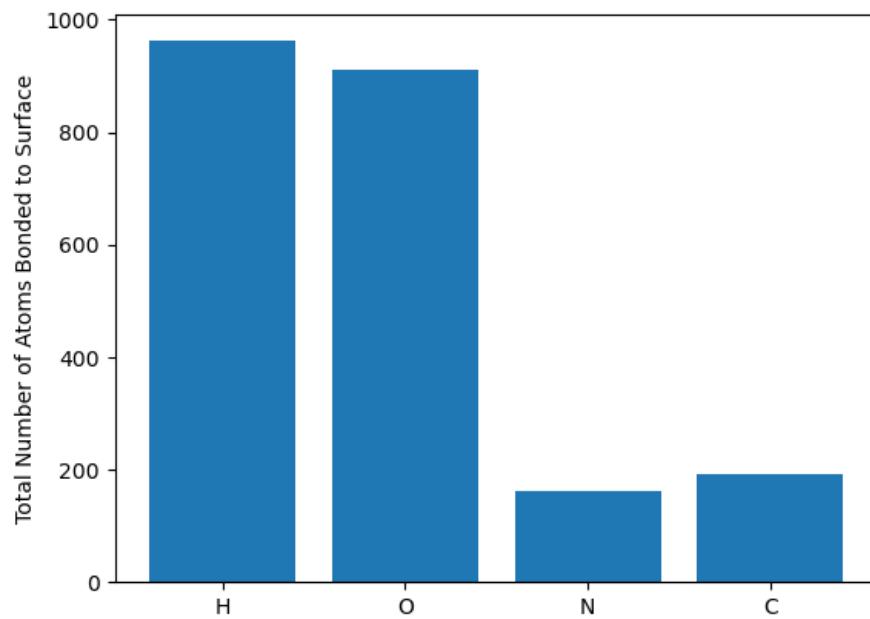


Figure 2: Number of atoms bonded to the surface in dataset configurations by element.

none of the co-adsorbates involve a bond between the surface and C, the significantly lower number of occurrences is unsurprising. However, the low number of N occurrences is much more unexpected. N does not occur in many of the isolated adsorbate and TSs considered (see Table 1), but it is a co-adsorbate, so one would expect it to exist in at least a quarter of proposed sample configurations. This discrepancy occurs because N is particularly prone toward causing surface restructuring on Cu.²⁸ In general, this makes it more difficult to converge geometries with N on the surface resulting in lower optimization success rates than O and HO on this surface.

In Figure 3 we examine the distribution of the differences between the activation barrier in co-adsorbed configurations and the corresponding isolated configuration for configurations in the dataset. Differences can be quite large in both negative and positive directions and range from approximately -0.8 to 1.0 eV. The center of the distribution for individual reactions (Figure 3b) appears to differ significantly from zero in at least a few cases. The range for individual reactions seems to always span at least about 0.5 eV, and $\text{HO}^* \rightarrow \text{H}^* + \text{O}^*$ spans approximately 1.5 eV.

Figure 4 shows all of the successfully optimized and accepted co-adsorbed transition states for the $\text{HO}^* + \text{H}^* \rightarrow \text{H}_2\text{O} + 2^*$ reaction and the associated activation barriers relative to that of the isolated reaction. Significant surface restructuring can be observed in several configurations, e.g., first and last image in the second row. Notably, the geometry of the atoms involved in the transition state varies significantly between co-adsorbed configurations. While larger geometry shifts can occur for adsorbates as well, we found that it is more common and pronounced more for transition states, making their energetics more challenging to learn.

Machine Learning

Our goal is to predict the change in energetics of stable and TS surface configurations relative to their isolated energetics due to the presence of co-adsorbates. We divided the prediction

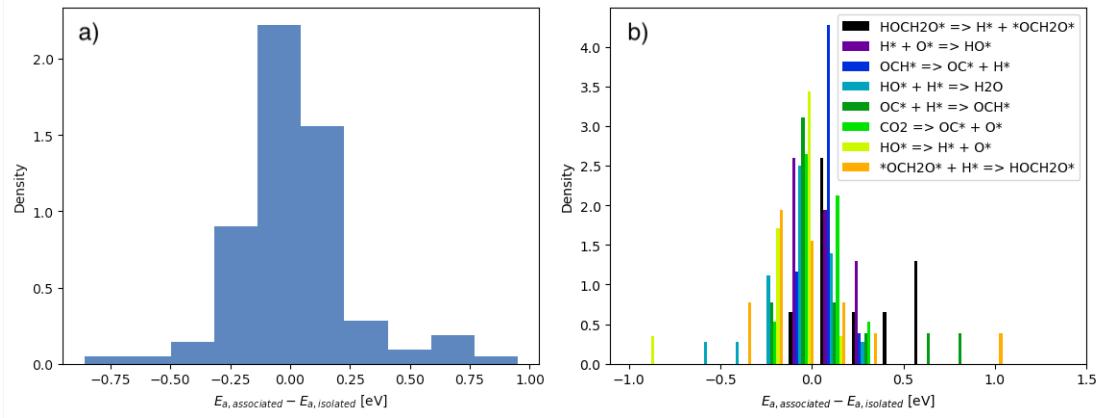


Figure 3: (a) Distribution of differences between the activation barrier of co-adsorbed configurations and the activation barrier of the corresponding isolated reaction, and (b) the same data shown for each reaction separately.

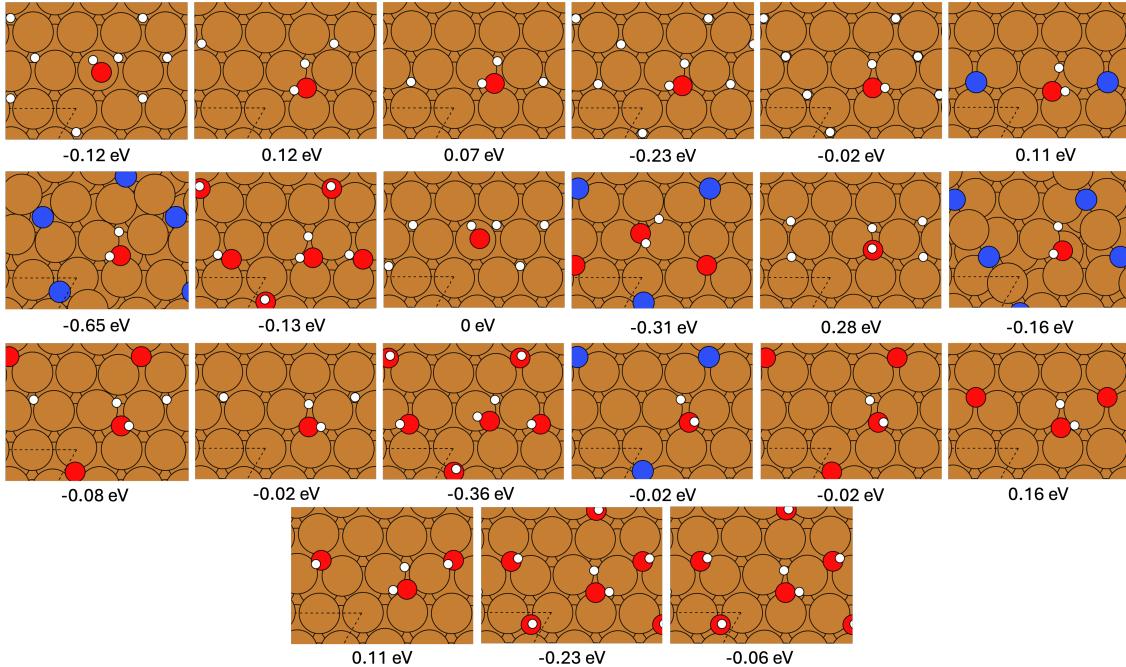


Figure 4: Final successful transition state configurations for the $\text{HO}^* + \text{H}^* \rightarrow \text{H}_2\text{O} + 2^*$ reaction. Values are the difference in energy between the co-adsorbed activation barrier and the isolated activation barrier. Images are periodic extensions of a 3x3 slab. White circles are H, red circles are O, blue circles are N, and brown circles are Cu atoms.

process into two primary steps: (i) decide whether a proposed surface configuration is stable and if so (ii) predict the association energy of the configuration. For the former we use multi-evaluation SIDT binary classifiers, and for the latter we use multi-evaluation SIDT regressors as implemented in our software, PySIDT.¹⁹ Both of these SIDT algorithms decompose a 2D graph representation of the chemical configuration into a set of chemical substructures contributing to the prediction. In the case of the classifier, each substructure is predicted to either be locally stable or unstable, and in this application if any prediction is unstable the configuration is classified as unstable. In the case of the regressor, an energetic contribution is predicted from each substructure and summing across all substructures gives the associated energy.

For our 2D representation we represented the periodic slab in its entirety resolving each site and adsorbate atom as nodes. Edges were included between all covalently bonded atoms, each site and adsorbate atom bound to that site, and between neighboring sites as defined by the ACAT software.²⁹ We used RMG’s cheminformatics engine and molecular representation software³⁰ within PySIDT¹⁹ for all operations on this representation.

When selecting decompositions of the 2D graphs into subgraphs there are two primary considerations: (i) we would like the set of decompositions to involve minimum redundancy in chemical information and (ii) we would like the decompositions to locally include all chemical information we need to make the prediction. The simplest decomposition one might think of is to look locally at each individual atom and site in the configuration. However, we are examining inter-adsorbate interactions so we would not necessarily expect every atom or site to have a significant unique contribution to the energetics. Looking at surface bonds should significantly reduce redundancy, however, since we study interactions between adsorbates that may be far apart in the 2D representation, we need to consider at least pairs of surface bonds. While pairs of surface bonds are sufficient to encode the interactions between co-adsorbates we are interested in resolving, we have allowed co-adsorbates to be adsorbed at sites that do not correspond to their lowest energy isolated configuration. (Note that most

studies only consider the isolated lowest energy sites). Therefore, we also need to predict the energetics of moving adsorbates between sites. For this reason our overall representation included both surface bonds alone (to account for energetics of adsorbates binding to individual sites) and pairs of surface bonds (to account for the lateral interactions). We did test including triad-wise (i.e., three-body) interactions, however, it did not improve performance in this work. It should be noted that SIDT does not need triad-wise decompositions to learn triad-wise interactions since the pair-wise decompositions can be grown to resolve more than two adsorbates in the learning processes.

To simplify and compartmentalize the training processes, we first train trees to predict the single surface-bond interactions on isolated data from Johnson et al.,²⁰ including not just the lowest energy, but all valid isolated structures. We then train a second tree on the co-adsorbed datasets generated in this work to add the contribution from the interactions of pairs of surface-bonds. For the classifier we only train the second tree on configurations the single-surface-bond tree predicts to be stable, and for the regressor we do delta learning, subtracting the single-surface-bond-tree prediction from the dataset before training. This architecture is shown in Figure 5.

Results

Stability Predictions for Co-Adsorbed Systems

Our dataset provided stability labels for 1720 unique and otherwise valid configurations. The configurations included unstable configurations that were proposed but then failed to preserve the original graph representation upon optimization, stable configurations that were proposed and successfully optimized, and stable configurations that were found in failed optimizations using DFT. Of these configurations, 91 were predicted to be unstable by the single surface-bond classifier, 75 of which incorrectly predicted to be unstable and 16 of which were correctly predicted to be unstable. The pair surface-bonds classifier was trained

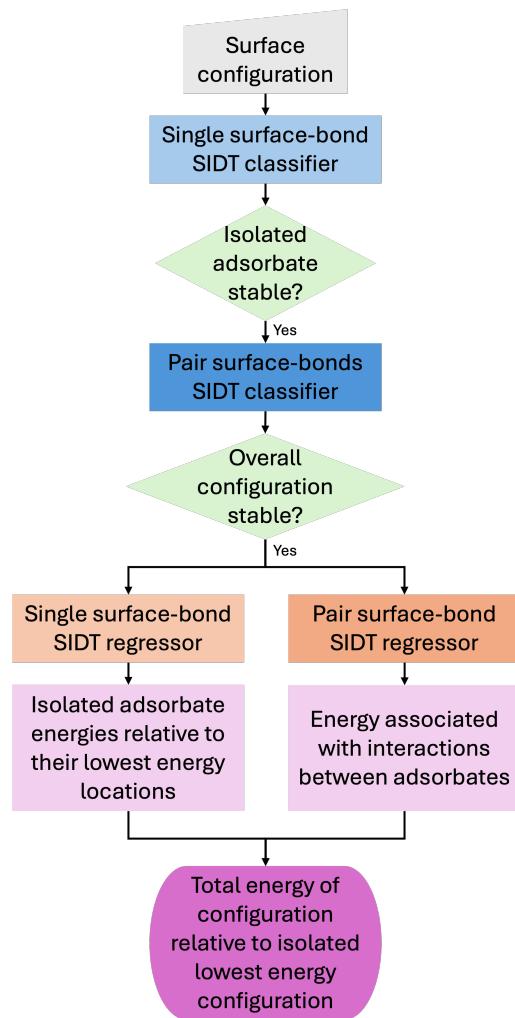


Figure 5: Diagram for computing the stability and energy of a given input surface configuration.

using a 8:1:1 train:validation:test split on the other 1629 configurations. Each iteration the SIDT computed the accuracy on the validation set and at the end of the run it reverted to the tree from the iteration with the best validation accuracy. The confusion matrix for the test set of the pair-surface-bonds classifier is available in Table 2. This implies an accuracy (fraction of classifications that are right) of 86% and a precision (fraction of predicted Trues that are correct) of 84%. Given the inherent challenge of determining the substructures that underlie all of the many different kinds of stability in these systems, this is good performance from a machine-learning perspective. From a more practical and problem oriented perspective, focusing on our ability to predict the association energies of the lower energy configurations, we do not need extremely high levels of accuracy for the stability classifier, because high association energy predictions and instability are inherently correlated. A co-adsorbed configuration that is unstable and cannot be occupied and a co-adsorbed configuration that is high in energy relative to other configurations at identical coverage and thus has very low occupation are both unimportant from a thermodynamic and kinetic perspective.

Table 2: Test set confusion matrix for stability of interactions between co-adsorbates.

	Predict True	Predict False
Value True	91	9
Value False	14	49

Association Energy Predictions for Co-Adsorbed Systems

For association energy prediction, as noted earlier, we only have 477(minima)+207(TS)= 684 valid configurations. For this smaller dataset we trained on the full dataset and computed leave-one-out errors for every training point. We first trained the tree out to 152 nodes only on configurations with two co-adsorbates and no TSs before training out to 475 nodes on the full dataset.

The resulting comparison plot is available in Figure 6a and the associated uncertainty calibration plot is available in Figure 6b. For a more detailed view, Figures 6c and 6d show the parity plot for the adsorbates and the TSs separately. The error analysis gives an overall MAE of 0.126 eV, a MAE on adsorbates of 0.106 eV and a MAE on TSs of 0.172 eV. The uncertainty calibration shows that the model is slightly underconfident at small confidence intervals and slightly overconfident at large confidence intervals, but in general, the model uncertainties appear to be a good representation of actual uncertainties. This is especially encouraging given that it is challenging to predict accurate uncertainties for some of the more unusual configurations in this dataset, especially for TSs.

Activation Barrier and Reaction Energy Correction Predictions for Co-Adsorbed Systems

While lateral interactions are often discussed in terms of the energies of specific configurations, this is not the most relevant quantity for kinetics, which, as discussed earlier, are only sensitive to relative configuration energies: activation barriers and reaction energies rather than the absolute energies. Our dataset offers a unique opportunity, allowing us to look directly at the relevant properties using the optimized endpoint configurations from the IRCs for each unique TS. We present parity plots for activation barriers and reaction energies in Figures 7a and 7b respectively.

We achieve an MAE of 0.180 eV for activation barriers and 0.130 eV for reaction energy.

Analysis of High Error Configurations

Figure 8 shows nine representative samples from the 20 configurations with the highest leave-one-out association energy errors. It is worth noting that the the high error configurations are dominated by transition states (a, b, c, e, f, g, h). In general, several different factors are associated with poor prediction: surface restructuring (c, e, g, h, i), transition states with co-adsorbates close to the reaction center (a,b,c,f), adsorbates optimizing to different

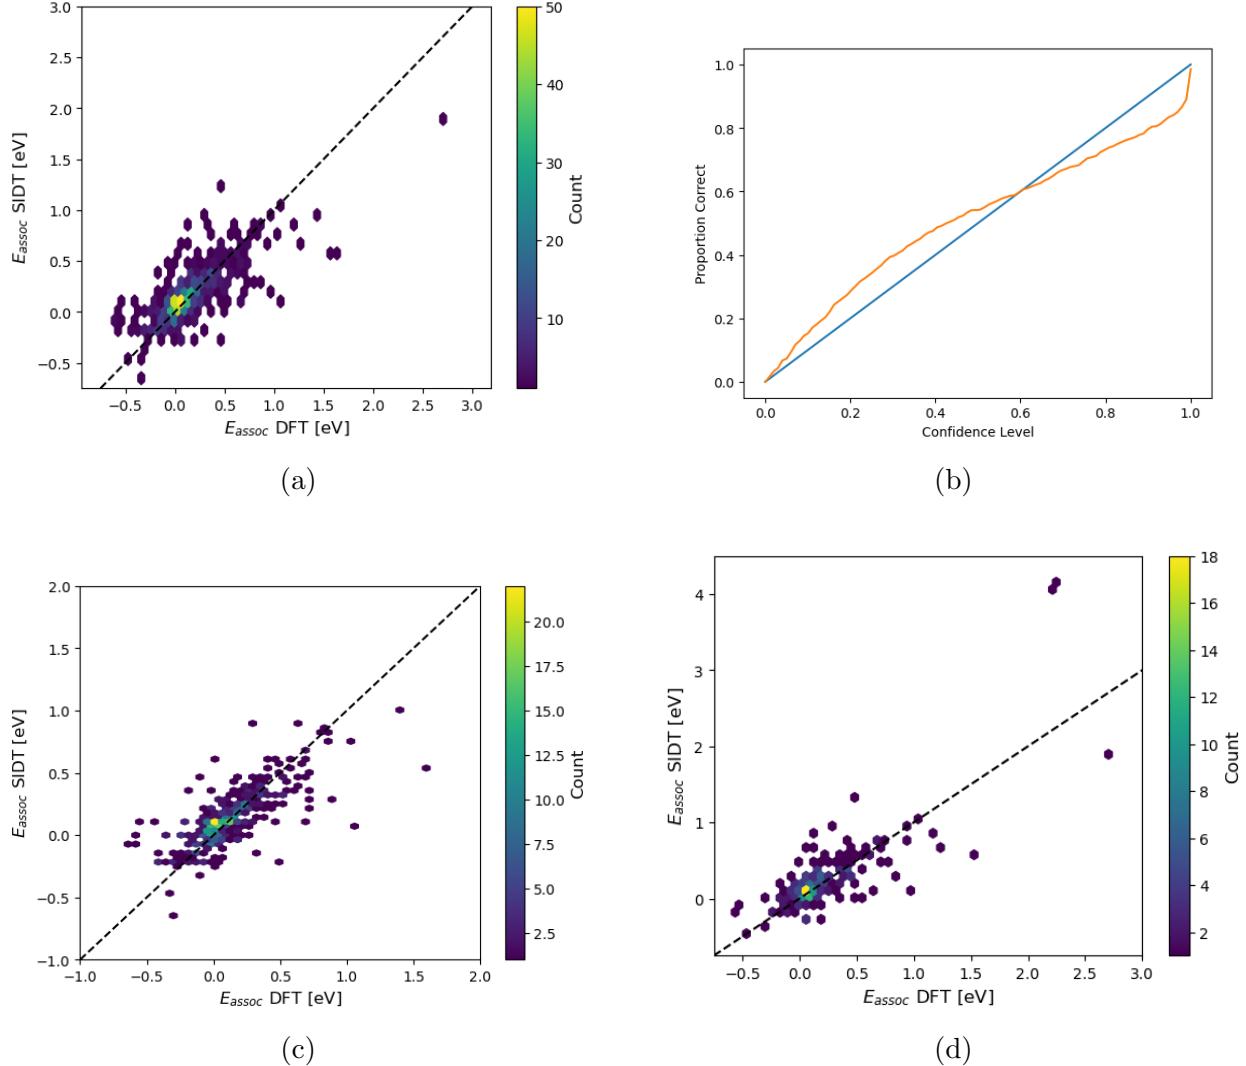


Figure 6: SIDT performance on absolute energies. (a) Parity plot for leave-one-out errors in association energy for all configurations. (b) Uncertainty calibration plot for leave-one-out and estimated errors in association energy for all configurations. (c) Parity plot for leave-one-out errors in association energy for adsorbate configurations. (d) Parity plot for leave-one-out errors in association energy for TS configurations.

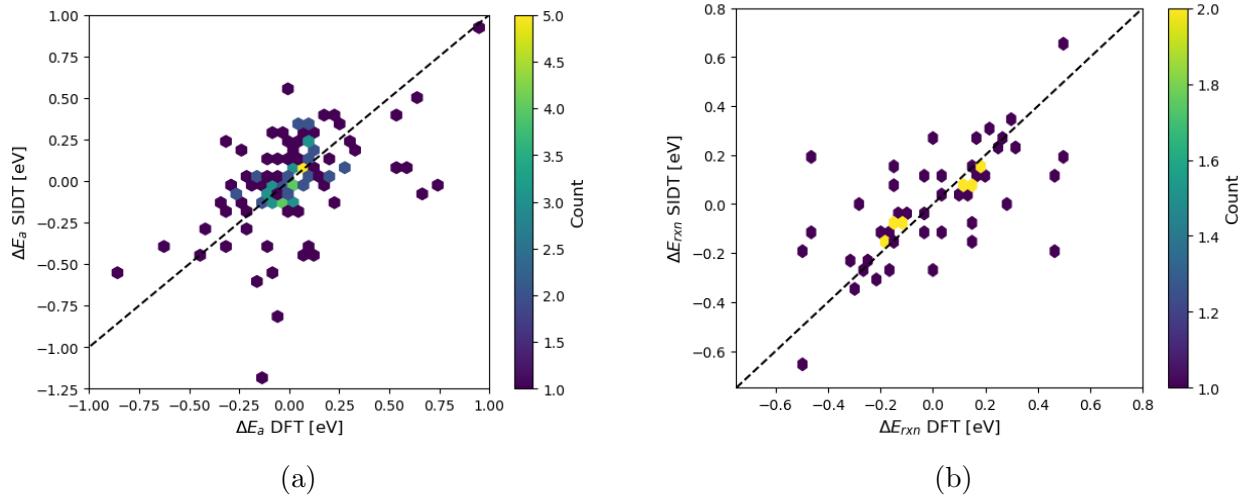


Figure 7: SIDT performance on relative energies. Parity plot for leave-one-out errors in activation barrier corrections (a) and in reaction energy corrections (b).

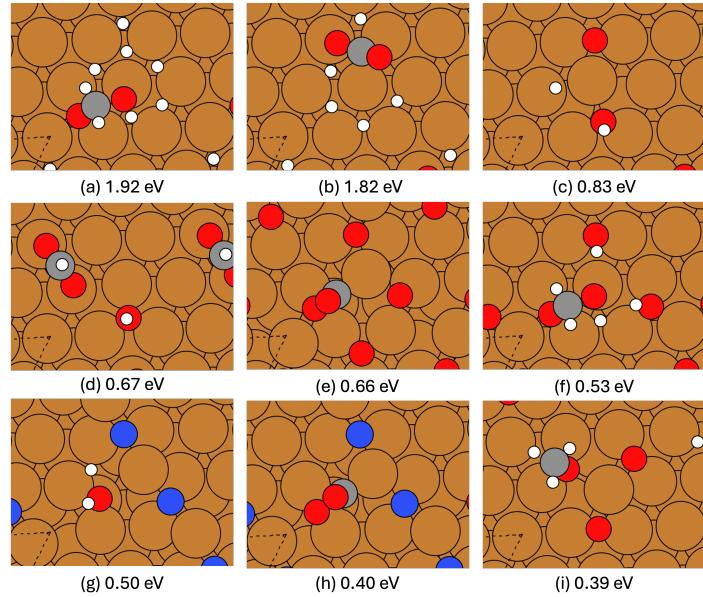


Figure 8: Selected samples from the 20 configurations with the worst leave-one-out association energy errors. Values are the difference in energy between the predicted and actual association energy. Subfigures a, b, c, e, f, g, and h are transition states while d and i are minima. Subfigure b has a hydrogen underneath the carbon atom.

conformers and/or site configurations (d), and hydrogen bonding (f). Figure 8c, a transition state for H^* diffusion, is particularly interesting because the changes in the transition state are less subtle than those for more complex reactions. Even without the involved surface restructuring, H^* is not stable at a hollow adjacent to O^* on a hollow. As a result, while the isolated diffusion reaction occurs between two adjacent hollows, this particular co-adsorbed diffusion reaction instead has to hop over the unstable hollow causing it to be between two non-adjacent hollows dramatically changing the barrier.

Most of these interactions are so specific or intricate that the generated dataset is unlikely to have enough similar samples to learn them. The 2D representation used in this work is not inherently aware of surface restructuring, conformers, or hydrogen bonding. In principle, such additional dimensions can be integrated into the 2D representation, but these interactions are likely to require significantly more data to learn properly. Note also that most studies of coverage dependence would likely have rejected these kinds of samples.

Interpreting Overall Errors

It is important to understand that the overall errors shown in this section represent errors across samples from the full space of co-adsorbed configurations. However, the accuracy of a rate coefficient prediction is determined by our ability to predict the lowest energy configurations of the relevant species at a given coverage. These lowest energy configurations tend to involve better spaced co-adsorbates with weaker and less complex interactions that are in general easier to predict than randomly sampled configurations. However, in order to know the lowest energy configuration at a given coverage, one needs to know the energies of all configurations of a given adsorbate/TS-co-adsorbate pair at that coverage. This makes it challenging to gauge the true error in the activation barrier and reaction energy one would experience.

Later in this work we will construct approximate baselines for a few cases to help gauge errors in predicting the lowest energy configurations.

Error Cancellation

Noting the MAEs of 0.106 eV and 0.172 eV for adsorbate and TS energy predictions, the low MAEs of the relative quantities suggest significant error cancellation. Assuming no correlation and normally distributed errors, one would expect that

$$\sigma_{Ea}^2 \approx \sigma_{TS}^2 + \sigma_{Ad}^2 \quad (4)$$

and

$$\sigma_{\Delta E_{rxn}}^2 \approx 2\sigma_{Ad}^2 \quad (5)$$

where σ denotes the standard deviation of the property, TS denotes the transition state energy correction, Ad denotes the adsorbate energy correction, Ea denotes the activation barrier energy correction and ΔE_{rxn} denotes the reaction energy correction. The above equations imply that the error in the relative properties should be significantly larger than the error in the absolute properties. However, for our model $\sigma_{Ea}^2/(\sigma_{TS}^2 + \sigma_{Ad}^2) = 0.561$ and $\sigma_{\Delta E_{rxn}}^2/(2\sigma_{Ad}^2) = 0.694$, implying that for the relative energies our model is significantly more accurate than one would expect from the the accuracy of the absolute predictions. This demonstrates the power of natural error cancellation inherent in the structure of the SIDT predictor.

In order to explain how the error cancellation occurs, let us consider inference for the pair surface-bonds SIDT regressor used above. As discussed earlier, inference occurs by finding all pairs of surface bonds, tagging the surface bonds in each, descending each down the SIDT summing the contribution from each node touched in the descent and then summing the contribution from each pair of surface bonds. With this in mind let us consider computing the difference in energy between a given reactant configuration and the corresponding TS or product configuration. Only a handful of bonds are created or formed during a reaction, so the two configurations are usually not very different. Especially at higher coverages many of the interactions may be unchanged and thus their contributions to the two configurations are

identical and cancel exactly, removing any contribution to the variance from that interaction. In many other cases the interaction is only slightly changed resulting in an SIDT descend that only differs deep into the tree. In these cases the contributions from the upper nodes, before the descents diverge, cancel exactly and only the nodes farther down the tree, where the energy contributions should be much smaller, contribute to the variance. This natural error cancellation allows our SIDT predictor to be much more accurate on these relative properties that are actually important for simulations than one would expect from a given level of absolute accuracy.

Coverage-Dependent Corrections to Microkinetic Parameters

Simply considering the difference in energy between reactant, TS, and product configurations is sufficient for KMC simulations. However, mean-field kinetics simulators^{31–33} do not resolve the exact configurations of the co-adsorbates. For mean-field parameterizations we need to predict energies as a function of average coverage. This is typically done for a given species or TS by taking the lowest energy configuration at each coverage (corresponding to an integer number of co-adsorbates).

To find the lowest energy configuration, we generate all unique stable 2D representations at a given coverage, and then make energy predictions on them. We start with a list of all stable isolated configurations of the target adsorbate or TS. For a given co-adsorbate we iterate through all of the sites on the surface (here we always used a 3×3 slab). For each configuration in our list, if the site is free and placing the co-adsorbate on the site results in a configuration that our SIDT stability classifiers predict to be stable and is unique compared to the configurations in our list, we add the new configuration to the list. Once we have iterated through all sites, we have a list of all feasible configurations. We then use our SIDTs to predict the association energy of each configuration and find the lowest energy corresponding to each integer number of co-adsorbates. One can also integrate association energy predictions into the feasibility search and include an energy-based criterion to reduce

the number of feasible configurations generated, but we did not need to do so for the cases discussed here.

For a given adsorbate or TS this algorithm gives us a minimum association energy at a sequence of coverage values for the given co-adsorbate. To obtain corrections to isolated energies and activation barriers we need to find the energy required to move the isolated adsorbate or TS to a surface at a given co-adsorbate coverage. For an individual adsorbate co-adsorbed with a distinct co-adsorbate we can then calculate the coverage-dependent corrections based on

$$\Delta E_{\text{ad},N_{\text{coad}}} = E_{\text{ad},N_{\text{coad}}} - E_{N_{\text{coad}}} \quad (6)$$

where $E_{\text{ad},N_{\text{coad}}}$ is the predicted lowest association energy of the adsorbate or TS species and N_{coad} co-adsorbates, $E_{N_{\text{coad}}}$ is the predicted lowest association energy of N_{coad} co-adsorbates on the surface and $\Delta E_{\text{ad},N_{\text{coad}}}$ is the energy correction for the adsorbate energy with N_{coad} co-adsorbates on the surface. Our initial instinct might be to also apply this equation to adsorbates that are identical to the co-adsorbate and to TSs, however, these cases require further considerations.

Applying Eq. 6 to a case where the adsorbate is the same as the co-adsorbate yields $E_{N_{\text{coad}}+1} - E_{N_{\text{coad}}}$ on the right hand side. While the second term corresponds to the coverage of interest, the former corresponds to a higher coverage. We do, however, know that $E_{N_{\text{coad}}}$ is the energy to move N_{coad} co-adsorbates from isolation to the coverage corresponding to N_{coad} so we can instead compute

$$\Delta E_{\text{coad},N_{\text{coad}}} = \bar{E}_{N_{\text{coad}}} = \frac{E_{N_{\text{coad}}}}{N_{\text{coad}}} \quad (7)$$

where $\Delta E_{\text{coad},N_{\text{coad}}}$ is the energy correction for the co-adsorbate with N_{coad} co-adsorbates on the surface and $\bar{E}_{N_{\text{coad}}}$ is the average (per co-adsorbate) association energy of N_{coad} co-adsorbates on the surface.

The energy correction for any TS for a reaction that does not involve a co-adsorbate can

be calculated correctly with Eq. 6. However, reactions involving co-adsorbates have nuances similar to the co-adsorbate case. Let us consider the associative desorption of nitrogen gas on a surface with a nitrogen coverage of $\theta_N = 6/9$



first in the forward direction. An initial instinct might be to treat the TS as in Eq. 6 as a distinct configuration from N^* . However, consider the crowded surface at $\theta_N = 6/9$ we are putting the TS on. In order to form this TS separately from the $6/9 N^*$ on the surface we need to put the reactants (two additional N^*) on the surface. A local $\theta_N = 8/9$ may not just be high energy, on many surfaces it may not be stable at all. Following this line of thought one might conclude that such a reaction is very slow or impossible because of the energy/instability of the reactants that need formed first. However, with the surface at $\theta_N = 6/9$, assuming that coverage is stable, and the TS is short-lived, one would expect this reaction to occur naturally between the co-adsorbates with two N^* participating and the remaining four co-adsorbed nearby occurring effectively at $\theta_N = 4/9$. Interestingly, though since this is a consideration of the reactants and not the reaction itself, this effect does not occur in the reverse direction where N_2 must adsorb at $\theta_N = 6/9$. We can account for this reactant-co-adsorbate effect by computing the TS energy correction in a particular direction from

$$\Delta E_{TS,N_{coad}} = E_{TS,(N_{coad}-N_{coad,reactant})} - E_{N_{coad}} \quad (9)$$

where $N_{coad,reactant}$ is the number of co-adsorbate participants in the reactants for the reaction direction of interest, $E_{TS,N}$ is the association energy of the transition state with N nearby co-adsorbates and $\Delta E_{TS,N_{coad}}$ is the energy correction for the TS.

Given the very large configurational space, computing the exact DFT correction for a single adsorbate or TS with respect to one co-adsorbate species would have required computational expense on par with generating the entire dataset in this work, making it com-

putationally too expensive to present parity plots. Instead, here we used a simple iterative refinement procedure to analyze the accuracy of our predictions. In each iteration we re-trained the SIDTs and predicted the lowest energy stable configurations at each coverage level according to the above procedure. We then took the predicted lowest energy configurations and attempted to calculate their energy using DFT and added them to the dataset for training the next iteration of the SIDT. We ran two refinement iterations for each case presented here.

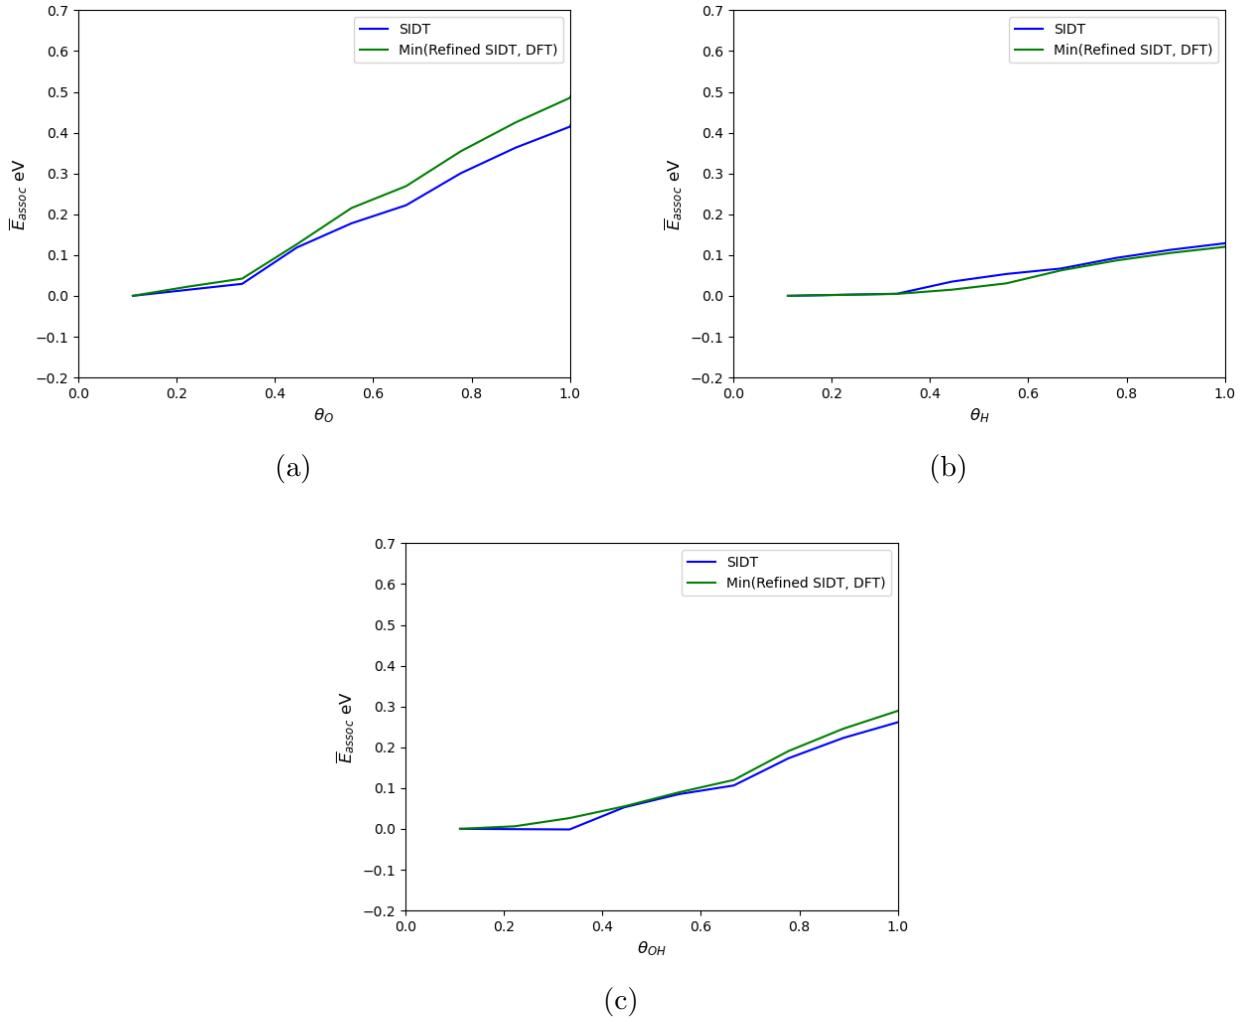


Figure 9: Average association energies for the lowest energy configurations for (a) O^* , (b) H^* , and (c) HO^* as a function of coverage of (a) O^* , (b) H^* , and (c) HO^* based on an iterative refinement process along with DFT calculations for the predicted lowest energy structures. SIDT corresponds to the original model. Min(Refined SIDT, DFT) is the minimum of the final SIDT after two refinement cycles and the associated DFT calculations.

We show the average association energy, $E_{N_{\text{coad}}}/N_{\text{coad}}$, for the lowest energy configurations for H^* , O^* , and HO^* in Figures 9a, 9b, and 9c respectively computed using the original SIDT model and the minimum of the refined SIDT model and the DFT calculations that we will refer to as the baseline. For all cases the SIDT agrees well with the baseline at lower coverages where significantly more training data is available for the larger co-adsorbates. For H^* , whose configurations have higher success rates at high coverages, this extends out to higher coverages. For the larger co-adsorbates O^* and HO^* it should be noted that stability is particularly difficult to predict on $\text{Cu}(111)$ at high coverages. We suspect that some of the higher coverages for these larger co-adsorbates may not be stable at all.

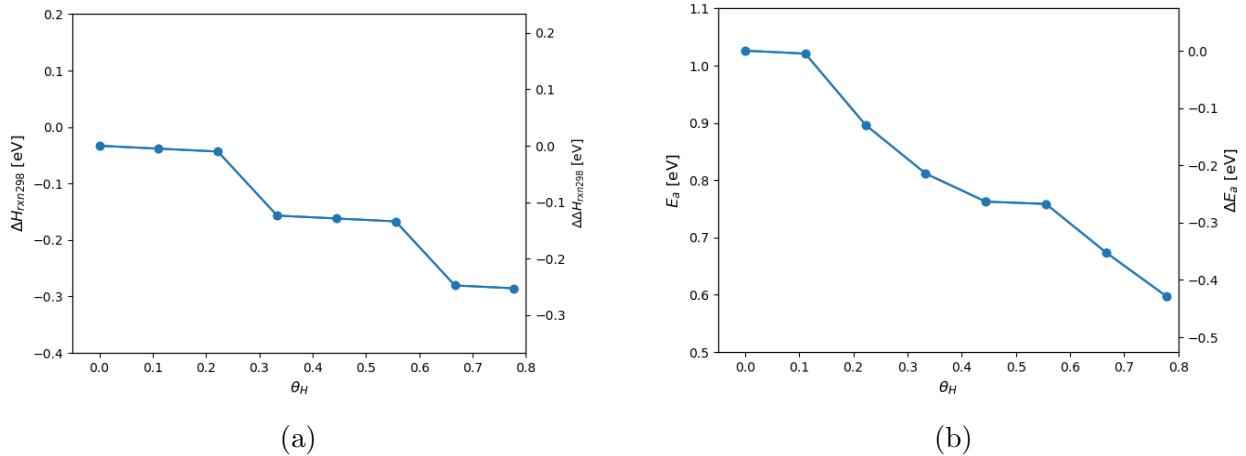


Figure 10: SIDT predicted H^* coverage dependence of the (a) enthalpy of reaction at 298 K and (b) activation barrier for the $\text{HO}^* + \text{H}^* \rightarrow \text{H}_2\text{O} + 2*$ reaction on $\text{Cu}(111)$. The secondary axes show the change in enthalpy of reaction at 298 K and activation barrier, as appropriate, relative to the isolated value.

In Figures 10 and 11 we apply our model to estimate coverage-dependent properties of the $\text{HO}^* + \text{H}^* \rightarrow \text{H}_2\text{O} + 2*$ reaction on $\text{Cu}(111)$ using Equations 6–9 and the isolated properties of the reaction from Johnson et al.²⁰ In Figure 10a we examine the enthalpy of reaction at 298 K. This reaction removes two adsorbates from the surface and the associated lateral interactions, so naturally we expect that at higher coverages where lateral interactions are more significant, the reactants will be higher in energy and thus the enthalpy of reaction will decrease, which is in agreement with Figure 10a. The activation barrier is shown in Figure

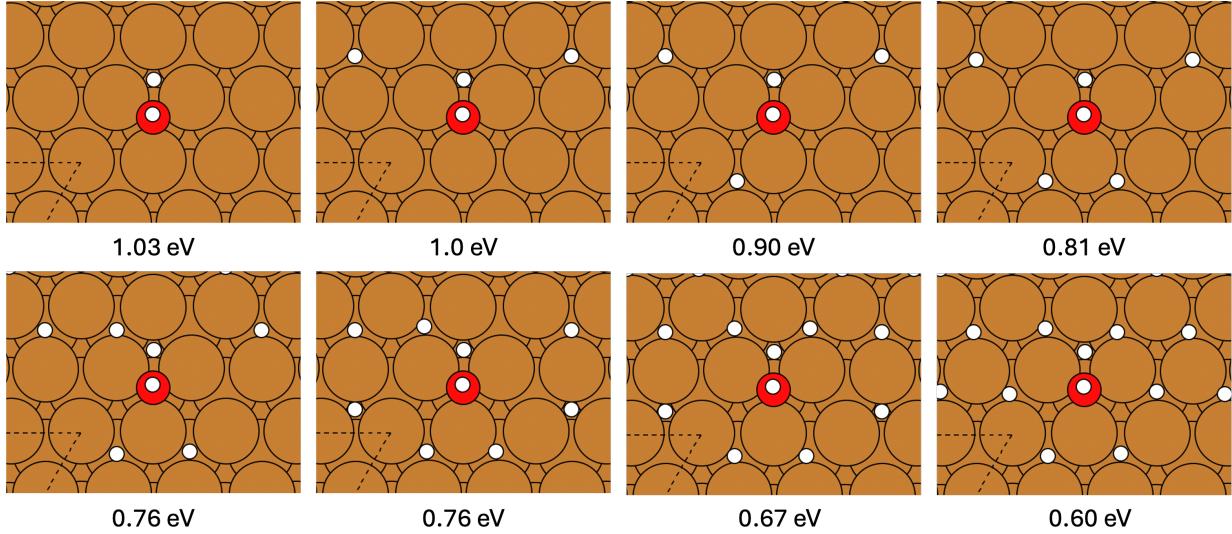


Figure 11: SIDT predicted lowest association energy TS configurations (generated initial guess geometries) for $\text{HO}^* + \text{H}^* \rightarrow \text{H}_2\text{O} + 2^*$. Energies are the SIDT predicted mean field activation barriers.

10b and the associated lowest energy TS configurations are shown in Figure 11. The barrier also decreases with coverage as a result of the stronger lateral interactions in the reactants at higher coverages. However, it is particularly interesting that the activation barrier is predicted to decrease by slightly more than the enthalpy of formation. For this particular reaction the product is in the gas phase, and thus does not contribute to the corrections, and for the activation barrier and enthalpy of formation the correction associated with the reactants is the same, so the only difference is the TS energy correction. Especially for larger co-adsorbates one would expect the TS correction energy to be positive and increase with coverage. However, here it appears to be slightly negative and relatively constant. The effect size is well within the leave-one-out errors on TS configurations so in principle this could just be a result of model error. However, if we revisit Figure 11, particularly the top fourth from the left configuration we see a -0.23 eV barrier correction for a $\theta_{\text{H}} = 4/9$ configuration (keep in mind Eq. 9). This number is from the end points of the reaction path (rather than the lowest energy TS and reactants) so it is not perfectly comparable, but it does strongly suggest that the mean field activation barrier correction prediction at $\theta_{\text{H}} = 4/9$ of -0.26 eV

is not unreasonable. If one looks carefully at Figure 4 configurations (row, column): (1,2), (1,3), (1,4), (1,5), (1,6), (2,1), (2,2), (2,4), (2,5), (3,1), (3,2), (3,3), (3,4), (3,5), (3,6), (4,1), (4,2), and (4,3) one can see that for the TSs with the O planted clearly in a hollow and the non-reactive H tilted the Cu atom adjacent to the O and opposite the tilt of the H is slightly displaced. If this displacement provides a slight stabilization effect for some of the adsorbed hydrogens, it could overcome the competing lateral interactions which are very weak because the co-adsorbate is a hydrogen and this TS is not very tightly bound to the surface since the product is a gas phase species.

Discussion

The Cu(111) Models

The presented Cu(111) models are a good starting point for general approximate estimation of adsorbate and reaction coverage dependence effects for initial mechanism construction, such as in RMG.^{30,34,35} Our models are sufficient for approximate estimation of adsorbates and similar transition states composed of H, C, O and N atoms and co-adsorbed with H*, O* or HO* at low to medium coverages on Cu(111) based on our results shown in Figures 9a-9c. While we did sample with N* as a co-adsorbate, we had too few valid samples and did not find the model able to make good predictions with N* as the co-adsorbate. We attribute this to N*'s tendency to cause restructuring on Cu, which makes it more likely samples will not converge and can significantly change the interactions between adsorbates

Calculation of Co-Adsorbed Configurations

Our presented approach to calculate and analyze co-adsorbed configurations is a highly effective and efficient way to examine the co-adsorbed configurational space based on combining direct DFT calculations with a low-data ML approach. It is however, worth discussing that our approach relies on the assumption that configurations that fail to optimize to the target

species (i.e., to the given minimum or saddle) are unstable and do not exist. Failure to optimize a configuration to a target does not imply in general that the target does not exist. If an initial guess is too far from the target, a configuration may optimize to a different species. However, configurations that fail to optimize, but do in fact exist, are likely to be high in energy. Since these configurations are high in energy and we are in general interested in and/or sensitive to the lowest energy configurations, the distinction between these edge case configurations being stable and unstable is unlikely to be very important.

Advantages of Using Subgraph Isomorphic Decision Trees

As discussed in the introduction and demonstrated here, our SIDT-based approach to coverage dependence is significantly more flexible, automatic, and powerful than current state-of-the-art cluster expansion techniques. In particular, SIDT is able to learn interactions CE has to be explicitly told how to resolve. Furthermore, although we did not demonstrate it in this manuscript, SIDT is able to run inference on arbitrary co-adsorbed configurations not just those adsorbates and co-adsorbates a CE scheme is fit for.

Simultaneously, SIDT is easier to apply, more flexible, and easier to interpret than possible DNN based approaches. Crucially, SIDT can be applied to much smaller datasets than is feasible for DNNs, and because of its interpretability it is much easier to analyze results and thus to improve performance.

Conclusions

Coverage dependence of chemical reactions is a key, but often ignored aspect of microkinetic model construction because of the computational expense and complexity that it requires to determine the necessary parameters.¹⁴ Comparable challenges such as rate coefficient pressure dependence^{36,37} in gas phase have readily available tools^{38–40} that are able to fully automate high accuracy computations using ab initio methods. The framework and tools

presented here open the door for decreasing the barrier to include coverage dependence routinely in future microkinetic models. Here we used our low-data machine-learning approach called the subgraph isomorphic decision tree (SIDT) to construct a predictor based on a fixed pre-generated dataset. For future applications it is easy to imagine using the workflow within an active learning scheme that automatically identifies what configurations should be calculated to improve the SIDT predictors. Our entire workflow is built to be automatic using Pynta, running the necessary calculations and post-processing them.

However, higher-level generalizations, such as the ones built into software such as RMG^{30,34,35} allow for efficiently approximating kinetic parameters without any ab initio calculations for instance for pressure-dependent reactions in the gas-phase.^{41,42} It is possible to imagine a similar, generalized approach, for instance for Cu(111) using the model presented in this work. However, in general it may be impractical to evaluate the SIDTs at every possible 2D configuration, as done in this work, to find the lowest energy configurations. Doing so is unlikely to be strictly necessary, but the scheme by which configurations are sampled must be considered carefully. Approximations across arbitrary or even a range of metals, however, is much more challenging. Constructing such a scheme might be best done by training a foundational SIDT model on one surface across a wide range of adsorbates, transition states, and co-adsorbates and then (e.g., by applying delta learning) training correction SIDTs on much smaller datasets to predict the difference between the foundational SIDT model and particular surfaces.

We have presented a framework for generating machine learning models and applying them to predict coverage dependent kinetic parameters for microkinetic models. Our toolkit enables automatic ab initio computation of co-adsorbed configurations and automatic post-processing including identification of the optimized configuration and TS validity evaluation for TSs. We demonstrated the training of SIDT on the generated dataset to predict the stability and association energy of co-adsorbed configurations. Lastly, we explained how to use the SIDTs to compute mean-field coverage dependent energy corrections for adsorbates

thermochemistry and reaction activation barriers.

On Cu(111), a challenging surface, we were able to achieve association energy MAEs of 0.106 eV on adsorbates and 0.172 eV on transition states and due to natural error cancellation in SIDTs on relative properties MAEs of 0.130 eV on reaction energies and 0.180 eV on activation barriers. We hope to extend these techniques to enable high accuracy and efficient calculation of coverage dependent kinetic parameters.

Supporting Information

The supporting information contains the generated datasets and the PySIDT models as described in the Supplementary Information document.

Acknowledgement

This work was done within the Exascale Catalytic Chemistry (ECC) Project, which is supported by the U.S. Department of Energy, Office of Science, Basic Energy Sciences, Chemical Sciences, Geosciences and Biosciences Division, as part of the Computational Chemistry Sciences Program.

This research used resources of the National Energy Research Scientific Computing Center (NERSC), a Department of Energy Office of Science User Facility using NERSC award BES-ERCAP0026789.

This article has been authored by employees of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employees co-own right, title and interest in and to the article and are responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government pur-

poses. The work at Argonne National Laboratory was supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, Chemical Sciences, Geosciences and Biosciences Division, under Contract No. DE-AC02-06CH11357). The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan <https://www.energy.gov/downloads/doe-public-access-plan>.

References

- (1) Kitchin, J. R. Correlations in coverage-dependent atomic adsorption energies on Pd(111). *Phys. Rev. B* **2009**, *79*, 205412.
- (2) Getman, R. B.; Schneider, W. F.; Smeltz, A. D.; Delgass, W. N.; Ribeiro, F. H. Oxygen-Coverage Effects on Molecular Dissociations at a Pt Metal Surface. *Phys. Rev. Lett.* **2009**, *102*, 076101.
- (3) Grabow, L. C.; Hvolbæk, B.; Nørskov, J. K. Understanding Trends in Catalytic Activity: The Effect of Adsorbate–Adsorbate Interactions for CO Oxidation Over Transition Metals. *Topics in Catalysis* **2010**, *53*, 298–310.
- (4) İnoğlu, N.; Kitchin, J. R. Simple model explaining and predicting coverage-dependent atomic adsorption energies on transition metal surfaces. *Phys. Rev. B* **2010**, *82*, 045414.
- (5) Lausche, A. C.; Medford, A. J.; Khan, T. S.; Xu, Y.; Bligaard, T.; Abild-Pedersen, F.; Nørskov, J. K.; Studt, F. On the effect of coverage-dependent adsorbate–adsorbate interactions for CO methanation on transition metal surfaces. *Journal of Catalysis* **2013**, *307*, 275–282.
- (6) Lu, J.; Behtash, S.; Faheem, M.; Heyden, A. Microkinetic modeling of the decarboxylation and decarbonylation of propanoic acid over Pd(111) model surfaces based on parameters obtained from first principles. *Journal of Catalysis* **2013**, *305*, 56–66.

(7) Majumdar, P.; Greeley, J. Generalized scaling relationships on transition metals: Influence of adsorbate-coadsorbate interactions. *Phys. Rev. Mater.* **2018**, *2*, 045801.

(8) Wu, C.; Schmidt, D.; Wolverton, C.; Schneider, W. Accurate coverage-dependence incorporated into first-principles kinetic models: Catalytic NO oxidation on Pt (111). *Journal of Catalysis* **2012**, *286*, 88–94.

(9) Xu, Z.; Kitchin, J. R. Probing the Coverage Dependence of Site and Adsorbate Configurational Correlations on (111) Surfaces of Late Transition Metals. *The Journal of Physical Chemistry C* **2014**, *118*, 25597–25602.

(10) Bronsted, J. N.; Sandved, K. H.; Lamer, V. K. Acid and Basic Catalysis. *Chemical Reviews* **1928**, *5*, 231–338.

(11) Evans, M. G.; Polanyi, M. Inertia and driving force of chemical reactions. *Transactions of the Faraday Society* **1938**, *34*, 11–24.

(12) Miller, S. D.; Kitchin, J. R. Relating the coverage dependence of oxygen adsorption on Au and Pt fcc(1 1 1) surfaces through adsorbate-induced surface electronic structure effects. *Surface Science* **2009**, *603*, 794–801.

(13) Lerch, D.; Wieckhorst, O.; Hammer, L.; Heinz, K.; Müller, S. Adsorbate cluster expansion for an arbitrary number of inequivalent sites. *Physical Review B - Condensed Matter and Materials Physics* **2008**, *78*, 121405.

(14) Nolen, M. A.; Farberow, C. A.; Kwon, S. Incorporating Coverage-Dependent Reaction Barriers into First-Principles-Based Microkinetic Models: Approaches and Challenges. *ACS Catalysis* **2024**, 14206–14218.

(15) Shan, B.; Kapur, N.; Hyun, J.; Wang, L.; Nicholas, J. B.; Cho, K. CO-coverage-dependent oxygen dissociation on Pt(111) surface. *Journal of Physical Chemistry C* **2009**, *113*, 710–715.

(16) Frey, K.; Schmidt, D. J.; Wolverton, C.; Schneider, W. F. Implications of coverage-dependent O adsorption for catalytic NO oxidation on the late transition metals. *Catalysis Science & Technology* **2014**, *4*, 4356–4365.

(17) Johnson, M. S.; Green, W. H. A machine learning based approach to reaction rate estimation. *Reaction Chemistry & Engineering* **2024**, *9*, 1364–1380.

(18) Pang, H. W.; Dong, X.; Johnson, M. S.; Green, W. H. Subgraph Isomorphic Decision Tree to Predict Radical Thermochemistry with Bounded Uncertainty Estimation. *Journal of Physical Chemistry A* **2024**, *128*, 2891–2907.

(19) Johnson, M. S.; Pang, H.-W.; Doner, A. C.; Green, W. H.; Zador, J. PySIDT: Subgraph Isomorphic Decision Trees for Molecular Property Prediction. *ChemRxiv* **2024**,

(20) Johnson, M. S.; Gierada, M.; Hermes, E. D.; Bross, D. H.; Sargsyan, K.; Najm, H. N.; Zádor, J. Pynta – An Automated Workflow for Calculation of Surface and Gas-Surface Kinetics. *Journal of Chemical Information and Modeling* **2023**, *63*, 5168.

(21) Wellendorff, J.; Lundgaard, K. T.; Møgelhøj, A.; Petzold, V.; Landis, D. D.; Nørskov, J. K.; Bligaard, T.; Jacobsen, K. W. Density functionals for surface science: Exchange-correlation model development with Bayesian error estimation. *Physical Review B - Condensed Matter and Materials Physics* **2012**, *85*, 235149.

(22) Giannozzi, P.; Baroni, S.; Bonini, N.; Calandra, M.; Car, R.; Cavazzoni, C.; Ceresoli, D.; Chiarotti, G. L.; Cococcioni, M.; Dabo, I. et al. QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials. *Journal of Physics: Condensed Matter* **2009**, *21*, 395502.

(23) Giannozzi, P.; Andreussi, O.; Brumme, T.; Bunau, O.; Nardelli, M. B.; Calandra, M.; Car, R.; Cavazzoni, C.; Ceresoli, D.; Cococcioni, M. et al. Advanced capabilities for materials modelling with Quantum ESPRESSO. *Journal of Physics: Condensed Matter* **2017**, *29*, 465901.

(24) Larsen, A. H.; Mortensen, J. J.; Blomqvist, J.; Castelli, I. E.; Christensen, R.; Dułak, M.; Friis, J.; Groves, M. N.; Hammer, B.; Hargus, C. et al. The atomic simulation environment—a Python library for working with atoms. *Journal of Physics: Condensed Matter* **2017**, *29*, 273002.

(25) Hermes, E. D.; Sargsyan, K.; Najm, H. N.; Zádor, J. Sella, an Open-Source Automation-Friendly Molecular Saddle Point Optimizer. *Journal of Chemical Theory and Computation* **2022**, *18*, 6974–6988.

(26) Hermes, E. D.; Sargsyan, K.; Najm, H. N.; Zádor, J. Accelerated Saddle Point Refinement through Full Exploitation of Partial Hessian Diagonalization. *Journal of Chemical Theory and Computation* **2019**, *15*, 6536–6549.

(27) Hermes, E. D.; Sargsyan, K.; Najm, H. N.; Zádor, J. Geometry optimization speedup through a geodesic approach to internal coordinates. *The Journal of Chemical Physics* **2021**, *155*, 094105.

(28) Muñoz-Márquez, M. A.; Parkinson, G. S.; Quinn, P. D.; Gladys, M. J.; Tanner, R. E.; Woodruff, D. P.; Noakes, T. C.; Bailey, P. N-induced pseudo-(1 0 0) reconstruction of Cu(1 1 1): One layer or more? *Surface Science* **2005**, *582*, 97–109.

(29) Han, S.; Lysgaard, S.; Vegge, T.; Hansen, H. A. Rapid mapping of alloy surface phase diagrams via Bayesian evolutionary multitasking. *npj Computational Materials* **2023**, *9*.

(30) Liu, M.; Dana, A. G.; Johnson, M. S.; Goldman, M. J.; Jocher, A.; Payne, A. M.; Grambow, C. A.; Han, K.; Yee, N. W.; Mazeau, E. J. et al. Reaction Mechanism Generator v3.0: Advances in Automatic Mechanism Generation. *Journal of Chemical Information and Modeling* **2021**, *61*, 2686–2696.

(31) Johnson, M. S.; Pang, H. W.; Payne, A. M.; Green, W. H. ReactionMechanismSimu-

lator.jl: A modern approach to chemical kinetic mechanism simulation and analysis. *International Journal of Chemical Kinetics* **2024**, *56*, 732–747.

(32) Goodwin, D. G.; Speth, R. L.; Moffat, H. K.; Weber, B. W. Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes. 2021; <https://www.cantera.org>.

(33) Ansys Chemkin-Pro — Chemical Kinetics Simulation Software. <https://www.ansys.com/products/fluids/ansys-chemkin-pro>.

(34) Johnson, M. S.; Dong, X.; Dana, A. G.; Chung, Y.; Farina, J. D.; Gillis, R. J.; Liu, M.; Yee, N. W.; Blondal, K.; Mazeau, E. et al. RMG Database for Chemical Property Prediction. *Journal of Chemical Information and Modeling* **2022**, *62*, 4906–4915.

(35) Johnson, M. S.; Pang, H. W.; Liu, M.; Green, W. H. Species selection for automatic chemical kinetic mechanism generation. *International Journal of Chemical Kinetics* **2024**,

(36) Klippenstein, S. J. From theoretical reaction dynamics to chemical modeling of combustion. *Proceedings of the Combustion Institute* **2017**, *36*, 77–111.

(37) Johnson, M. S.; Green, W. H. Examining the accuracy of methods for obtaining pressure dependent rate coefficients. *Faraday Discussions* **2022**, *238*, 380–404.

(38) de Vijver, R. V.; Zádor, J. KinBot: Automated stationary point search on potential energy surfaces. *Computer Physics Communications* **2020**, *248*, 106947.

(39) Zádor, J.; Martí, C.; Van de Vijver, R.; Johansen, S. L.; Yang, Y.; Michelsen, H. A.; Najm, H. N. Automated reaction kinetics of gas-phase organic species over multiwell potential energy surfaces. *J. Phys. Chem. A* **2023**, *127*, 565–588.

(40) Elliott, S. N.; Moore, K. B.; Copan, A. V.; Keçeli, M.; Cavallotti, C.; Georgievskii, Y.; Schaefer, H. F.; Klippenstein, S. J. Automated theoretical chemical kinetics: Predicting

the kinetics for the initial stages of pyrolysis. *Proceedings of the Combustion Institute* **2021**, *38*, 375–384.

(41) Matheu, D. M.; Lada, T. A.; Green, W. H.; Dean, A. M.; Grenda, J. M. Rate-based screening of pressure-dependent reaction networks. *Computer Physics Communications* **2001**, *138*, 237–249.

(42) Johnson, M. S.; Dana, A. G.; Green, W. H. A workflow for automatic generation and efficient refinement of individual pressure-dependent networks. *Combustion and Flame* **2023**, *257*, 112516.

TOC Graphic

