Competing adsorption of H and CO on Pd-alloy surfaces: Mechanistic insight into the mitigating effect of Cu on CO poisoning

Pernilla Ekborg-Tanner, Paul Erhart*

Department of Physics, Chalmers University of Technology, SE 412 96 Gothenburg, Sweden

Abstract

Multi-component alloys offer broad tunability for addressing challenges in materials science, but their vast configurational space makes their surface chemistry highly sensitive to operating conditions, for example through adsorption and segregation. Here, we study Pd–Au–Cu alloy surfaces in H_2 and CO environments motivated by their use in H technologies, in particular plasmonic H_2 sensing, where alloying can mitigate limitations intrinsic to Pd such as hysteresis and CO poisoning. Modeling multicomponent surfaces with multiple adsorbate species under realistic conditions is challenging. To this end, we establish an accurate and efficient framework that combines machine-learned interatomic potentials trained on density-functional theory data to generate training data for cluster expansions with effectively no limitations on training set size.

By constructing continuous surface phase diagrams for H–CO coadsorption we find that coadsorption under operating conditions is governed primarily by the H coverage during annealing. Au-rich surfaces, formed under H-poor conditions, suppress both CO and H adsorption, while H-rich conditions yield Pd-rich surfaces that maintain higher H coverages compared to Pd at relevant CO partial pressures, indicating improved CO poisoning resistance. This effect is insensitive to relative amounts of Au and Cu, despite experimental evidence of the mitigating effect of specifically Cu on CO poisoning. Kinetic barriers for dilute alloy surfaces indicate that absorption pathways near Au are highly unfavorable, while Cu leave the energetics unchanged compared to pure Pd. This finding suggests that Cu in the surface region provides viable pathways to shuttle H into the material when Pd-dominated paths are blocked by CO.

Keywords: Alloy surfaces, Surface phase diagrams, Coadsorption, Hydrogen sensing, Multiscale modeling

1. Introduction

Multi-component materials underpin applications ranging from chemical sensing and catalysis to energy conversion and structural engineering. Under realistic operating conditions, such as finite temperatures and reactive gas environments, the surface properties of these materials are often changed due to adsorption and segregation [1, 2]. Such processes critically influence materials performance by altering kinetic, thermodynamic, electronic, optical, and mechanical properties. Understanding and controlling these phenomena is, therefore, essential for the design of functional multi-component materials. Yet, current modeling approaches still struggle to capture the immense combinatorial and structural complexity of these systems, particularly when adsorbates, surfaces, and thermodynamic effects all come into play.

Pd-based nanoalloys have received considerable attention, especially in the context of optical hydrogen sensing, which exploits changes in the position and shape of the localized surface plasmon resonance (LSPR) of Pd nanoparticles upon hydrogen sorption [3]. While the basic working principle is well established, further improvements are needed to make this technology robust under realistic conditions. Pure Pd sensors suffer from hysteresis between hydrogen absorption and desorption, as well as carbon monoxide (CO) poisoning caused by the preferential adsorption of CO over H [4]. These effects directly compromise the reliability and stability of LSPR-based hydrogen sensors. Alloying is an effective mitigation strategy, as the addition of approximately 25 % Au suppresses hysteresis, while introducing only 5 % Cu has been found to reduce CO poisoning [5, 6, 7]. While the effect of Au on hysteresis is well understood [8], the mechanism behind the mitigating effect of Cu is not yet understood, especially since Cu rarely resides directly at the surface [2].

Introducing multiple alloying elements and adsorbate species, however, substantially enlarges the design space and creates additional complexity due to the nontrivial coupling between adsorption and segregation. This system thus provides a technologically relevant and scientifically rich case that demands advancing models of multi-component materials under reactive conditions.

To comprehensively model such a system, a framework is required that (i) captures subtle energy differences between competing atomic arrangements in the bulk and at surfaces, (ii) treats the interactions among adsorbates and between adsorbates and surfaces with sufficient accuracy, and (iii) enables efficient sampling across a range of thermodynamic boundary conditions, including temperature and gas pressure. Here, we

*Corresponding author
Email address: erhart@chalmers.se (Paul Erhart)

address these challenges by combining machine-learned interatomic potentials (MLIPs) trained on density-functional theory (DFT) calculations with cluster expansions (CEs) and Monte Carlo (MC) sampling and apply this approach to study the competing adsorption between H and CO on PdAuCu alloy surfaces.

We find that the surface coverage of H and CO depends on the chemical configuration of the surface region, which is determined by the preparation conditions. In addition, we find that alloying increases the H-to-CO ratio compared to pure Pd, regardless of the exact alloy composition. The adsorption thermodynamics alone does not, however, explain the experimental observation that Cu is necessary to mitigate CO poisoning. To gain further insight into the CO poisoning mechanism, we therefore study the kinetics of H sorption into the bulk by analyzing the kinetic barriers associated with H absorption of dilute PdAu and PdCu alloys. We find that paths close to Au are highly unfavorable, while Cu displays similar energetics as pure Pd. This leads us to propose that Cu mitigates CO poisoning by providing viable pathways in situations where the most favorable paths are blocked by CO.

2. Methodology

Electronic-structure calculations, most commonly based on DFT, remain the common standard for accuracy when modeling extended systems. However, their high computational cost renders them impractical for systems with multiple alloying components, low symmetries, and finite-temperature effects such as mixing, ordering, and segregation.

CEs provide an efficient and well-established framework for modeling configurational complexity in alloys [9]. They have been successfully applied to multi-component bulk alloys [10] as well as surfaces with adsorbates [2, 11, 12]. In low-symmetry systems, such as surfaces, the number of interaction parameters grows rapidly, raising the demand for training data, which is still most commonly generated using DFT calculations. While Bayesian approaches help manage this growth by incorporating physical insight via priors [12], the combinatorial explosion that occurs when increasing the number of components remains challenging [10, 13, 14]. When multiple surface orientations are relevant, it is furthermore desirable for the different surface-specific CEs to yield consistent behavior in the bulk limit. Without such consistency, predictions of surface segregation across facets may become unreliable, which poses the additional challenge of reconciling models for different surface orientations.

In the last decade, MLIPs have emerged as powerful tools that bridge the accuracy of electronic-structure methods with the efficiency of empirical potentials. They open the possibility of serving as intermediaries for constructing CEs, offering accuracy at scale [15]. While MLIPs have already been deployed in hybrid molecular dynamics (MD)-MC simulations to model ordering and mixing in bulk alloys [16, 17], these approaches are not well suited for sampling adsorbate–gas equilibria. For this purpose, lattice-based CE-MC simulations remain the most effective strategy.

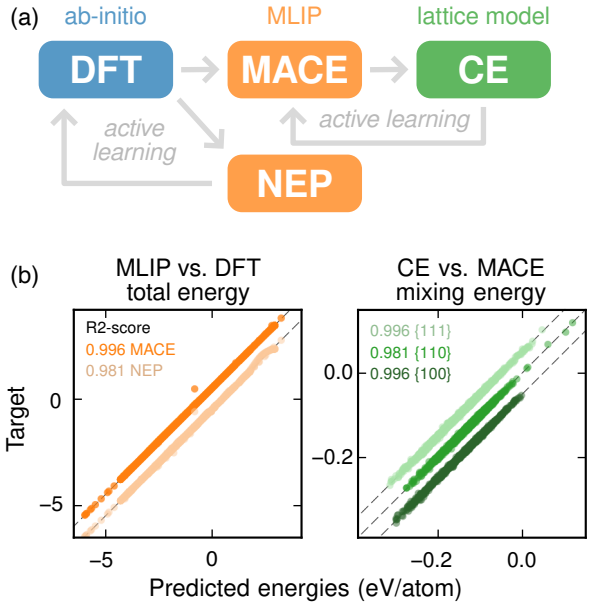

Figure 1: **Model development.** (a) Schematic illustration of the computational framework. (b) Model errors for the MLIPs and the CE. Note that the data has been systematically shifted along the y-axis to allow for clearer visualization of the different models.

To tackle the above challenges, we adopt the following approach. To cut the cost of DFT calculations, we build MLIPs based on both the neuroevolution potential (NEP) and the MACE architecture. Specifically, we use NEP models, which are up to three orders of magnitude faster than MACE (Fig. S2), to build a comprehensive set of reference data via active learning, and MACE models, which yield higher accuracy, to generate reference data for training CE models. The CEs are constructed for {111}, {110}, and {100} surfaces that are coupled to the same underlying bulk model through constraints, and include Au, Pd, and Cu on the metal lattice as well as H and CO on the adsorbate lattices. Training data for these models are generated from structural relaxation using the MACE model, such that relaxation effects are implicitly included, in an active learning loop. In taking this approach, we exploit that different types of predictions require different levels of accuracy.

2.1. *Machine-learned interatomic potentials*

We construct MLIPs using the fourth-generation NEP framework [17, 18], the separable neuroevolution strategy [19] and the iterative procedure described in Ref. 20, utilizing the GPUMD [18, 21] and CALORINE [22] packages, as detailed in Supp. Note 1. A combination of structure enumeration [23, 24], nudged elastic band (NEB) [25, 26, 27] pathways and MD trajectories are used to generate bulk and surface structures in an active training loop, using the NEP models and the ASE [28], HIPHIVE [29], and ICET packages [30, 31]. The final set of reference structures comprises 18 403 structures, corresponding to a total of 586 006 atoms (see Table S3 for an overview). Using the final set of reference structures, we construct a MACE model [32] using the MACE-TORCH package [33] as described in Supp. Note 2.

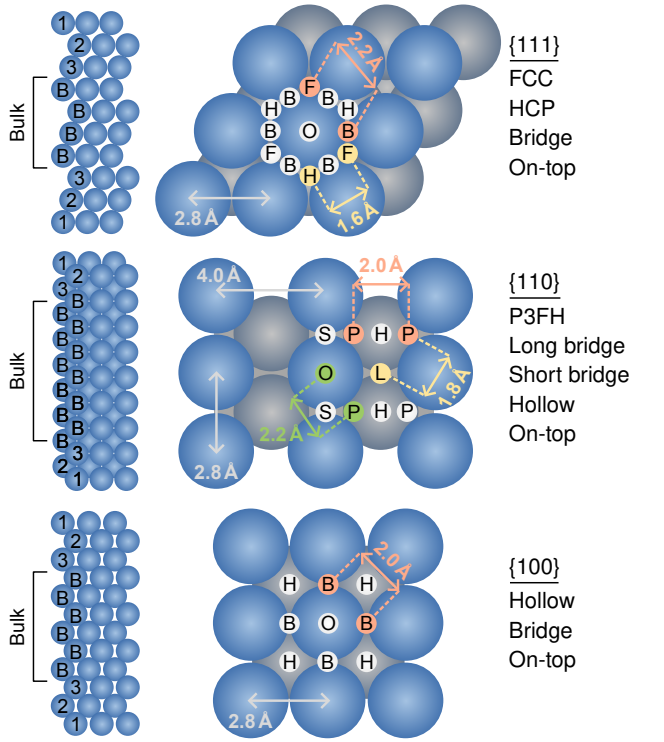

Figure 2: **Schematic illustration of the surface slabs and their respective adsorption sites.** Side views of the surface slabs (to the left) with the surface layers and bulk region indicated. Top views of the surfaces (to the right) with blue atoms for the top layer, gray atoms for the subsequent layers. The smaller atoms show the adsorption sites labeled by their first letter. The arrows between sites indicate their three-dimensional assuming a 4.0 Å lattice parameter.

2.2. *Density functional theory calculations*

The MLIPs are fitted using forces, energies, and virials obtained from DFT calculations, employing the van-der-Waals density functional with the consistent exchange (vdW-DF-cx) method [34, 35] as implemented in the Vienna ab-initio simulation package [36, 37] using projector-augmented wave setups [38, 39] with a plane wave energy cutoff of 520 eV. The Brillouin zone is sampled with automatically generated **k**-point grids with a maximum spacing of 0.1/Å.

2.3. *Cluster expansion construction*

CEs are constructed for face-centered cubic (FCC) surface-adsorbate slabs with {111}, {110} and {100} orientation using the ICET package [30] following the procedure outlined in Ref. 31. The CEs predict the mixing energy based on training data obtained by atomic relaxation and energy calculation using the MACE model. The mixing energies are calculated by using the energy per atom of a pure metal slab for the metal species and the molecular ground state energy (divided by 2 for H_2) for the adsorbate species. Figure 2 shows all atomic sites of the slab and adsorbates. The number of layers varies between the slabs to achieve a similar slab thickness, with 10 layers for {111}, 18 for {110} and 12 for {100}.

The adsorbate sublattice represents an open system and can be occupied by H, CO, or vacancies. All relevant adsorption sites are considered, which allows for configurations with coverages far beyond the relevant limit. Since we do not take into account H sorption into the bulk, only surface coverages up to the limit where entering the subsurface becomes favorable are of interest. This corresponds to a 1:1 ratio between adsorbate and surface metal atoms for {111} and {100} and 2:1 for {110} (Fig. S4). The larger allowed coverage on {110} surfaces is a consequence of the large surface area of {110} where almost the entire second atomic layer is exposed (Fig. 2). For this reason, we define 100 % surface coverage on {110} to mean a 2:1 ratio, as opposed to 1:1 for the other facets, which allows for a fairer comparison between the surface orientations.

Cutoffs are selected to include pairs and triplets up to distances of 8.0 Å and 5.0 Å, respectively, for a 4.0 Å lattice parameter. The number of interaction parameters depends on the number of layers and adsorption sites and ranges from 3300 ({100}) to 9600 ({110}) for the systems considered here. In the context of CE models, these are very large parameter spaces which would be difficult to fit properly using conventional approaches. To overcome this challenge, we exploit pruning and merging of clusters.

Because the adsorbate lattice contains many close-lying, symmetrically inequivalent sites, their interactions will contribute to a very large number of inequivalent clusters. Many of these correspond to interaction lengths shorter than the minimal distance between adsorbates at the maximum coverage limit. For this reason, we introduce a lower cutoff radius of 1.5 Å for {111} and {100}, and 1.2 Å for {110} and prune all clusters below this limit. We also make the pragmatic choice of pruning all adsorbate-adsorbate interactions longer than the FCC nearest-neighbor distance (2.8 Å) as well as all adsorbate-slab interactions extending below the third layer.

Merging, in this context, refers to grouping together clusters that are not strictly symmetrically equivalent but can be assumed to be similar [2, 31]. This includes clusters far from the surface region, which should be similar to the equivalent cluster of a bulk system. Here, we consider the three outermost atomic layers to be the surface region and all inner layers to be the bulk, and merge all clusters with the same radius that consist of bulk sites only, as described in Ref. 31. To ensure the CE models for the different surface orientations yield a correct and consistent description of the bulk region, we first construct a bulk CE for the AuCuPd system based on approximately 900 enumerated structures with up to 9 atoms in the unit cell and up to 50 % Au and 20 % Cu. The resulting bulk interaction parameters are used in place of the merged bulk interactions for the surface slabs.

Pruning, merging and fixing the bulk interactions cuts the number of orbits by half, resulting in 2000 to 4500 parameters, which can be further reduced by regularization. Because these numbers are still too large to efficiently apply common advanced linear regression methods, a two step approach is adapted. First, an initial feature selection is performed using the least absolute shrinkage and selection operator (LASSO) [40] resulting in 1000 to 1400 parameters. Second, automatic relevance detection regression (ARDR) [41] is used to fit the final model after tuning the hyperparameters to achieve the low-

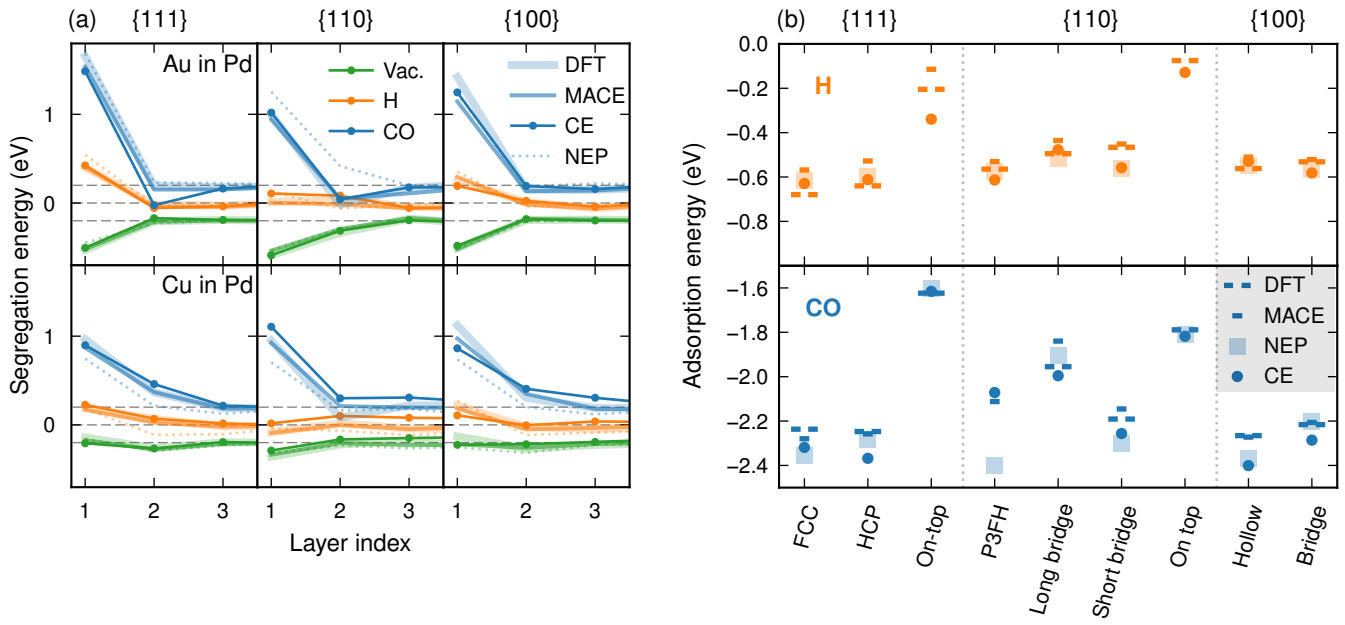

Figure 3: **Model performance for surface properties.** (a) Segregation energies from DFT, NEP, MACE, and CE for dilute Pd alloys in vacuum, H_2 and CO (100 % coverage). Negative (positive) values indicate that the minority species prefers (avoids) residing in the surface region. Note that the curves are shifted along the y-axis for better visualization. (b) Adsorption energies for H and CO on Pd (25 % coverage) for all models. Additional sites and adsorption energies on Cu and Au can be found in Fig. S3.

est number of features with a converged R2-score (Fig. S5). We use LASSO and ARDR as implemented in SCIKIT-LEARN [42] via the TRAINSTATION package [43].

Atomic structures for the training set are generated and selected using an active learning approach and relaxed using the MACE model, keeping the in-plane lattice parameter fixed based on interpolation of lattice parameters from relaxed bulk structures. An initial dataset, consisting of 300 structures with approximately orthogonal cluster vectors (following the method outlined in Ref. 31), is used to train an ensemble of CEs. A CE corresponding to the mean of the ensemble is used to run 300 short annealing MC simulations. The ensemble of CEs is used to calculate the uncertainty in energy across these trajectories, and the structure with the highest uncertainty for each trajectory is added to the training set. This process is repeated four times, at which point the root mean square error (RMSE) is no longer improving (Fig. S5).

For the final models, we also include the structures generated specifically for calculating the adsorption and segregation energies as well as similar sets of structures generated for adsorbate-free CEs in parallel, resulting in about 2800 structures for each orientation. The final models show excellent agreement with the MACE training data (Fig. 1b). For the {111} surface, the final model has 323 features (i.e., non-zero parameters) and achieves a validation RMSE of 3.8 meV/metal atom, for the {110} surface the final model has 281 features and achieves a validation RMSE of 6.4 meV/metal atom, and finally, for the {100} surface the final mode has 195 features and yields a validation RMSE of 3.7 meV/metal atom. The models are further validated by comparing the predicted segregation and adsorption energies with DFT and MLIP data (Fig. 3), which demostrates excellent agreement. Here, the segregation energy is calculated for a single Au (Cu) atom in different layers of a 2×2 Pd slab, compared to the energy of placing the minority atom in the bulk.

2.4. *Monte Carlo simulations*

To study the surface-adsorbate system under various conditions, MC simulations are performed using the MCHAMMER package [30]. The structures studied typically consist of $6{\times}6{\times}n_{\text{layers}}$ metal atoms and up to 1 (for {111} and {100}) or 2 (for {110}) monolayers of H or CO and the simulation runs for 100 MC cycles.

MC simulations of an adsorbate lattice need to account for the fact that many unfavorable configurations will not be represented well by the CE, since the adsorbates can move to more favorable configurations during relaxation. Especially when using the full adsorbate lattice (i.e., all adsorption sites), distances between neighboring sites are typically too short to allow for occupation of both sites in atomic relaxation as well as in reality. To account for this, we impose a minimal distance between occupied adsorbate sites (1.9 Å for {111} and {110}, 2.1 Å for {100} Fig. 2), which still allows for reaching the maximum coverage before absorption to the subsurface becomes favorable (Fig. S4). In addition, adsorbate sites not well represented by the CEs (bridge for {111}, hollow for {110} and on-top for {100}, see Fig. S3) are excluded in the MC simulations.

The MC simulations are performed in multiple steps to mimic the experimentally studied systems, which were annealed in a H_2 environment (corresponding to 40 mbar H partial pressure) at 773 K before the optical measurements were performed in ambient conditions and varying H_2 and CO partial pressures. First, the slab configuration is obtained from MC

simulations at 773 K and a varying H coverage. Here, separate canonical ensembles are used for the bulk, surface and adsorbate subsystems. The equilibrium composition of the surface region, in relation to the bulk and adsorbates, is found by interpolating results from preceding simulations where bulk and surface are treated as one subsystem. Then, the co-adsorption phenomena during the experimental measurements is simulated via MC simulations at 300 K. Due to the different timescales of adsorption and chemical reordering of the slab configuration, were the latter is much slower, we keep the slab configuration fixed (from the previous simulation) while the adsorbate sublattices equilibrates within a semi-grand canonical (SGC) ensemble controlled by the H and CO chemical potentials.

To prevent the adsorbate configuration from freezing in, swaps in the SGC ensemble are mixed with swaps in a canonical ensemble for the adsorbate sublattice, which enables swaps between two sites. This scheme enables moving an adsorbate from, e.g., a bridge site to a neighboring FCC site without having to first switch the site to vacuum which rarely will be an accepted move.

3. Results and discussion

3.1. Surface phase diagram construction

Traditionally, SPDs are constructed by calculating the free energy per primitive cell (or area) as a function of chemical potential at a few selected surface coverages and identifying the minimum energy phase. In Fig. 4, we show this process for the Pd {111} surface. In Fig. 4a, the free energy in the 0 K limit,

$$G(\theta_{\mathrm{H}}, \theta_{\mathrm{CO}} = 0) = E^{\mathrm{mix}}(\theta_{\mathrm{H}}, \theta_{\mathrm{CO}}) - \theta_{\mathrm{H}}\mu_{\mathrm{H}}, \tag{1}$$

is calculated from the CE mixing energy (from which the gas phase H and CO energies in relation to their respective coverages θ have already been subtracted) and the chemical potential. The convex hull indicates the stable adsorbate coverage in a certain chemical potential interval. By repeating this for multiple combinations of H and CO coverages, a two-dimensional SPD can be constructed (Fig. 4b). While this example uses a CE, it is straightforward to construct this kind of surface diagram using DFT calculations for pure metal surfaces and certain high-symmetry alloys (see Supp. Note 4, Fig. S6). For more complex system, however, the number of configurations to consider quickly becomes unmanageable. In addition, finite temperatures require treatment of the configurational entropy. Using a CE in combination with MC simulations allows for efficient thermodynamic sampling, including configurational entropy contributions, in the SGC ensemble where the chemical potentials are controlled, which results in a continuous SPD (Fig. 4c).

3.2. Pressure conversion

In most cases, the partial pressure p is more relevant than the chemical potential μ. If ideal gas behavior can be assumed,

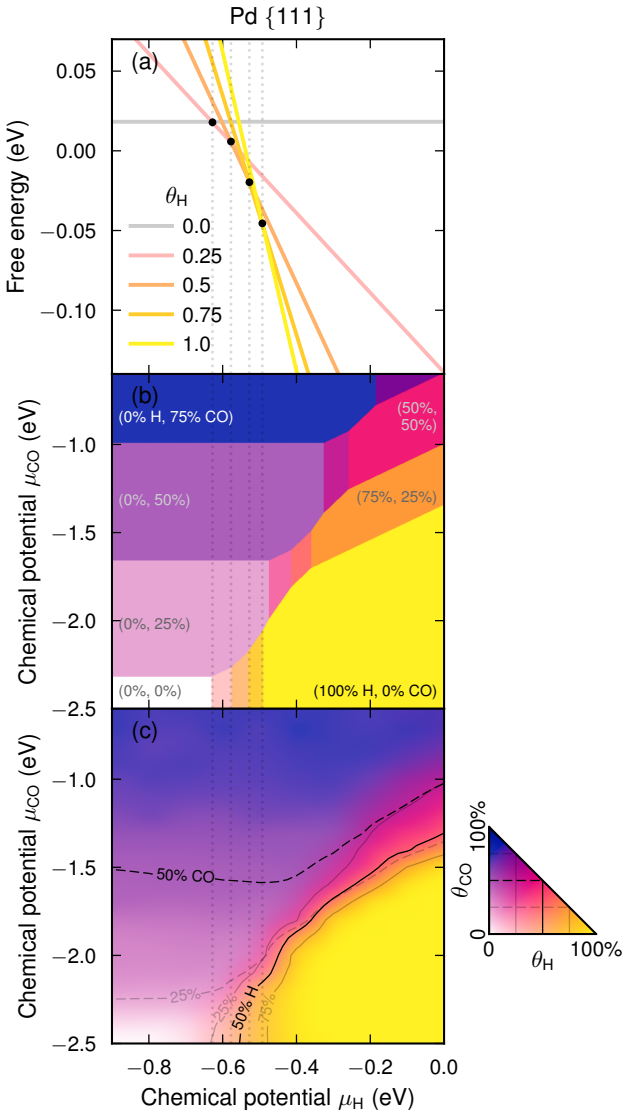

Figure 4: **SPD construction for Pd {111}.** Traditionally, SPDs are constructed by calculating the free energy as a function of the chemical potential for specific coverages (a) and translating the convex hull to a discrete surface diagram (b). MC sampling of CEs allows for efficient sampling of surface coverage(s) as a function of chemical potential(s) which results in continuous surface diagrams (c).

the two are connected via the relations

$$\mu_{\mathrm{H}}(T, p) = \frac{1}{2}\left(\mu^0_{\mathrm{H_2}}(T) + kT \ln \frac{p_{\mathrm{H_2}}}{p^0}\right) \tag{2}$$

$$\mu_{\mathrm{CO}}(T, p) = \mu^0_{\mathrm{CO}}(T) + kT \ln \frac{p_{\mathrm{CO}}}{p^0} \tag{3}$$

for H and CO, respectively. The reference chemical potential $\mu^0(T)$, defined at a reference partial pressure p^0, can, in principle, be obtained from thermodynamic tables [47]. This approach, however, suffers from the fact that adsorption energies are generally associated with a significant error in DFT calculations, leading to an effective shift of the chemical potential. In the context of our modeling framework, this approach suffers from the fact that the temperature dependence of the adsorbed

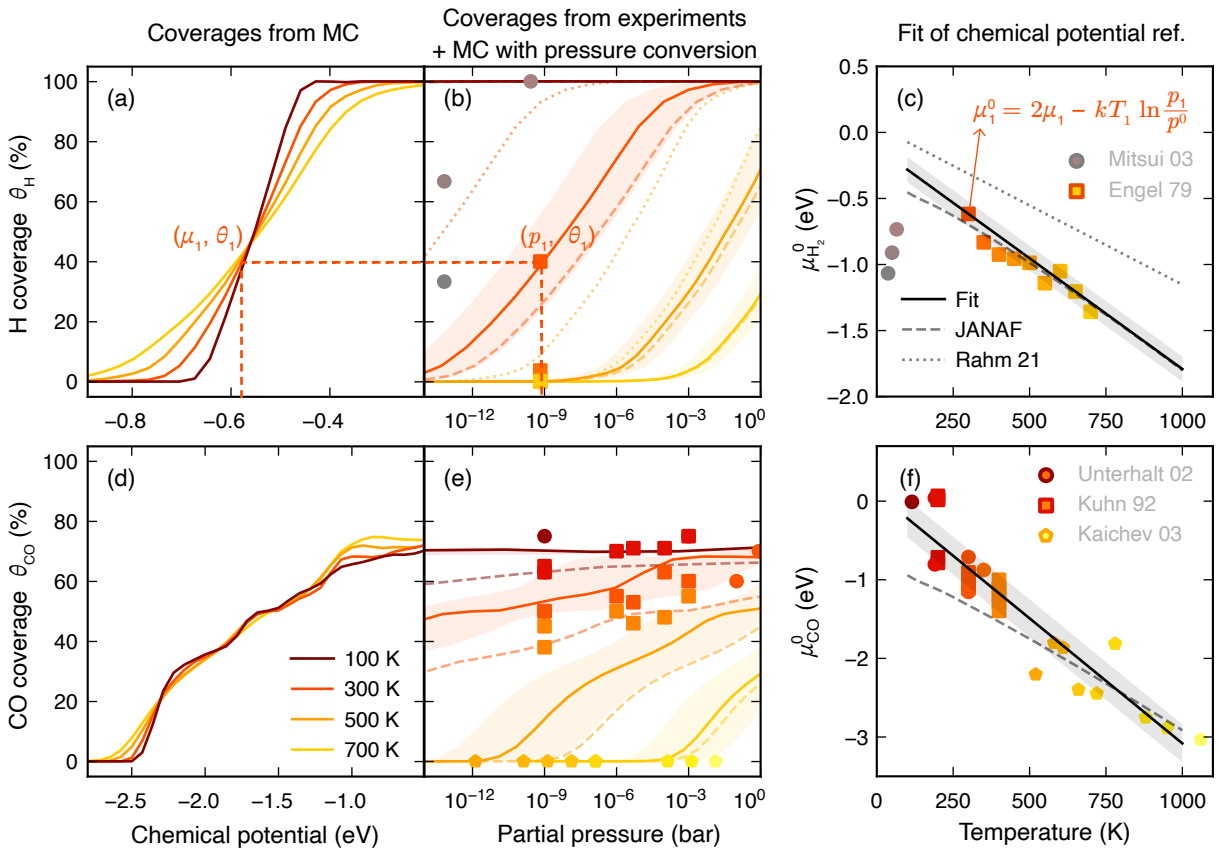

Figure 5: **Conversion between chemical potential and pressure for Pd.** (a) H coverage vs. chemical potential obtained from MC simulations. (b) Experimental records of H coverage vs. partial pressure from Mitsui 03 [44], Engel 79 [45], and additional sources (see Fig. S7) [46]. (c) Fit of the reference chemical potential based on the data in (a) and (b) as well as the calculated reference chemical potential based on tabulated thermodynamic data (JANAF [47]) or computational models (Rahm 21 [8]). Once the reference chemical potential is known, the MC data (a) can be converted from chemical potential to partial pressure, as shown in (b) for the different references. (d–f) show the same process for CO with experimental data for Pd:CO from Unterhalt 02 [48], Kuhn 92 [49], Kaichev 03 [50], and additional sources (see Fig. S7) [51, 52, 53]. Note that the legends are shared across each row.

molecules is limited to configurational entropy, while the full temperature dependence is accounted for with regard to the gas phase.

In some cases, one can find an observable that can be measured experimentally and calculated in the modeling framework and use this as a basis for calculating the reference. In the present work, the adsorbate coverage at a given temperature can be used for this purpose by linking the coverage as a function of chemical potential from MC simulations to experimental observations of coverage as a function of partial pressure. With this approach, the full temperature dependence is implicitly taken into account, but the availability of experimental data at relevant conditions as well as uncertainty in the measurements are potential issues.

In Fig. 5 we show our approach to pressure conversion. The H coverage on pure Pd is obtained as a function of the H chemical potential (Fig. 5a) from MC simulations. Experimental records of H coverage at known H partial pressure and temperature [45, 46, 44] (Fig. 5b) are then matched to the MC results and used to calculate μ^0, which can fitted by a simple temperature-dependent function. In Fig. 5c we show this (linear) fit as well as $\mu^0(T)$ calculated from the NIST-JANAF thermodynamic tables [47] and a fit based on the bulk hydride formation from Ref. 8. The shaded regions correspond to an offset equal to the RMSE of the linear fit in the 200 K to 500 K interval. The NIST-JANAF data is corrected by a constant shift of −0.2 eV based on the difference between the CE H adsorption energy (−0.63 eV) and experimental reports [54, 4], (∼ −0.43 eV). Except for the low-temperature limit, the experimental data is well represented. It should, however, be noted that most of the experimental data temperatures indicate very low (sub-percent) coverage with H diffusion to the bulk [45]. For this reason, the only data point with substantial coverage at a relevant temperature (highlighted in Fig. 5a–c) was given a larger weight in the fit. (The fitting procedure included additional sources of experimental data for other surface orientations, see Fig. S7.)

In Fig. 5d–f the analysis of the pressure conversion procedure is shown for CO. The MC simulations display a stair-like behavior for the CO coverage, indicating a preference for ordered phases at 33, 50 and 67 % CO coverage (see Fig. S8 for site-specific coverages) and saturation at 75 % in line with experimental observations [50]. CO adsorption on Pd has been well-studied for the past decades and experimental reports of

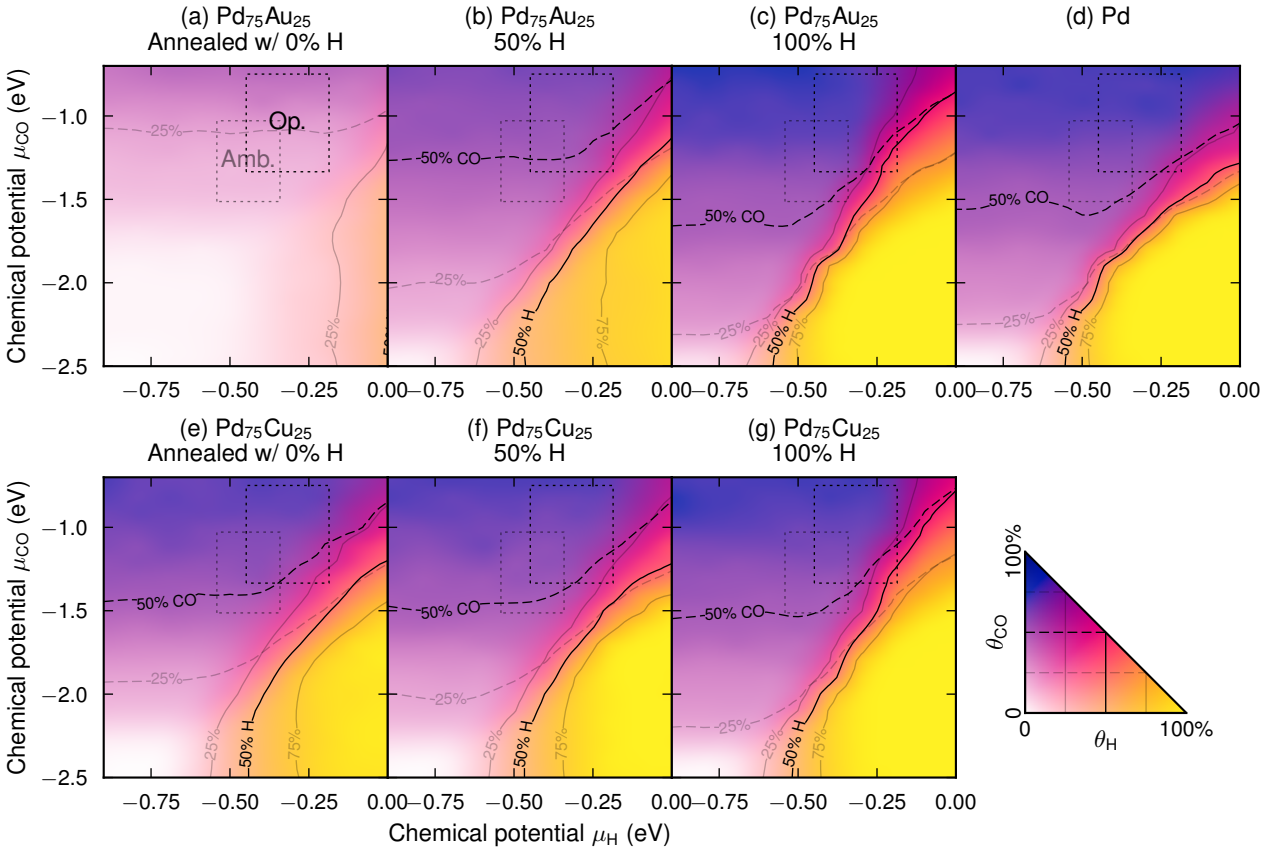

Figure 6: **SPD for Pd, $Pd_{75}Au_{25}$ and $Pd_{75}Cu_{25}$ prepared in different conditions.** The slab configurations were obtained from MC simulations at 773 K with (a) 0, (b) 50 or (c) 100 % H coverage and then covered with CO and H via MC simulations at 300 K while keeping the chemical configuration of the slab fixed. The surface coverage is represented by the color according to the colormap in the lower right corner and the contour lines, where the solid lines correspond to H and the dashed to CO. The dotted boxes indicate regions approximately corresponding to operating vs. ambient conditions.

the CO coverage over a wide range of pressures and temperatures are available [48, 49, 50, 51, 52, 53]. The calculated values for $\mu^0(T)$ have some spread, but overall seem to follow a linear trend that is captured by the fit. The calculation based on the NIST-JANAF thermodynamic tables [47] include a correction shift of −0.8 eV based on the approximated difference between the CE H adsorption energy (∼ −2.3 eV) and experimental reports [53, 55], (∼ −1.5 eV). The NIST-JANAF reference differs from the fitted function in slope and offset.

In Fig. 5b, e we show how the pressure conversion changes with the different approaches to calculating μ^0. Clearly, relatively small differences in μ^0 can lead to shifts of several orders of magnitude in the pressure. This finding highlights the inherent difficulty in accurate conversion between chemical potential and partial pressure, rather than a flaw in our methodology. To account for the uncertainty in partial pressure, we indicate relevant pressure regions in SPDs below instead of exact partial pressures. Due to the lack of experimental data for H, we use all three approaches to determining μ^0 to span the pressure region. For CO, we rely on the fit with error bars due to the well-corroborated experimental data and lacking treatment of vibrations in the NIST-JANAF approach.

3.3. Influence of preparation conditions

In the following we present SPDs for Pd-alloy surfaces and study how the preparation conditions affect the H and CO coadsorption. By preparation conditions we refer to the environment in which the slab configuration equilibrates, which in the present work is set up to mimic the experimental procedure of annealing the samples at 773 K in a H_2 environment [6, 7]. Due to the large uncertainty associated with the pressure conversion, we consider a range of fixed H coverages during annealing rather than specifying the H_2 partial pressure via the H chemical potential. The resulting compositions of the surface region are presented in Fig. S9 to S11.

We then analyze the coadsorption under operation, i.e. in conditions similar to the experimental measurements [6, 7], by constructing SPDs for the H and CO coverage as a function of their respective chemical potentials while keeping the slab configuration fixed. The regions corresponding to ambient (0.6 µbar H_2, 0.1 µbar CO [56]) and operating conditions (1 mbar to 100 mbar H_2 [57, 58], 0.1 mbar to 5 mbar CO [6, 7]) are indicated, where the width is associated with the uncertainty in partial pressure conversion rather than the considered pressure intervals.

Figure 6 shows SPDs for Pd, $Pd_{75}Au_{25}$, and $Pd_{75}Cu_{25}$

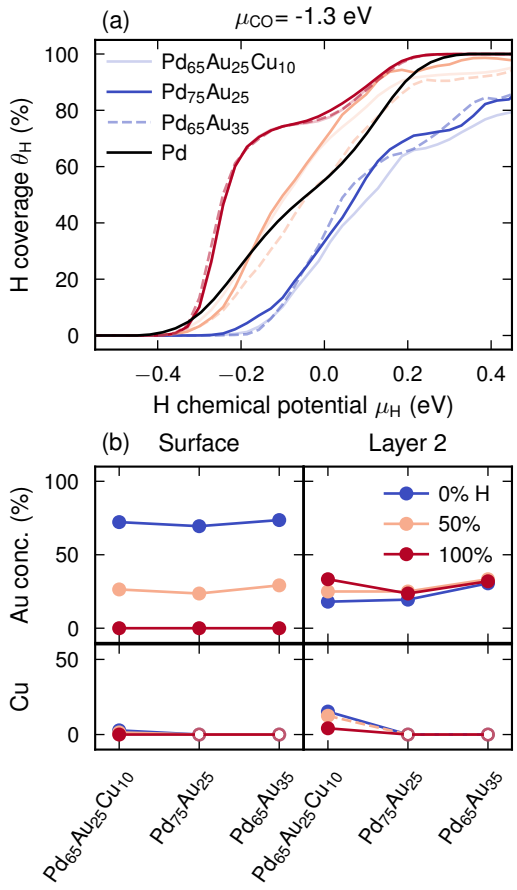

Figure 7: **H coverage and composition on Pd and PdAuCu-alloys.** (a) H coverage as a function of the H chemical potential μ_H at fixed CO chemical potential $\mu_{CO} = -1.3$ eV for Pd, $Pd_{65}Au_{25}Cu_{10}$, $Pd_{75}Au_{25}$ and $Pd_{65}Au_{35}$ obtained from MC simulations at 300 K and with the alloy configuration kept frozen. (b) Concentration of the surface layer and subsequent layer for the alloys, obtained from MC simulations at 773 K with varying H coverage as indicated by the legend (used also in a).

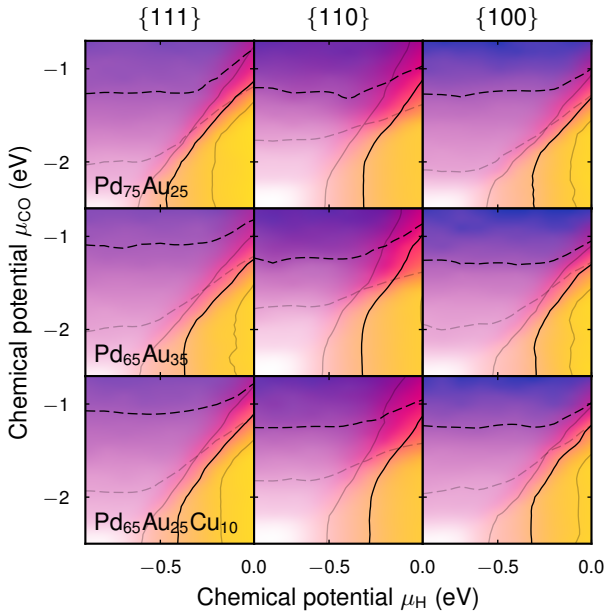

Figure 8: **Composition and orientation-dependent SPDs.** SPD for $Pd_{75}Au_{25}$, $Pd_{65}Au_{35}$ and $Pd_{65}Au_{35}Cu_{10}$ at varying H and CO chemical potentials obtained from MC simulations at 300 K with alloy configuration kept frozen. The slab configurations are annealed at 773 K with 50 % H coverage. The coverages follow the same colormap as Fig. 6, with solid and dashed contour lines for H and CO, respectively. Note that for {110}, 100 % coverage corresponds to 2 ML of adsorbates due to the large surface area per metal atom.

{111}-surfaces prepared by annealing at 773 K with 0 %, 50 % or 100 % H coverage. For $Pd_{75}Au_{25}$, the surface coverage clearly depends on the preparation conditions. Without adsorbed H during annealing, Au segregates to the surface (Fig. S9). This increases the adsorption energy of both CO and H, leading to low coverages (Fig. 6a) and, in particular, almost no H at the surface under operating conditions which would hinder the H_2 sensing ability. With 50 % H during annealing, the surface composition is similar to the bulk which leads to increased adsorption (Fig. 6b). Lastly, with 100 % H, the outermost surface layer is made up almost entirely of Pd which leads to a further increase of the overall coverages (Fig. 6c) resulting in a similar SPD as to the case of pure Pd (Fig. 6d). There is, however, a notable increase in H coverage in the lower right corner of the operating region, suggesting that $Pd_{75}Au_{25}$ with mostly Pd at the surface is superior to Pd in terms of CO adsorption blocking the surface.

For $Pd_{75}Cu_{25}$, the change in H and CO coverages with preparation conditions is less significant (Fig. 6e–g) and the SPD falls in between the 50 % and 100 % H (coverage during annealing) SPDs for $Pd_{75}Au_{25}$. This is because the surface segregation tendency of Cu in PdCu is much weaker than for Au in the case of PdAu, leading to a majority of Pd at the surface under all conditions (Fig. S9). Based on the H coverage contour lines in the vicinity of the operating region, $Pd_{75}Cu_{25}$ allows for slightly more H at the surface compared to Pd.

In summary, alloys containing Au are greatly influenced by the preparation conditions in terms of their ability to adsorb H. Figure 7 shows that the H adsorption as a function of the H chemical potential shifts by at least 0.2 eV with the annealing H coverage, corresponding to several orders of magnitude for the pressure (Fig. 5a–b). Figure 7 also indicates that the alloy composition plays a secondary role, as will be further discussed in the next section.

3.4. *Influence of alloy composition*

Figure 8a shows SPDs for $Pd_{75}Au_{25}$, $Pd_{65}Au_{35}$, and $Pd_{65}Au_{35}Cu_{10}$ prepared with 50 % H coverage (the upper right corner corresponds to Fig. 6b). As noted also in Fig. 7a, the CO and H coadsorption behavior is seemingly independent of the alloy composition. In particular, there is no distinct effect of introducing Cu compared to Au, which is unexpected given the strong experimental evidence for Cu, specifically, reducing CO poisoning [6, 7]. This (lack of a) trend remains across alloy compositions and preparation conditions (Fig. S16–S18).

Figure 9 shows how the 50 % H adsorption contour line (from the SPDs) shifts with alloy composition, which can be interpreted as an indicator for the resistance to CO poisoning

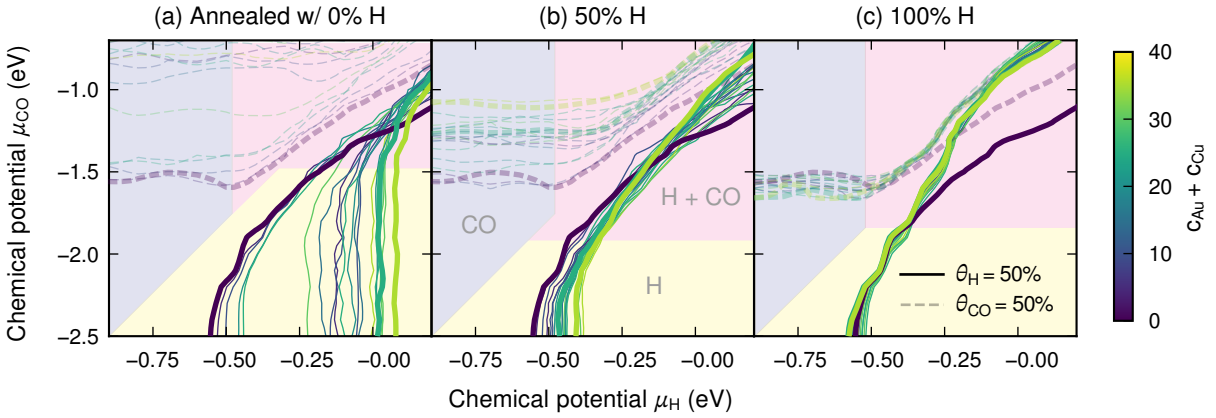

Figure 9: **50% coverage contour lines for Pd and various PdAuCu alloy {111} surfaces.** The slab configurations are obtained from MC simulations at 773 K with (a) 0 %, (b) 50 %, and (c) 100 % H coverage. The H and CO coverages are obtained from subsequent MC simulations at 300 K with fixed slab configuration. All alloy compositions considered are represented, with Pd, $Pd_{75}Au_{25}$, and $Pd_{65}Au_{25}Cu_{10}$ highlighted by thicker lines (see Fig. S9 for details concerning the compositions). The colored region indicate the limits of H, CO and H–CO adsorption.

(with respect to adsorption thermodynamics). Three regions can be identified: the H adsorption limit at the bottom, the CO adsorption limit to the left, and the H–CO coadsorption limit in the upper right corner. In the H adsorption limit (lower right quadrant)., Pd generally adsorbs more H than the alloys and the contour lines move to higher H chemical potentials with decreasing Pd content. A similar trend can be identified for CO, in the CO adsorption limit (upper left quadrant).

In the coadsorption limit (upper right quadrant), however, we see a shift in the trend where the 50 % H contour line for Pd crosses the corresponding contour lines for the alloys, meaning that the alloys adsorb more H at a given chemical potential. This suggests that in the coadsorption limit, alloying, in general, is beneficial for CO poisoning reduction. Depending on the preparation condition, this crossing of the contour lines happens inside (100 % H) or outside (0 % H) of the operating region ($\mu_H < -0.2$ eV). This suggests that a high H coverage during preparation is necessary to achieve the CO poisoning mitigating effects. We emphasize, however, again that there is a large uncertainty associated with the pressure conversion, and these predictions should be viewed as qualitative rather than quantitative.

For the systems prepared in 100 % H, the spread in the contour lines is generally narrow, which is reasonable since all systems have close to 100 % Pd at the surface (Fig. S9). In the coadsorption limit, however, all contour lines for the alloys distinctly deviate from the contour lines for Pd. This is surprising since their concentrations vary over a wide range *except* for the surface layer, which is nearly identical to the pure Pd case. It thus appears that any amount of alloyant in the bulk and subsurface region can distinctly change the adsorption behavior without entering the surface layer and without any dependence on the concentration.

While these findings explain why alloying might mitigate CO poisoning, they do not provide a rationale for the benefit of adding Cu, for which there is strong experimental evidence [6, 7]. This suggests that *adsorption thermodynamics* are not

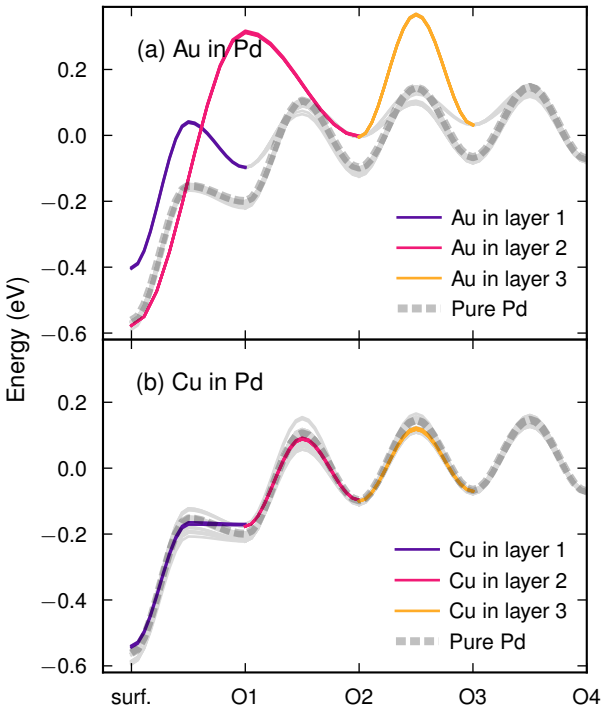

Figure 10: **Kinetic barriers for dilute Pd-alloys.** Migration barriers for a H atom moving from the surface into the material via octahedral absorption sites, for pure Pd (dashed) and structures with a single (a) Au or (b) Cu atom placed in one of the three outermost atomic layers i.e., the surface region). The paths closest to the Au/Cu atom are highlighted by the colors indicated by the legend. Note that for the case of Au in layer 2, the O1 site is not stable and the H atom relaxes to the surface.

enough to fully understand CO poisoning. Given the results above, we can hypothesize that under operating conditions, CO blocks the surface entirely for Pd and only partially for the alloys. The unoccupied areas will correspond to less favorable adsorption sites, i.e., less Pd than in the covered regions. It is then possible that the presence of Cu, especially in the second

atomic layer, might facilitate the sorption of H into the bulk. To investigate this, we look at migration paths and associated kinetic barriers via NEB calculations using the MACE model.

We consider {111} surfaces consisting of Pd with a single Au or Cu atom placed in atomic layer 1, 2, or 3, with a single H in the FCC adsorption site or one of the octahedral absorption sites, and mapped out all paths through the slab along with their associated kinetic barriers (Fig. 10). For PdAu, paths close to the Au atom are associated with a large increase in the energy of the initial and final states as well as the barrier. For PdCu, on the other hand, the migration energetics remain mostly unchanged or slightly improved compared to Pd. This suggests that the presence of Cu in the surface region could improve the *sorption kinetics*, especially in comparison to Au, and thereby enable H sorption in situations where the most favorable paths through the surface are blocked by CO.

At first glance, this negative impact of Au on H sorption appears to contradict prior studies, which demonstrate that Au enhances sorption kinetics in Pd-based alloys [5, 59]. We argue, however, that while Au *globally* improves sorption kinetics, through lattice expansion and a reduction in the total energy barrier for H sorption (mainly by increasing the adsorption energy) [59], it *locally* inhibits the kinetics for paths that passes an Au atom. This local effect becomes critical when the most energetically favorable H sorption paths are blocked by CO.

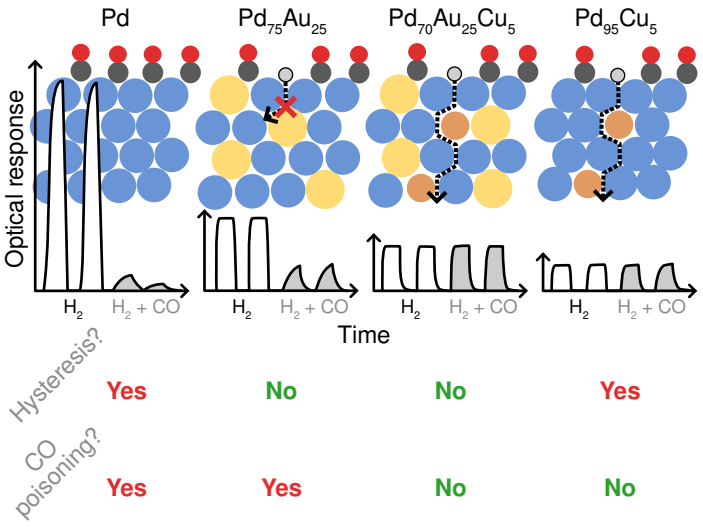

Figure 11: **Proposed mechanism for Cu-induced CO poisoning resistance.** The upper panel shows the (absolute) shift in optical response of Pd, $Pd_{75}Au_{25}$, $Pd_{70}Au_{25}Cu_5$, and $Pd_{95}Cu_5$ when subjected to pulses of H_2 (40 mbar) vs. H_2 (40 mbar) + CO (5 mbar), adapted from [6]. The slab representations show the corresponding proposed slab–CO–H interplay, where Pd is completely blocked by CO, $Pd_{75}Au_{25}$ suffers from poor kinetics due to Au in the surface region, while alloys with Cu are able to shuttle H into the material via Cu in the surface region. The lower panel summarizes the corresponding hysteresis and CO poisoning characteristics.

3.5. Influence of surface orientation

In the discussion above, we have mainly focused on the {111} surface, since it is the minimum energy surface for all three alloyants [60]. Figure 8 shows SPDs for the {110} and {100}, and additional results for the other surface orientation are available in the Supporting Information (Fig. S10 – S19). There are some minor changes of the SPDs between the orientations, most notable for {110} compared to the other surfaces, but overall the conclusions from the previous sections hold, namely that the coadsorption behavior depends primarily on the preparation conditions.

Some orientation-dependent effects can be pointed out. {110} alloy surfaces show less benefit compared to Pd in the coadsorption limit, in terms of the ability to adsorb H (Fig. S19). In addition, the surface segregation behavior of {110} PdCu, specifically, differs from the other orientations in so far that Cu segregates to the surface in vacuum (Fig. S10) to a much larger degree compared to the other orientations. {100} alloy surfaces display very similar trends to {111} with a slightly larger benefit compared to pure Pd in the coadsorption limit (Fig. S19). In the lower right corner of what we consider operation conditions ($\mu_H = -0.2$ eV, $\mu_{CO} = -1.3$ eV), we find an increase in H adsorption with Cu concentration for the alloy prepared in 100 % H which is the only clear example of Cu positively impacting CO poisoning characteristics (Fig. S13).

4. Conclusions

In the present work, we have studied the coadsorption of H and CO on PdAuCu alloy surfaces with 0 % to 35 % Au and 0 % to 25 % Cu to understand the mitigating effect on CO poisoning of such alloys, compared to Pd, in the context of H_2 sensing. We have found that by tuning the surface composition via the preparation conditions, specifically the H coverage, the adsorption of H and CO during operation can be tuned. Increasing the amount of Au at the surface, which can be achieved by preparing the surfaces at low H coverages, significantly reduces the CO adsorption tendency. Unfortunately, this also reduces the H adsorption such that higher H partial pressures are necessary to adsorb a comparable amount of H. On the other hand, when H coverage is high during preparation, one obtains Pd-rich surfaces that are prone to adsorb both H and CO, but with a higher H to CO ratio than pure Pd. This suggests that PdAuCu alloys prepared in a H-rich environment lead to improved CO poisoning resistance based on *adsorption thermodynamics*.

Adsorption thermodynamics alone are, however, not sufficient to explain the experimentally observed mitigating effect of Cu on CO poisoning [6, 7]. The experimentally observed optical shifts in response to H_2 vs. a mixture of H_2 and CO (Fig. 11), reveal that the primary difference between the samples during CO exposure is not the absolute H sorption, but rather the relative H sorption compared to the CO-free case and, importantly, the *sorption kinetics*.

With regard to the kinetic aspect, we have shown that in the dilute limit, introducing Au in the surface region of a Pd slab results in a significant increase in the kinetic barriers associated with H absorption, while Cu does not affect the H absorption energetics compared to pure Pd. Based on this finding, we suggest the following rationale for the observed CO poisoning behavior of the different systems (Fig. 11). For pure Pd, the

surface is almost completely blocked by CO, resulting in slow sorption of small amounts of H. For the alloys, the surface is only partially covered by CO, leading to some H adsorption. Crucially, while Au in the near-surface hinders H sorption, Cu has no pronounced effect on the sorption kinetics compared to Pd. We thus suggest that the introduction of Cu creates viable pathways for H sorption in cases where the most favorable pathways are blocked by CO at the surface.

Further studies of the bulk absorption process in realistic systems are necessary to further support the suggested mechanism regarding the Cu-induced improvement in H sorption kinetics. The MLIPs developed in the present work for the PdAuCu:(H,CO) system provide an ideal starting point for this endeavor.

Acknowledgments

We gratefully acknowledge funding from the Swedish Research Council (Nos. 2020-04935 and 2021-05072) and the Area of Advance at Nano at Chalmers as well as computational resources provided by the National Academic Infrastructure for Supercomputing in Sweden at NSC, PDC, and C3SE partially funded by the Swedish Research Council through grant agreement No. 2022-06725, as well as the Berzelius resource provided by the Knut and Alice Wallenberg Foundation at NSC, as well as the Berzelius resource provided by the Knut and Alice Wallenberg Foundation at NSC.

Data Availability

The NEP and MACE models as well as the database of DFT calculations used to train these models have been deposited in the Zenodo database under accession code 10.5281/zenodo.17670908.

References

[1] O. M. Løvvik, S. M. Opalka, Reversed surface segregation in palladium-silver alloys due to hydrogen adsorption, Surface Science 602 (17) (2008) 2840–2844. doi:10.1016/j.susc.2008.07.016.

[2] P. Ekborg-Tanner, P. Erhart, Hydrogen-Driven Surface Segregation in Pd Alloys from Atomic-Scale Simulations, The Journal of Physical Chemistry C 125 (31) (2021) 17248–17260. doi:10.1021/acs.jpcc.1c00575.

[3] C. Langhammer, E. M. Larsson, B. Kasemo, I. Zorić, Indirect Nanoplasmonic Sensing: Ultrasensitive Experimental Platform for Nanomaterials Science and Optical Nanocalorimetry, Nano Letters 10 (9) (2010) 3529–3538. doi:10.1021/nl101727b.

[4] J. Wellendorff, T. L. Silbaugh, D. Garcia-Pintos, J. K. Nørskov, T. Bligaard, F. Studt, C. T. Campbell, A benchmark database for adsorption bond energies to transition metal surfaces and comparison to selected DFT functionals, Surface Science 640 (2015) 36–44. doi:10.1016/j.susc.2015.03.023.

[5] C. Wadell, F. A. A. Nugroho, E. Lidström, B. Iandolo, J. B. Wagner, C. Langhammer, Hysteresis-Free Nanoplasmonic Pd-Au Alloy Hydrogen Sensors, Nano Letters 15 (5) (2015) 3563–3570. doi:10.1021/acs.nanolett.5b01053.

[6] I. Darmadi, F. A. A. Nugroho, S. Kadkhodazadeh, J. B. Wagner, C. Langhammer, Rationally Designed PdAuCu Ternary Alloy Nanoparticles for Intrinsically Deactivation-Resistant Ultrafast Plasmonic Hydrogen Sensing, ACS Sensors 4 (5) (2019) 1424–1432. doi:10.1021/acssensors.9b00610.

[7] I. Darmadi, S. Z. Khairunnisa, D. Tomeček, C. Langhammer, Optimization of the composition of PdAuCu ternary alloy nanoparticles for plasmonic hydrogen sensing, ACS Applied Nano Materials 4 (2021) 8716. doi:10.1021/acsanm.1c01242.

[8] J. M. Rahm, J. Löfgren, E. Fransson, P. Erhart, A tale of two phase diagrams: Interplay of ordering and hydrogen uptake in Pd–Au–H, Acta Materialia 211 (2021) 116893. doi:10.1016/j.actamat.2021.116893.

[9] J. M. Sanchez, F. Ducastelle, D. Gratias, Generalized cluster description of multicomponent systems, Physica 128 (1-2) (1984) 334.

[10] D. C. Beaudry, M. J. Waters, G. M. Valentino, D. L. Foley, E. Anber, Y. Rakita, C. J. Brandenburg, J.-P. Couzinié, L. Perrière, T. Aoki, K. E. Knipling, P. G. Callahan, B. W. Redemann, T. M. McQueen, E. J. Opila, J. M. Rondinelli, M. L. Taheri, Exceptional hardness in multi-principal element alloys via hierarchical oxygen heterogeneities, Science Advances 10 (38) (2024) eado9697. doi:10.1126/sciadv.ado9697.

[11] S. Müller, M. Stöhr, O. Wieckhorst, Structure and stability of binary alloy surfaces: Segregation, relaxation, and ordering from first-principles calculations, Applied Physics A 82 (3) (2006) 415–419. doi:10.1007/s00339-005-3362-6.

[12] T. Mueller, G. Ceder, Bayesian approach to cluster expansions, Physical Review B 80 (2) (2009) 024103. doi:10.1103/PhysRevB.80.024103.

[13] J. Zhang, X. Liu, S. Bi, J. Yin, G. Zhang, M. Eisenbach, Robust data-driven approach for predicting the configurational energy of high entropy alloys, Materials & Design 185 (2020) 108247. doi:10.1016/j.matdes.2019.108247.

[14] Y. L. Müller, A. R. Natarajan, Constructing multicomponent cluster expansions with machine-learning and chemical embedding, npj Computational Materials 11 (1) (2025) 60. doi:10.1038/s41524-025-01543-3.

[15] M. F. Shojaei, J. Holber, S. Das, G. H. Teichert, T. Mueller, L. Hung, V. Gavini, K. Garikipati, Bridging scales with Machine Learning: From first principles statistical mechanics to continuum phase field computations to study order–disorder transitions in Li_xCoO_2, Journal of the Mechanics and Physics of Solids 190 (2024) 105726. doi:10.1016/j.jmps.2024.105726.

[16] B. Sadigh, P. Erhart, A. Stukowski, A. Caro, E. Martinez, L. Zepeda-Ruiz, Scalable parallel Monte Carlo algorithm for atomistic simulations of precipitation in alloys, Physical Review B 85 (18) (2012) 184203. doi:10.1103/PhysRevB.85.184203.

[17] K. Song, R. Zhao, J. Liu, Y. Wang, E. Lindgren, Y. Wang, S. Chen, K. Xu, T. Liang, P. Ying, N. Xu, Z. Zhao, J. Shi, J. Wang, S. Lyu, Z. Zeng, S. Liang, H. Dong, L. Sun, Y. Chen, Z. Zhang, W. Guo, P. Qian, J. Sun, P. Erhart, T. Ala-Nissila, Y. Su, Z. Fan, General-purpose machine-learned potential for 16 elemental metals and their alloys, Nature Communications 15 (2024) 10208. doi:10.1038/s41467-024-54554-x.

[18] Z. Fan, Y. Wang, P. Ying, K. Song, J. Wang, Y. Wang, Z. Zeng, K. Xu, E. Lindgren, J. M. Rahm, A. J. Gabourie, J. Liu, H. Dong, J. Wu, Y. Chen, Z. Zhong, J. Sun, P. Erhart, Y. Su, T. Ala-Nissila, Gpumd: A package for constructing accurate machine-learned potentials and performing highly efficient atomistic simulations, Journal of Chemical Physics 157 (2022) 114801. doi:10.1063/5.0106617.

[19] D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters, J. Schmidhuber, Natural evolution strategies, Journal of Machine Learning Research 15 (27) (2014) 949–980. URL http://jmlr.org/papers/v15/wierstra14a.html

[20] E. Fransson, J. Wiktor, P. Erhart, Phase transitions in inorganic halide perovskites from machine-learned potentials, The Journal of Physical Chemistry C 127 (2023) 13773. doi:10.1021/acs.jpcc.3c01542.

[21] K. Xu, H. Bu, S. Pan, E. Lindgren, Y. Wu, Y. Wang, J. Liu, K. Song, B. Xu, Y. Li, T. Hainer, L. Svensson, J. Wiktor, R. Zhao, H. Huang, C. Qian, S. Zhang, Z. Zeng, B. Zhang, B. Tang, Y. Xiao, Z. Yan, J. Shi, Z. Liang, J. Wang, T. Liang, S. Cao, Y. Wang, P. Ying, N. Xu, C. Chen, Y. Zhang, Z. Chen, X. Wu, W. Jiang, E. Berger, Y. Li, S. Chen, A. J. Gabourie, H. Dong, S. Xiong, N. Wei, Y. Chen, J. Xu, F. Ding, Z. Sun, T. Ala-Nissila, A. Harju, J. Zheng, P. Guan, P. Erhart, J. Sun, W. Ouyang, Y. Su, Z. Fan, Gpumd 4.0: A high-performance molecular dynamics package for versatile materials simulations with machine-learned potentials, Materials Genome Engineering Advances 3 (2025) e70028. doi:10.1002/mgea.70028.

[22] E. Lindgren, J. M. Rahm, E. Fransson, F. Eriksson, N. Österbacka, Z. Fan, P. Erhart, calorine: A python package for constructing and sampling neuroevolution potential models, Journal of Open Source Software 9 (2024) 6264. doi:10.21105/joss.06264.

[23] G. L. W. Hart, R. W. Forcade, Algorithm for generating derivative structures, Physical Review B 77 (2008) 224115. doi:10.1103/PhysRevB.77.224115.

[24] G. L. Hart, R. W. Forcade, Generating derivative structures from multilattices: Algorithm and application to hcp alloys, Physical Review B 80 (1) (2009) 014120. doi:10.1103/PhysRevB.80.014120.

[25] G. Henkelman, H. Jonsson, Improved tangent estimate in the neb method for finding minimum energy paths and saddle points, Journal of Chemical Physics 113 (2000) 9978–9985. doi:10.1063/1.1323224.

[26] G. Henkelman, B. P. Uberuaga, H. Jonsson, A climbing-image neb method for finding saddle points and minimum energy paths, Journal of Chemical Physics 113 (2000) 9901–9904. doi:10.1063/1.1329672.

[27] P. Lindgren, G. Kastlunger, A. A. Peterson, Scaled and dynamic optimizations of nudged elastic bands, Journal of Chemical Theory and Computation 15 (11) (2019) 5787–5793. doi:10.1021/acs.jctc.9b00633.

[28] A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. Bjerre Jensen, J. Kermode, J. R. Kitchin, E. Leonhard Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. Bergmann Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, K. W. Jacobsen, The atomic simulation environment—a Python library for working with atoms, Journal of Physics: Condensed Matter 29 (27) (2017) 273002. doi:10.1088/1361-648X/aa680e.

[29] F. Eriksson, E. Fransson, P. Erhart, The hiphive package for the extraction of high-order force constants by machine learning, Advanced Theory and Simulations 2 (2019) 1800184. doi:10.1002/adts.201800184.

[30] M. Ångqvist, W. A. Muñoz, J. M. Rahm, E. Fransson, C. Durniak, P. Rozyczko, T. H. Rod, P. Erhart, ICET – A python library for constructing and sampling alloy cluster expansions, Advanced Theory and Simulations 0 (2) (2019) 1900015. doi:10.1002/adts.201900015.

[31] P. Ekborg-Tanner, P. Rosander, E. Fransson, P. Erhart, Construction and Sampling of Alloy Cluster Expansions—A Tutorial, PRX Energy 3 (4) (2024) 042001. doi:10.1103/PRXEnergy.3.042001.

[32] I. Batatia, D. P. Kovacs, G. N. C. Simm, C. Ortner, G. Csanyi, MACE: Higher order equivariant message passing neural networks for fast and accurate force fields, in: A. H. Oh, A. Agarwal, D. Belgrave, K. Cho (Eds.), Advances in Neural Information Processing Systems, Curran Associates, Inc., 2022, arXiv:2206.07697. doi:10.48550/arXiv.2206.07697.

[33] I. Batatia, S. Batzner, D. P. Kovács, A. Musaelian, G. N. C. Simm, R. Drautz, C. Ortner, B. Kozinsky, G. Csányi, The design space of E(3)-equivariant atom-centered interatomic potentials (2022). arXiv:2205.06643, doi:10.48550/arXiv.2205.06643.

[34] M. Dion, H. Rydberg, E. Schröder, D. C. Langreth, B. I. Lundqvist, Van der waals density functional for general geometries, Physical Review Letters 92 (2004) 246401. doi:10.1103/PhysRevLett.92.246401.

[35] K. Berland, P. Hyldgaard, Exchange functional that tests the robustness of the plasmon description of the van der Waals density functional, Physical Review B 89 (2014) 035412. doi:10.1103/PhysRevB.89.035412.

[36] G. Kresse, J. Furthmüller, Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set, Computational Materials Science 6 (1) (1996) 15–50. doi:10.1016/0927-0256(96)00008-0.

[37] G. Kresse, J. Furthmüller, Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set, Physical Review B 54 (16) (1996) 11169–11186. doi:10.1103/PhysRevB.54.11169.

[38] P. E. Blöchl, Projector augmented-wave method, Physical Review B 50 (24) (1994) 17953–17979. doi:10.1103/PhysRevB.50.17953.

[39] G. Kresse, D. Joubert, From ultrasoft pseudopotentials to the projector augmented-wave method, Physical Review B 59 (3) (1999) 1758–1775. doi:10.1103/PhysRevB.59.1758.

[40] R. Tibshirani, Regression shrinkage and selection via the Lasso, Journal of the Royal Statistical Society: Series B (Methodological) 58 (1) (1996) 267–288. doi:10.1111/j.2517-6161.1996.tb02080.x.

[41] D. J. C. MacKay, Bayesian Non-linear Modelling for the Prediction Competition, ASHRAE Transactions, V.100, Pt.2 (1994) 1053–1062.

[42] 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, E. Duchesnay, Scikit-learn: Machine learning in python, Journal of Machine Learning Research 12 (2011) 2825–2830.

[43] E. Fransson, F. Eriksson, P. Erhart, Efficient construction of linear models in materials modeling and applications to force constant expansions, npj Computational Materials 6 (2020) 135. doi:10.1038/s41524-020-00404-5.

[44] T. Mitsui, M. K. Rose, E. Fomin, D. F. Ogletree, M. Salmeron, Dissociative hydrogen adsorption on palladium requires aggregates of three or more vacancies, Nature 422 (6933) (2003) 705–707. doi:10.1038/nature01557.

[45] T. Engel, H. Kuipers, A molecular-beam investigation of the scattering, adsorption and absorption of H_2 and D_2 from/on/in Pd(111), Surface Science 90 (1) (1979) 162–180. doi:10.1016/0039-6028(79)90017-7.

[46] M. Wilde, M. Matsumoto, K. Fukutani, T. Aruga, Depth-resolved analysis of subsurface hydrogen absorbed by Pd(100), Surface Science 482-485 (2001) 346–352. doi:10.1016/S0039-6028(01)00727-0.

[47] M. W. Chase, N. I. of Standards, Technology, NIST-JANAF thermochemical tables, fourth edition, https://janaf.nist.gov/, NIST Standard Reference Database Number 13 (1998). doi:10.18434/T42S31.

[48] H. Unterhalt, G. Rupprechter, H.-J. Freund, Vibrational Sum Frequency Spectroscopy on Pd(111) and Supported Pd Nanoparticles: CO Adsorption from Ultrahigh Vacuum to Atmospheric Pressure, The Journal of Physical Chemistry B 106 (2) (2002) 356–367. doi:10.1021/jp013000+.

[49] W. K. Kuhn, J. Szanyi, D. W. Goodman, CO adsorption on Pd(111): the effects of temperature and pressure, Surface Science 274 (3) (1992) L611–L618. doi:10.1016/0039-6028(92)90834-S.

[50] V. V. Kaichev, I. P. Prosvirin, V. I. Bukhtiyarov, H. Unterhalt, G. Rupprechter, H.-J. Freund, High-Pressure Studies of CO Adsorption on Pd(111) by X-ray Photoelectron Spectroscopy and Sum-Frequency Generation, The Journal of Physical Chemistry B 107 (15) (2003) 3522–3527. doi:10.1021/jp021992t.

[51] J. W. He, U. Memmert, K. Griffiths, W. N. Lennard, P. R. Norton, Absolute coverages for the various surface phases of CO and O adsorbed on Pd(110), Surface Science 202 (3) (1988) L555–L558. doi:10.1016/0039-6028(88)90031-3.

[52] J. He, P. R. Norton, Interaction of CO with a Pd(110) surface, studied by low energy electron diffraction, thermal desorption spectroscopy, and $\Delta\Phi$, The Journal of Chemical Physics 89 (2) (1988) 1170–1176. doi:10.1063/1.455225.

[53] R. J. Behm, K. Christmann, G. Ertl, M. A. Van Hove, Adsorption of CO on Pd(100), The Journal of Chemical Physics 73 (6) (1980) 2984–2995. doi:10.1063/1.440430.

[54] O. M. Løvvik, R. A. Olsen, Adsorption energies and ordered structures of hydrogen on Pd(111) from density-functional periodic calculations, Physical Review B 58 (16) (1998) 10890–10898. doi:10.1103/PhysRevB.58.10890.

[55] J. Szanyi, W. K. Kuhn, D. W. Goodman, CO adsorption on Pd(111) and Pd(100): Low and high pressure correlations, Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films 11 (4) (1993) 1969–1974. doi:10.1116/1.578532.

[56] N. N. W. Service, Jetstream: An online school for weather, https://www.noaa.gov/jetstream/atmosphere, accessed: 2025-02-15 (2025).

[57] U.S. Department of Energy, Energy Efficiency and Renewable Energy (EERE), Fuel Cell Technologies Office, Multi-Year Research, Development, and Demonstration Plan, 2011–2020: Section 3.7 Hydrogen Safety, Codes and Standards (2015).

[58] I. Darmadi, F. A. A. Nugroho, C. Langhammer, High-Performance Nanostructured Palladium-Based Hydrogen Sensors-Current Limitations and Strategies for Their Mitigation, ACS Sensors 5 (11) (2020) 3306–3327. doi:10.1021/acssensors.0c02019.

[59] K. Namba, S. Ogura, S. Ohno, W. Di, K. Kato, M. Wilde, I. Pletikosić, P. Pervan, M. Milun, K. Fukutani, Acceleration of hydrogen absorption by palladium through surface alloying with gold, Proceedings of the National Academy of Sciences 115 (31) (2018) 7896–7900. doi:10.1073/pnas.1800412115.

[60] R. Tran, Z. Xu, B. Radhakrishnan, D. Winston, W. Sun, K. A. Persson, S. P. Ong, Surface energies of elemental crystals, Scientific Data 3 (1) (Sep. 2016). doi:10.1038/sdata.2016.80.

Supporting Information

Competing Adsorption of H and CO on Pd-Alloy Surfaces: Mechanistic Insight into the Mitigating Effect of Cu on CO Poisoning

Pernilla Ekborg-Tanner[1] and Paul Erhart[*,1]

[1] Department of Physics, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden
[*] erhart@chalmers.se

Contents

Supplementary Notes

Supplementary Note 1: Construction of neuroevolution potential models

We constructed machine-learned interatomic potentials (MLIPs) based on the fourth-generation neuroevolution potential (NEP) framework (1, 2) using the iterative procedure described in Ref. 3 utilizing the GPUMD (1, 4) and CALORINE packages (5). For the NEP models, radial and angular cutoffs were set to 7 Å and 4 Å, respectively, and angular descriptors included both three- and four-body components, corresponding to $l^3_{\max} = 4$ and $l^4_{\max} = 2$. The neural network architecture consisted of 44 descriptor nodes in the input layer (9 radial, 35 angular) and one hidden layer with 40 fully connected neurons. Training was carried out using the separable neuroevolution strategy (6).

The final set of models was trained for 4×10^5 generations using the separable neuroevolution strategy (6), after which the loss, based on the root mean square errors (RMSEs), was deemed converged (Fig. S1). We trained one model using the complete set of reference data ("full") as well as five models using 90% and 10% of the reference data for training and validation, respectively. ("splits"). During this stage, we used weights of 10, 1, and 0.1 for the loss terms associated with energies, forces, and virials. In the case of the "full" model, we then added a second stage during which we used weights of 5, 5, and 0.1 for energies, forces, and virials, leading to a marked improvement in the energies at a very slight deterioration of the accuracy of the forces. The coefficient of determination R^2, RMSE, and mean absolute error (MAE) for the prediction of energies, forces, and stresses are shown in Table S1 for both "full" model as well as the "split" models. The latter models can be used as an ensemble to obtain uncertainty estimates.

The NEP models constructed here yield a computational speed of up to 8.7×10^6 atom step s^{-1} (equivalent to a cost of 0.1×10^{-6} s atom^{-1} step) on Nvidia A100 cards (Fig. S2). For illustration, using a time step of 0.5 fs, this allows one to simulate up to 9.4 ns per day and GPU for a system size of 40 000 atoms.

Supplementary Note 2: Construction of MACE models

MACE models (7) were constructed using the MACE-TORCH package (8). The architecture employs a message-passing neural network built with 64 scalar channels and `max_L`= 0 corresponding to a hidden representation of 64×0e. The models thus comprise two interaction layers, each operating with correlation order 3, which corresponds to an effective body order of 4, and using spherical harmonics up to $\ell = 3$. The radial representation is constructed from eight radial functions combined with five basis functions, and interactions are restricted to a radial cutoff of 6 Å, giving each atom a 12 Å receptive field. No additional distance transformation is applied to the radial basis. Altogether, the architecture contains 171 472 trainable parameters.

Training used the ADAM optimizer with a batch size of 100 and included an exponential moving average of the parameters with decay 0.99. The optimization started with a learning rate of 0.01 and a weight decay of 5×10^{-7}. During the initial training stage, the loss function weighted energies and forces with factors of 1 and 100, respectively. After 200 epochs, the procedure switched to a second stage that strongly emphasized energetic accuracy, employing a loss with weights of 10 000 for energies and 550 for forces, together with a reduced learning rate of 0.001. Training proceeded for in total 1000 epochs.

Using the above parameters we trained five models ("splits") using five different shuffle-splits of the reference set using a 90% and 10% of the data for training and validation, respectively. The coefficient of determination R^2, RMSE, and MAE for the prediction of energies, forces, and stresses are shown in Table S2 for all five models. The model "split 1" was used for generating reference data for the construction of cluster expansion (CE) models. The metrics split by category and structure type for the "split 1" model are shown in Table S4.

The MACE models constructed here yield a computational speed of up to 0.018×10^6 atom step s^{-1} (equivalent to a cost of 55×10^{-6} s atom^{-1} step) on Nvidia A100 cards (Fig. S2). For illustration, using a time step of 0.5 fs, this allows one to simulate up to 20 ps per day and GPU for a system size of 40 000 atoms.

Supplementary Note 3: Construction of reference structure set

Structures were generated and manipulated using the ASE (9) and HIPHIVE packages (10). The initial set of structures comprised a range of bulk, surface, and molecular structures relaxed at the density-functional theory (DFT) level as well as structures with random displacements ("rattled structures"). To include information about the mixed systems, we generated bulk and surface structures obtained by systematic enumeration of decorations of the metal lattice using the procedure described in Ref. 11, 12 as implemented in the ICET package (13). In this step, we applied a linear interpolation of the lattice parameters between the end members to account for chemical expansion, but did not relax the structures.

The dataset was then augmented with structures from several iterations of active learning using NEP models (Supplementary Note 1). To this end, we trained an ensemble of five NEP models by randomly splitting the

available data into training and validation sets. The ensemble was used to estimate the model error. Molecular dynamics (MD) simulations were then carried out using the GPUMD package (1, 4) at a range of temperatures and pressures, considering compositions over the entire concentration range. The ensemble was used to select structures with high uncertainty, quantified by the standard deviations of energies and forces over the ensemble. Subsequently, we computed energies, forces, and stresses via DFT for these configurations and included them when training the next-generation NEP model. During this procedure, we also added information about saddle point configurations for hydrogen migration that were generated by interpolation as well as using the nudged elastic band method (14–16). Both the unrelaxed and relaxed midpoint configurations were added to the reference set. The final reference set comprised 18 403 structures, corresponding to a total of 586 006 atoms (see Table S3 for an overview).

Supplementary Note 4: Surface phase diagram construction

To construct surface phase diagrams using DFT, one has to carefully consider which atomic configurations to consider. For the case of pure Pd, we use $1 \times 1 \times 10$, $1 \times 2 \times 10$ or $2 \times 2 \times 10$ {111} surface slabs, depending on the coverage. We choose the smallest possible cell for each coverage and only a single calculation for each coverage (neglecting any ordering effects).

For the alloys, we use $3 \times 3 \times 10$ slabs with atomic configurations obtained from MC simulated annealing where the temperature is incrementally decreased down to 100 K with 0% or 100% H coverage. For the top 3 layers, the concentrations are restricted to even thirds to promote ordering while the bulk is fixed at 25% Au. This procedure results in the following ordered structures:

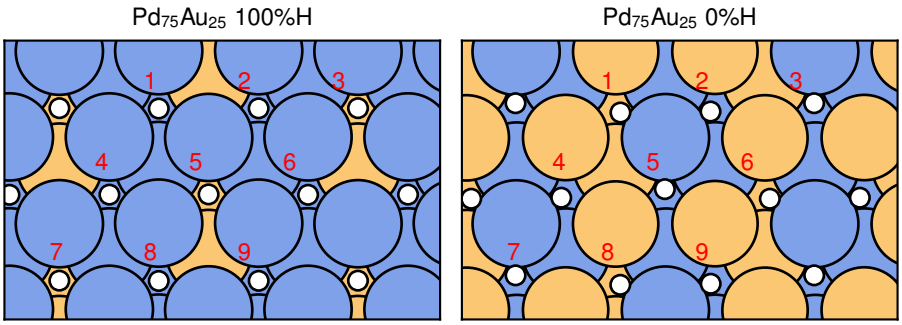

For the Pd-rich structure to the left, we can group the H sites into two groups, {1, 2, 4, 6, 8, 9} and {3, 5, 7}, with the same nearest neighbors. Similarly for Au-rich structures, {1, 6, 8}, {2, 4, 0} and {3, 5, 7}. This means that for a single adsorbate, only 2 or 3 sites have to be considered. It is, in principle, possible to enumerate the configurations for higher coverages as well but would result in a large number of configurations. Instead, we generate 3 random configurations for each considered combination of H and CO coverage.

Supplementary Tables

Table S1: Overview of metrics for neuroevolution potential models. R^2: Coefficient of determination. RMSE: Root mean square error. MAE: Mean absolute error. The RMSEs and MAEs are given units of meV atom^{-1} for the energy, meV Å^{-1} for the force, and MPa for the stress.

Model	Energy			Force			Stress		
	R^2	RMSE	MAE	R^2	RMSE	MAE	R^2	RMSE	MAE
Full									
Training (stage 1)	0.99979	16.6	6.6	0.9983	52	29	0.9816	690	219
Training (stage 2)	0.99990	11.8	4.8	0.9982	54	30	0.9775	763	214
Split 1									
Training	0.99976	17.9	6.2	0.9984	52	29	0.9857	607	241
Validation	0.91752	337.0	14.1	0.9960	82	29	0.9612	1005	248
Split 2									
Training	0.99970	20.0	7.0	0.9983	52	29	0.9902	507	207
Validation	0.99973	18.8	6.3	0.9970	74	31	0.9913	454	199
Split 3									
Training	0.99973	19.1	7.9	0.9984	51	29	0.9903	502	212
Validation	0.99983	15.1	7.3	0.9981	56	29	0.9907	486	207
Split 4									
Training	0.99957	23.9	7.8	0.9982	54	30	0.9806	705	226
Validation	0.99953	24.6	7.8	0.9967	72	32	0.9902	527	220
Split 5									
Training	0.99962	22.4	7.0	0.9983	52	29	0.9773	770	232
Validation	0.99942	28.1	8.0	0.9965	75	29	0.9868	567	230

Table S2: Overview of metrics for MACE potential models. R^2: Coefficient of determination. RMSE: Root mean square error. MAE: Mean absolute error. The RMSEs and MAEs are given units of meV atom^{-1} for the energy, meV Å^{-1} for the force, and MPa for the stress.

Model	Energy			Force			Stress		
	R^2	RMSE	MAE	R^2	RMSE	MAE	R^2	RMSE	MAE
Split 1									
Training	0.99997	6.3	1.1	0.9997	23	13	0.9954	341	106
Validation	1.00000	2.0	1.1	0.9996	26	14	0.9970	301	109
Split 2									
Training	0.99998	5.5	1.6	0.9996	26	15	0.9893	525	120
Validation	0.99988	12.7	1.8	0.9994	30	16	0.9977	247	116
Split 3									
Training	0.99997	6.4	1.2	0.9996	24	14	0.9930	426	108
Validation	1.00000	2.2	1.2	0.9993	34	15	0.9980	226	105
Split 4									
Training	0.99997	6.2	1.0	0.9997	22	13	0.9882	554	97
Validation	0.99999	2.8	1.0	0.9997	23	13	0.9959	322	98
Split 5									
Training	0.99998	4.7	1.1	0.9997	23	13	0.9934	415	107
Validation	0.99988	12.8	1.3	0.9996	24	14	0.9614	972	117

Table S3: Overview of metrics for the 'full" neuroevolution potential (NEP) model measured against the entire reference set (also see Table S1). R^2: Coefficient of determination. RMSE: Root mean square error. MAE: Mean absolute error. The RMSEs and MAEs are given units of meV atom^{-1} for the energy, meV Å^{-1} for the force, and MPa for the stress.

Type	Number of		Energy			Force			Stress		
	structures	atoms	R^2	RMSE	MAE	R^2	RMSE	MAE	R^2	RMSE	MAE
Atoms	6	6									
Ground states											
bulk	12	18	0.99986	10.9	7.7						
dimer	6	12	0.99998	10.8	7.1		60	28			
{111} surfaces	24	2184	0.99747	3.9	3.2		26	19	0.9956	136	87
Dimers	574	1148	0.99932	45.7	24.8	0.9982	89	31	0.9939	4	1
Unrelaxed structures (cut from bulk)											
{100} surfaces	25	1120	0.99988	7.9	5.5		31	13	0.8147	286	147
{110} surfaces	30	1400	0.99976	11.6	7.2		30	16	0.7374	359	150
{111} surfaces	98	5668	0.99979	6.8	5.0	0.9692	32	19	0.9910	225	142
Energy-volume curves											
bulk	1781	3454	0.99993	7.3	4.9	0.0000	0	0	0.9910	872	459
Rattled											
bulk	203	29185	0.99980	11.5	8.7	0.9909	41	30	0.9980	632	347
bulk alloys	404	18401	0.99968	6.0	4.9	0.9969	64	45	0.9986	638	328
{100} surfaces	272	8573	0.99989	11.1	7.2	0.9946	97	60	0.9914	306	154
{110} surfaces	288	13523	0.99988	8.1	5.8	0.9949	93	60	0.9955	323	139
{111} surfaces	236	9231	0.99995	5.5	4.0	0.9961	90	56	0.9947	465	216
Enumerated structures											
bulk hydrides	1612	16176	0.99991	7.5	4.9	0.9863	25	13	0.7782	446	282
alloyed hydrides	4647	38970	0.99961	7.3	5.1	0.9536	37	22	0.7744	591	335
Active learning											
bulk	196	20625	0.99997	4.5	3.2	0.9953	56	38	0.9819	415	264
bulk alloy	52	6779	0.99995	4.0	2.7	0.9911	88	65	0.9901	331	221
bulk hydrides	22	248	0.99656	8.2	7.4	0.9869	53	25	0.9841	1427	497
alloyed hydrides	36	392	0.99949	5.6	4.8	0.9776	35	23	0.6567	603	424
{100} surfaces	691	39485	0.99996	4.3	3.0	0.9931	54	31	0.9627	146	79
{110} surfaces	841	45539	0.99989	6.8	3.8	0.9901	59	35	0.9773	103	61
{111} surfaces	749	33996	0.99989	7.0	3.7	0.9924	60	35	0.9689	173	101
gas	33	3241	0.99301	58.9	35.2	0.9874	98	47	0.8173	147	82
ice	105	6390	0.82626	11.4	3.5	0.9818	101	67	0.8915	794	443
Controlled displacements											
{100} surfaces	336	41868	0.99999	2.6	1.9	0.9971	24	9	0.9120	121	71
{110} surfaces	693	57677	0.99999	2.9	2.1	0.9967	33	15	0.9905	47	28
{111} surfaces	813	55444	0.99999	2.7	2.1	0.9982	35	16	0.9781	89	56
Hydrogen migration											
bulk	162	5346	0.99994	6.4	4.6	0.9958	31	14	0.9964	800	395
bulk alloys	371	12243	0.99970	3.5	2.7	0.9963	40	21	0.9990	424	230
{100} surfaces	66	2178	0.99999	3.0	2.6	0.9782	28	15	0.9557	132	81
{110} surfaces	66	2178	0.99997	4.1	4.0	0.9059	27	16	0.9722	102	64
{111} surfaces	39	975	0.99994	4.7	4.5	0.9986	41	18	0.9776	136	88
Water											
bulk	960	92493	0.97875	2.4	1.1	0.9995	61	43	0.9972	129	94
droplet	14	1134	0.89732	5.1	4.7	0.9945	57	43	0.9677	80	60
molecular dimers	962	5772	0.99501	7.0	5.1	0.9986	120	70	0.9972	2	1
monomers	978	2934	0.99982	2.1	1.4	0.9998	42	23	0.9995	1	0

Table S4: Overview of metrics for the "split 1" MACE model measured against the entire reference set (also see Table S2). R^2: Coefficient of determination. RMSE: Root mean square error. MAE: Mean absolute error. The RMSEs and MAEs are given units of for the energy, meV Å^{-1} for the force, and MPa for the stress.

Type	Number of		Energy			Force			Stress		
	structures	atoms	R^2	RMSE	MAE	R^2	RMSE	MAE	R^2	RMSE	MAE
Atoms	6	6									
Ground states											
bulk	12	18	0.99987	10.5	4.5						
dimer	6	12	1.00000	3.1	2.5		56	22			
{111} surfaces	24	2184	0.99999	0.3	0.3		11	9		38	25
Dimers	574	1148	0.99999	5.7	3.2	0.9994	52	14	0.9986	2	1
Unrelaxed structures (cut from bulk)											
{100} surfaces	25	1120	0.99999	1.9	1.2		12	6	0.9642	126	69
{110} surfaces	30	1400	1.00000	1.0	0.9		16	9	0.9643	132	73
{111} surfaces	98	5668	0.99999	1.8	1.2	0.9930	15	10	0.9967	136	88
Energy-volume curves											
bulk	1781	3454	1.00000	0.8	0.6				0.9994	226	105
Rattled											
bulk	203	29185	1.00000	0.8	0.6	0.9983	18	13	1.0000	94	60
bulk alloys	404	18401	0.99999	1.3	1.0	0.9994	28	20	0.9999	143	92
{100} surfaces	272	8573	0.99999	3.1	2.1	0.9990	42	29	0.9987	117	67
{110} surfaces	288	13523	0.99999	1.9	1.4	0.9992	38	27	0.9995	102	53
{111} surfaces	236	9231	0.99999	2.0	1.4	0.9993	37	26	0.9993	171	74
Enumerated structures											
bulk hydrides	1612	16176	1.00000	1.2	0.8	0.9973	11	6	0.9106	283	186
alloyed hydrides	4647	38970	0.99998	1.7	1.3	0.9920	15	9	0.8925	408	188
Active learning											
bulk	196	20625	1.00000	1.0	0.8	0.9986	30	20	0.9987	110	75
bulk alloy	52	6779	0.99999	1.5	1.1	0.9975	47	36	0.9991	100	74
bulk hydrides	22	248	0.99886	4.7	2.2	0.9968	26	14	0.9953	774	400
alloyed hydrides	36	392	0.99998	1.1	1.0	0.9932	19	13	0.8136	445	323
{100} surfaces	691	39485	1.00000	1.1	0.9	0.9981	28	17	0.9807	105	48
{110} surfaces	841	45539	0.99999	2.0	0.9	0.9971	32	18	0.9862	80	46
{111} surfaces	749	33996	0.99999	1.8	0.8	0.9981	30	19	0.9946	72	40
gas	33	3241	0.99997	3.7	2.1	0.9977	42	10	0.9280	92	43
ice	105	6390	0.98375	3.5	1.4	0.9983	31	20	0.9429	576	367
Controlled displacements											
{100} surfaces	336	41868	1.00000	1.1	0.7	0.9995	10	5	0.9528	89	51
{110} surfaces	693	57677	1.00000	1.2	1.0	0.9994	14	8	0.9881	53	28
{111} surfaces	813	55444	1.00000	1.3	1.0	0.9997	14	8	0.9926	52	31
Hydrogen migration											
bulk	162	5346	1.00000	1.4	1.2	0.9997	9	6	0.9998	200	131
bulk alloys	371	12243	0.99999	0.6	0.5	0.9997	11	8	0.9998	162	102
{100} surfaces	66	2178	1.00000	0.5	0.4	0.9951	13	8	0.9754	98	59
{110} surfaces	66	2178	1.00000	1.4	1.3	0.9760	14	9	0.9724	102	67
{111} surfaces	39	975	1.00000	0.8	0.8	0.9999	13	8	0.9953	62	37
Water											
bulk	960	92493	0.99834	0.7	0.3	0.9999	19	14	0.9935	197	143
droplet	14	1134	0.98784	1.8	1.4	0.9988	26	20	0.9616	88	65
molecular dimers	962	5772	0.99978	1.5	1.1	0.9999	29	16	0.9993	1	1
monomers	978	2934	0.99999	0.4	0.2	1.0000	5	3	0.9996	1	0

Supplementary Figures

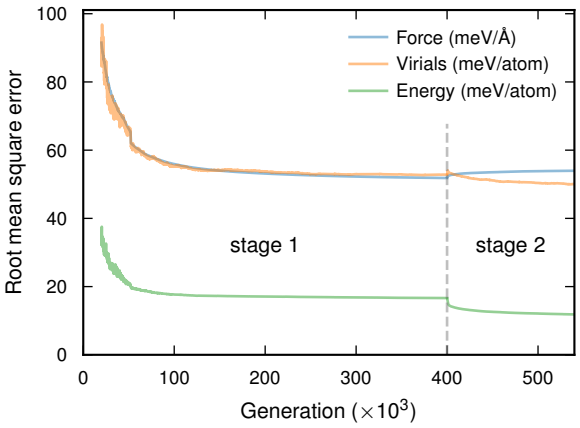

Figure S1: Learning curves for "full" neuroevolution potential model trained in two stages. In the first stage the model was trained using weights of 10, 1, and 0.1 for forces, energies, and virials, while in the second stage the weights were set to 5, 5, and 0.1, respectively.

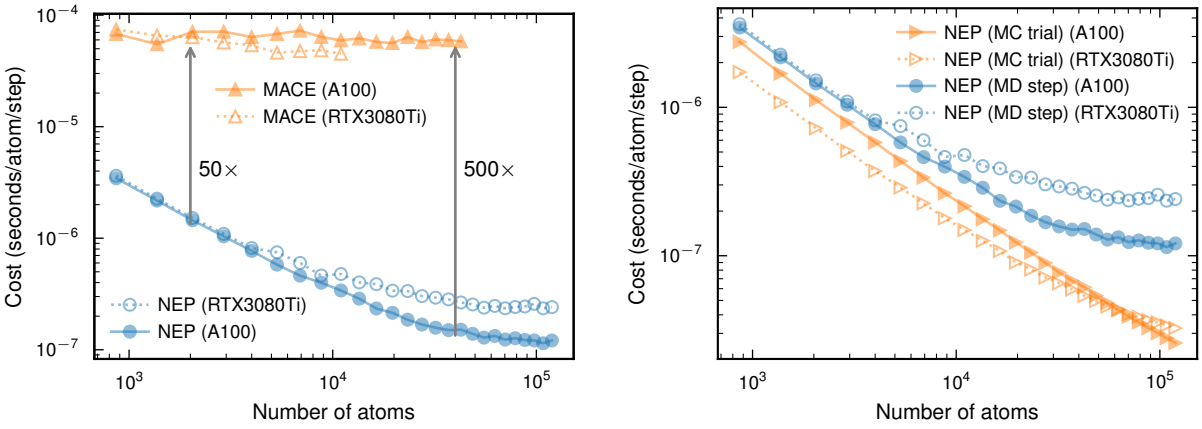

Figure S2: Computational cost of MLIP models, comparing the performance of NEP and MACE models in MD simulations (left) as well as the cost of MD and Monte Carlo (MC) steps when using NEP models (right) on different GPUs.

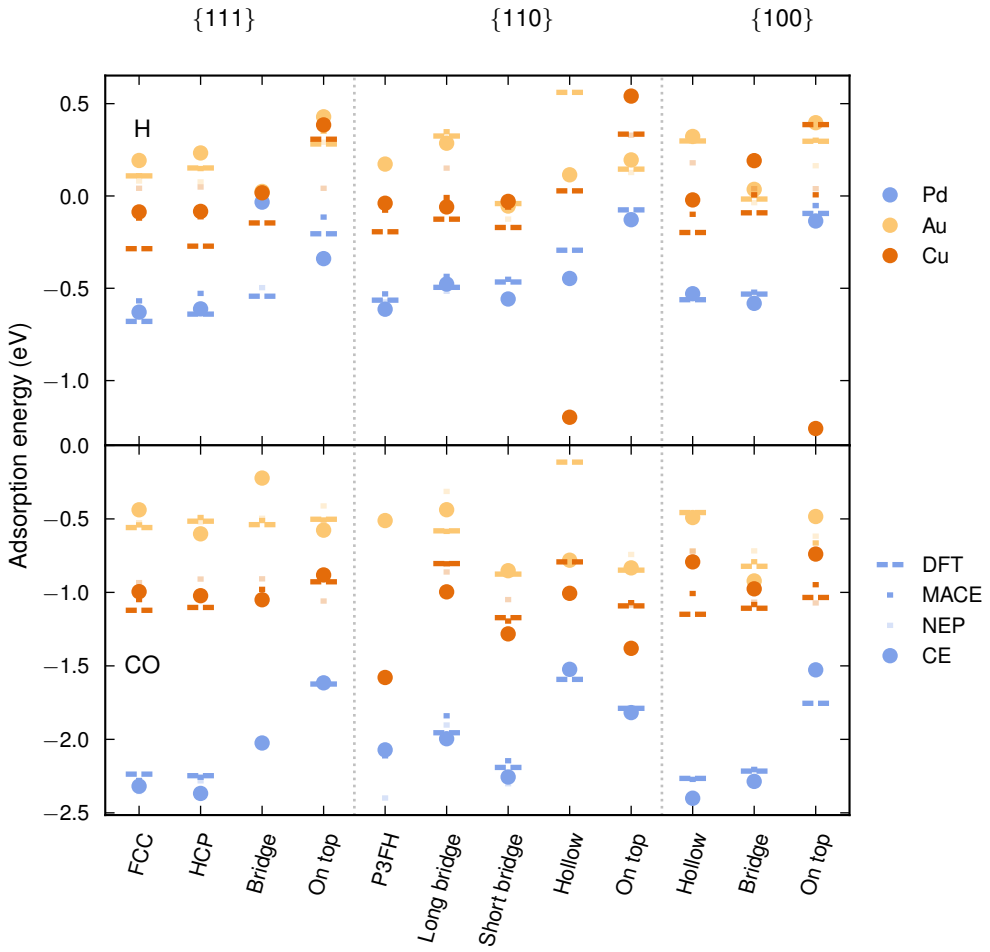

Figure S3: Adsorption energies for all adsorption sites on pure Pd (blue), Cu (orange), and Au (yellow).

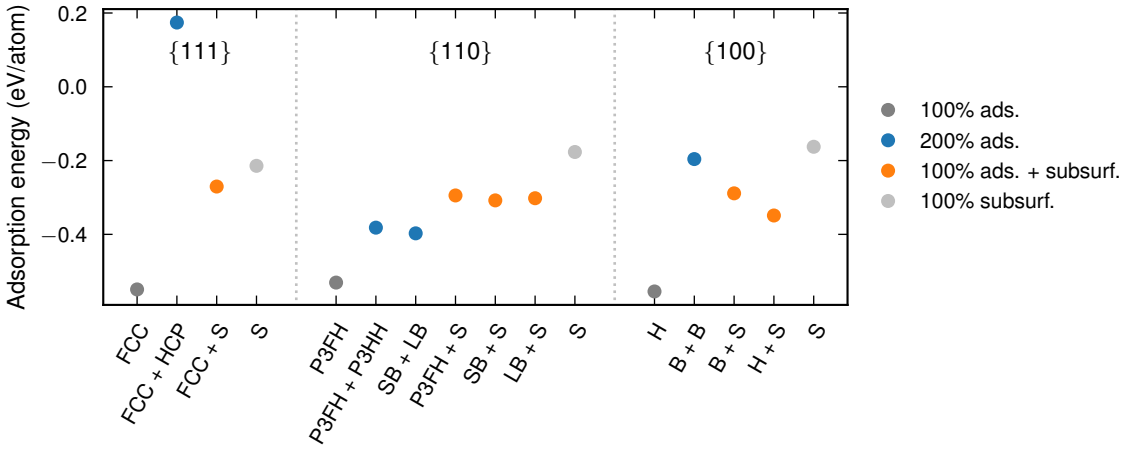

Figure S4: Adsorption energy for high-coverage configurations compared to the corresponding energy of one of the adsorbates is moved to the subsurface. For {111} and {100}, it is energetically favorable for an H atom to move to the subsurface than surpassing 100% surface coverage, while {110} allows 200 % surface coverage before occupying the subsurface. Note that here, 200% coverage refers to a 2:1 ratio between the adsorbate and surface metal atoms, while everywhere else in this work this will be referred to as 100% coverage for {110} since it is 100% of the maximum coverage.

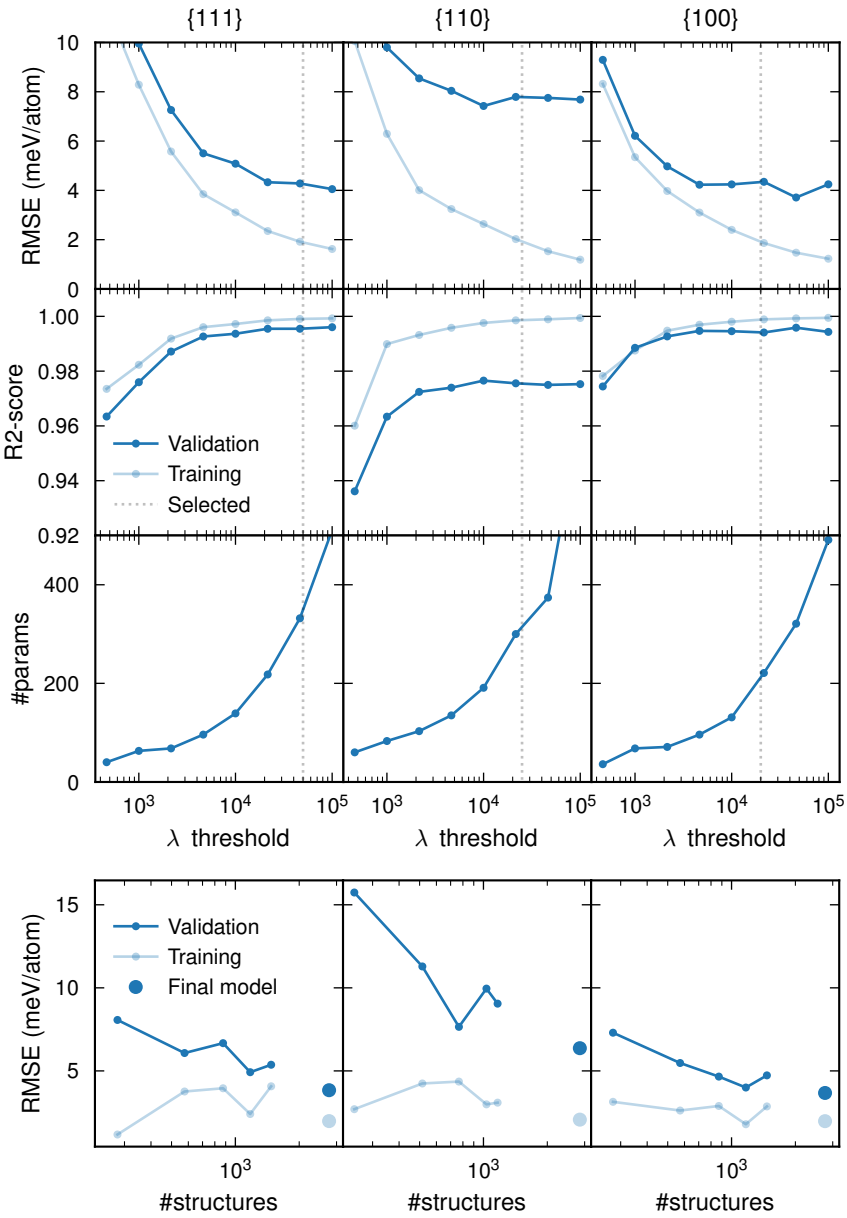

Figure S5: The upper panel shows the tuning of the ARDR hyperparameter λ threshold for the CEs via the RMSE and R2-score from validation and training and the corresponding number of features. The selected value is indicated by the dotted lines. The lower panel shows the RMSE obtained during active learning. The final model has an increase in training set size because datasets from adsorbate-free CEs developed in parallel were included. The RMSEs are reported in units of meV per metal atom.

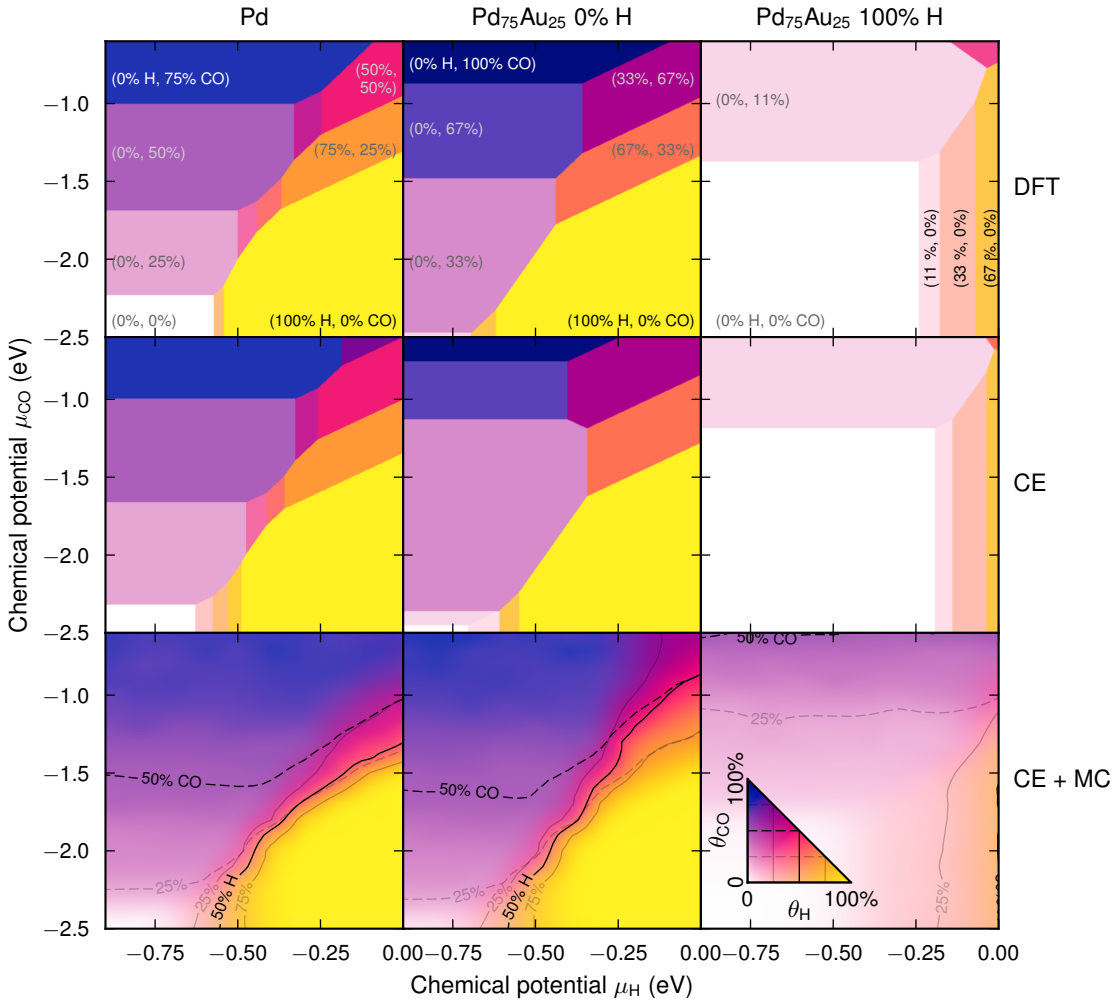

Figure S6: Surface phase diagrams for Pd and $Pd_{75}Au_{25}$ where the alloy configuration is obtained from MC simulations with 0% or 100% H coverage. The alloy configurations differ between the discrete and continuous phase diagrams. For the discrete case, the procedure is described in Supplementary Note 4 and the continuous case uses the same structures as in the paper.

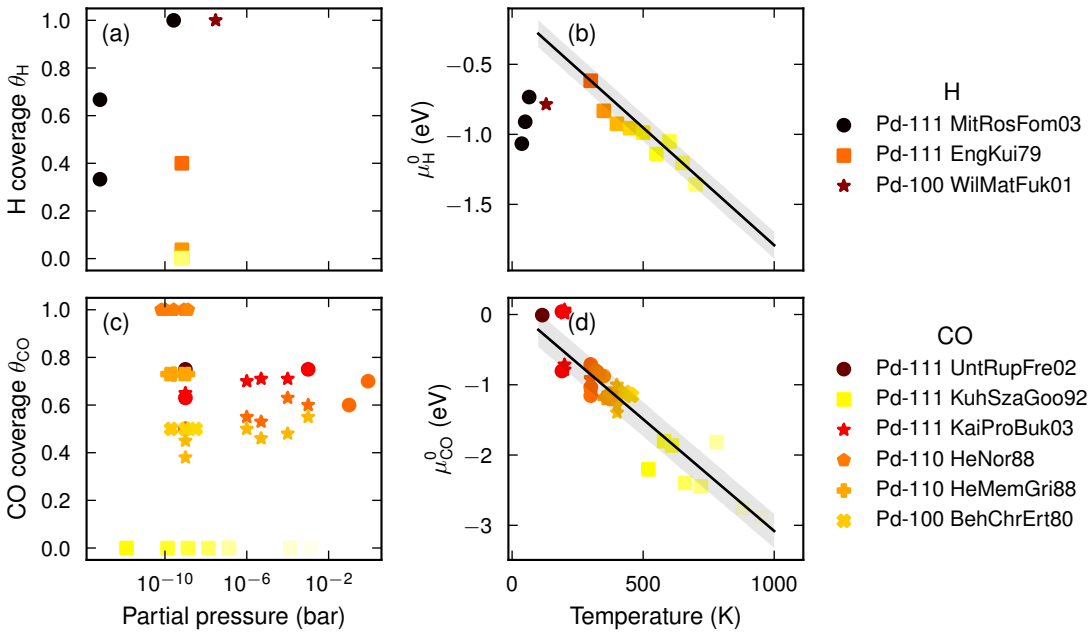

Figure S7: Experimental records of H (17–19) or CO (20–25) coverages on Pd {111}, {110}, and {100} surfaces. These records are combined with our MC simulations to calculate the reference chemical potential.

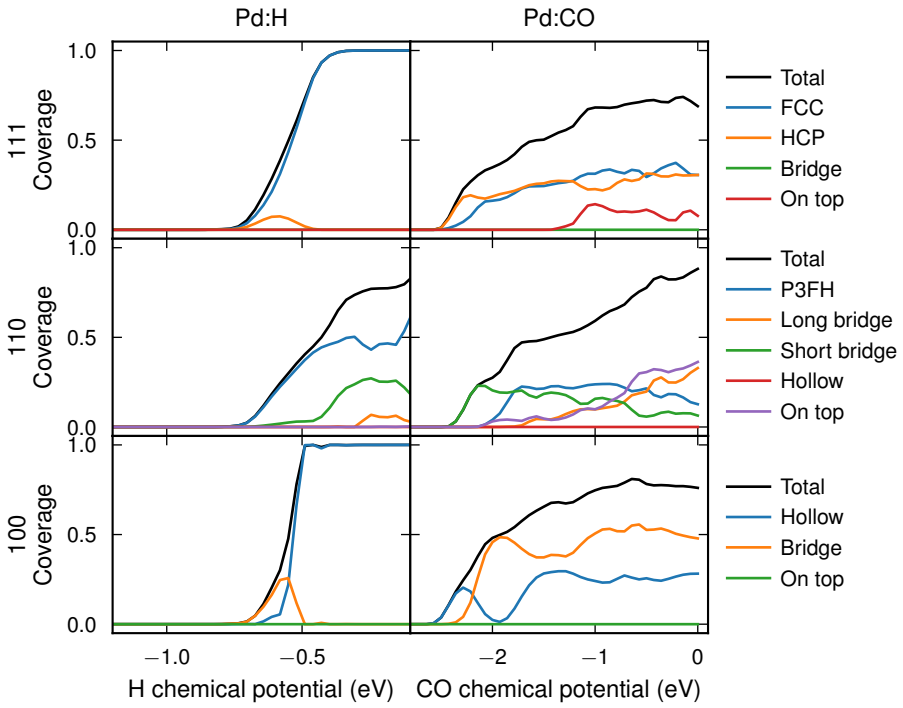

Figure S8: Site-specific H and CO coverages on Pd.

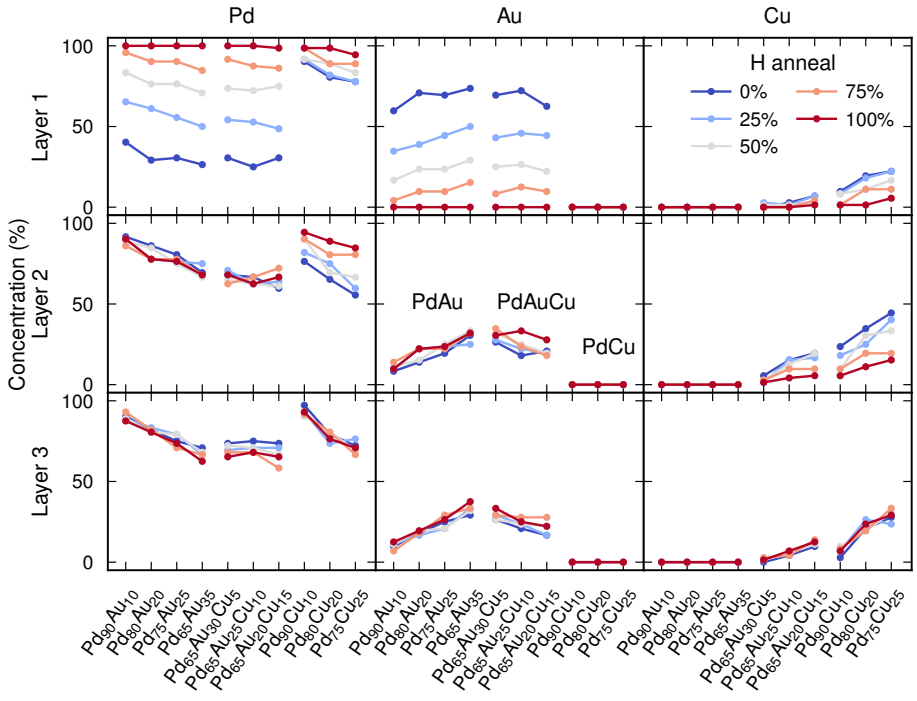

Figure S9: Surface compositions for PdAuCu {111} surface slabs from MC simulations at 773 K with varying H coverage.

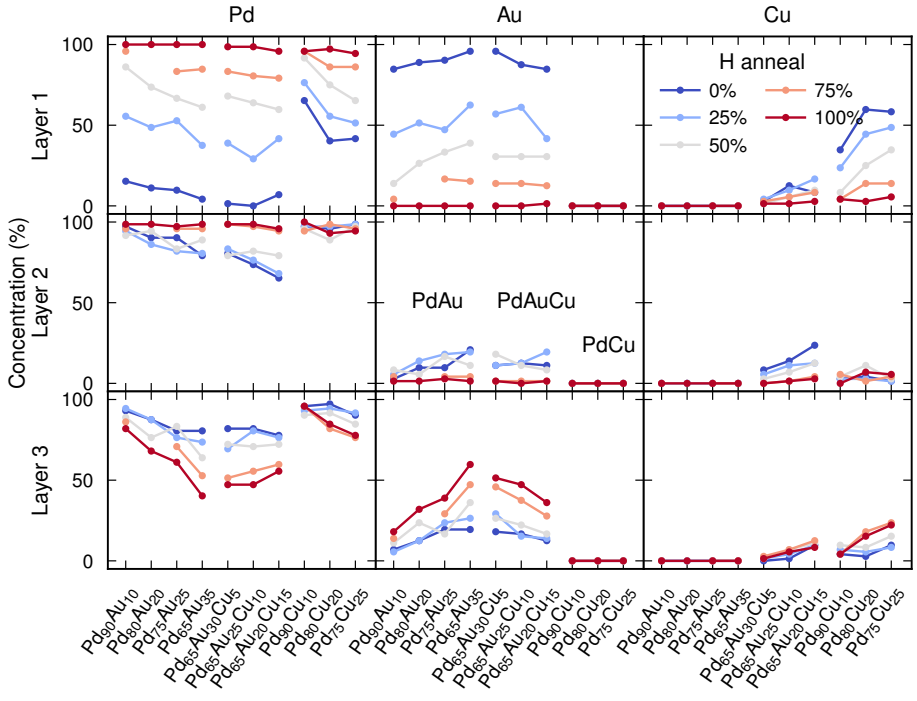

Figure S10: Surface compositions for PdAuCu {110} surface slabs from MC simulations at 773 K with varying H coverage.

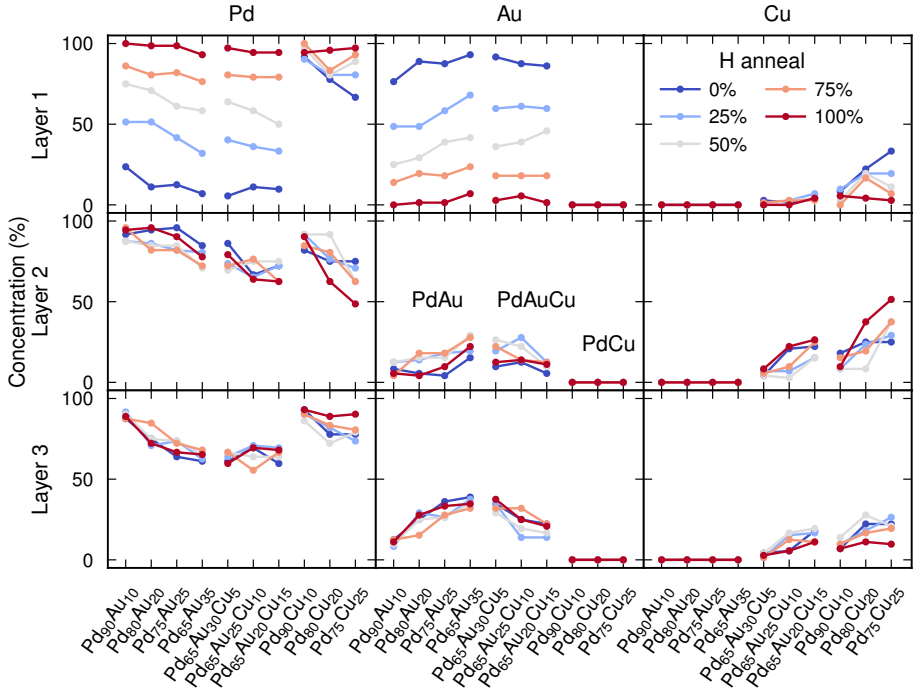

Figure S11: Surface compositions for PdAuCu {100} surface slabs from MC simulations at 773 K with varying H coverage.

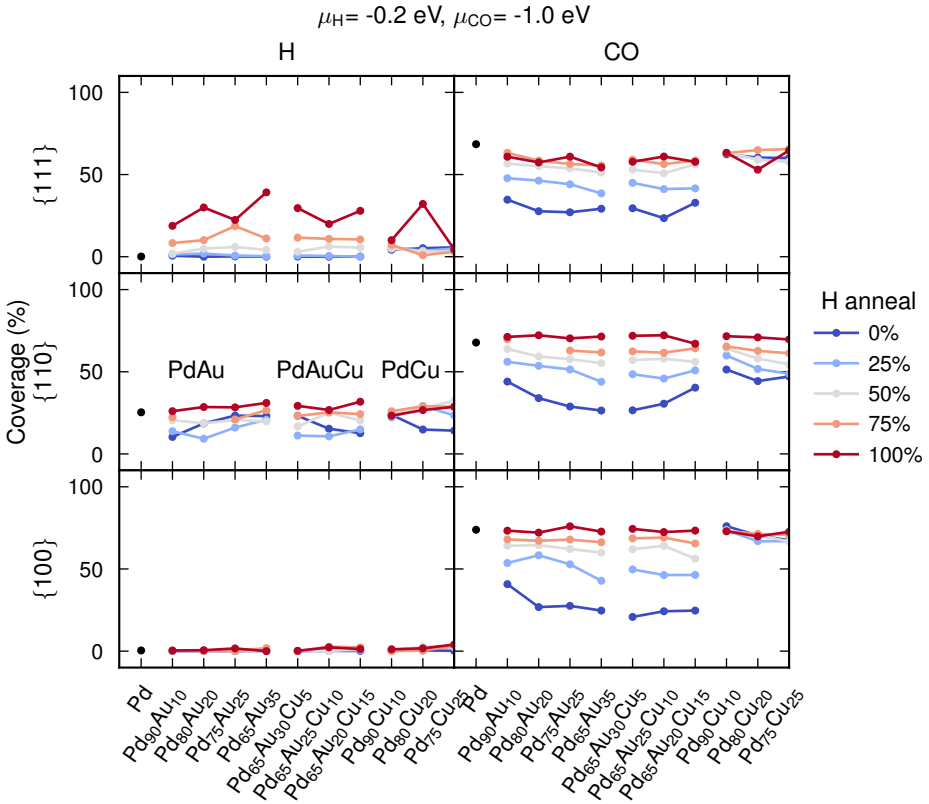

Figure S12: Adsorbate coverages on PdAuCu alloys at $\mu_H = -0.2$ eV and $\mu_{CO} = -1.0$ eV. The slab configuration is obtained from MC simulations at 773 K with varying H coverage as indicated by the legend.

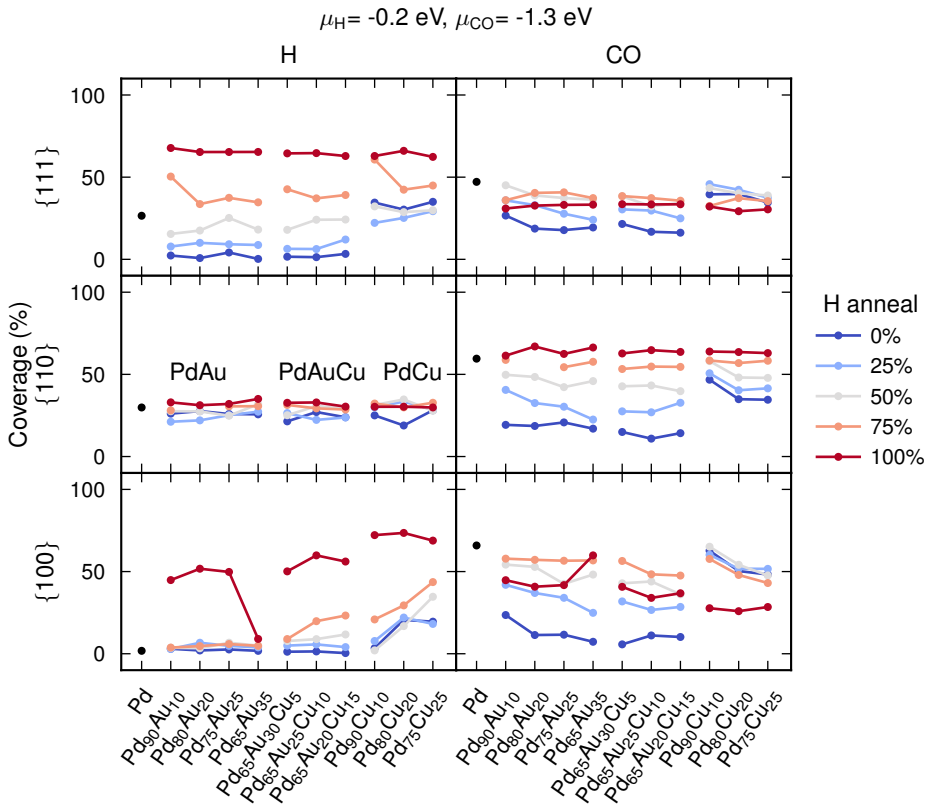

Figure S13: Adsorbate coverages on PdAuCu alloys at $\mu_H = -0.2\,eV$ and $\mu_{CO} = -1.3\,eV$. The slab configuration is obtained from MC simulations at 773 K with varying H coverage as indicated by the legend.

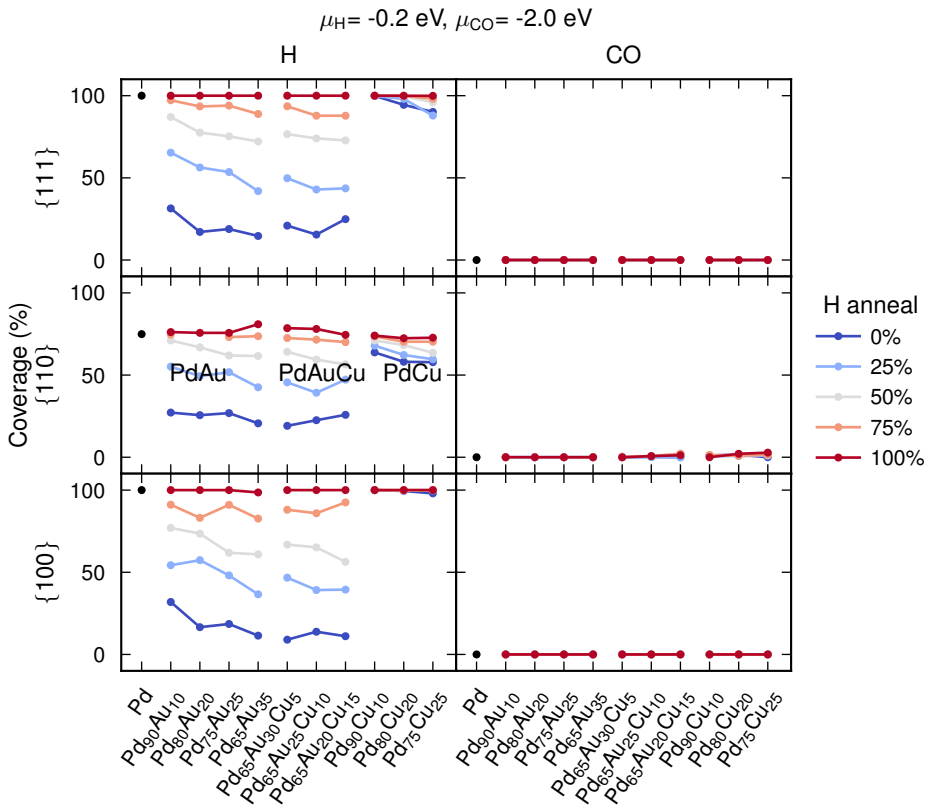

Figure S14: Adsorbate coverages on PdAuCu alloys at $\mu_H = -0.2\,eV$ and $\mu_{CO} = -2.0\,eV$. The slab configuration is obtained from MC simulations at 773 K with varying H coverage as indicated by the legend.

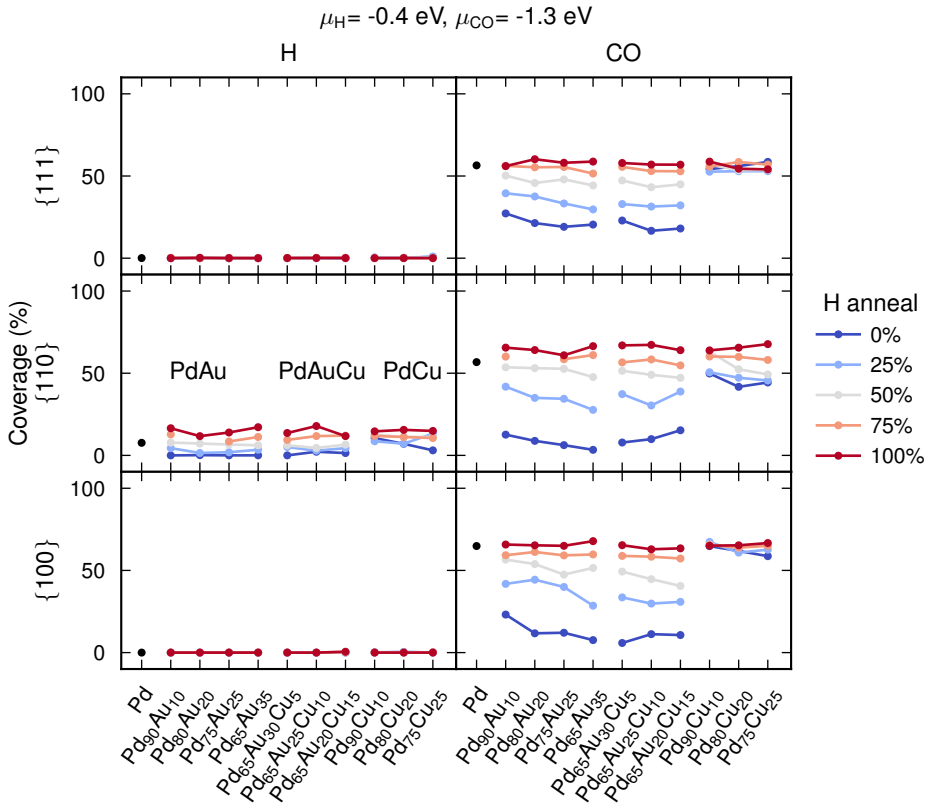

Figure S15: Adsorbate coverages on PdAuCu alloys at $\mu_H = -0.4\,eV$ and $\mu_{CO} = -1.3\,eV$. The slab configuration is obtained from MC simulations at 773 K with varying H coverage as indicated by the legend.

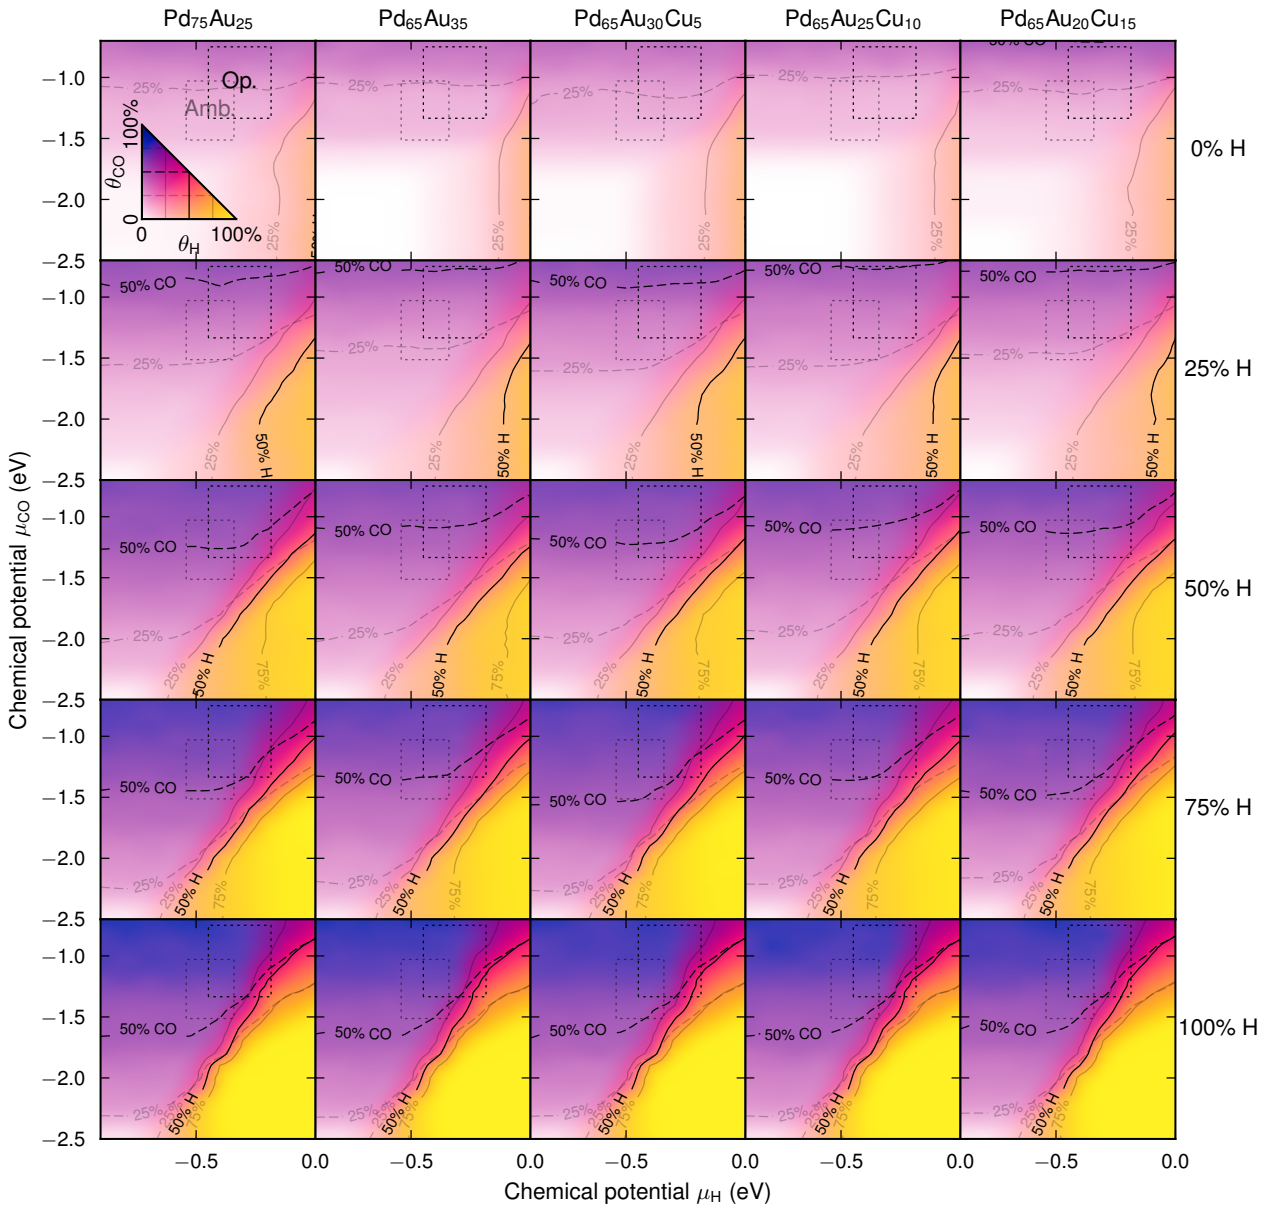

Figure S16: Surface phase diagrams for PdAuCu {111} surface slabs at 300 K. The slab configuration is obtained from MC simulations at 773 K with varying H coverage as indicated by the labels to the right.

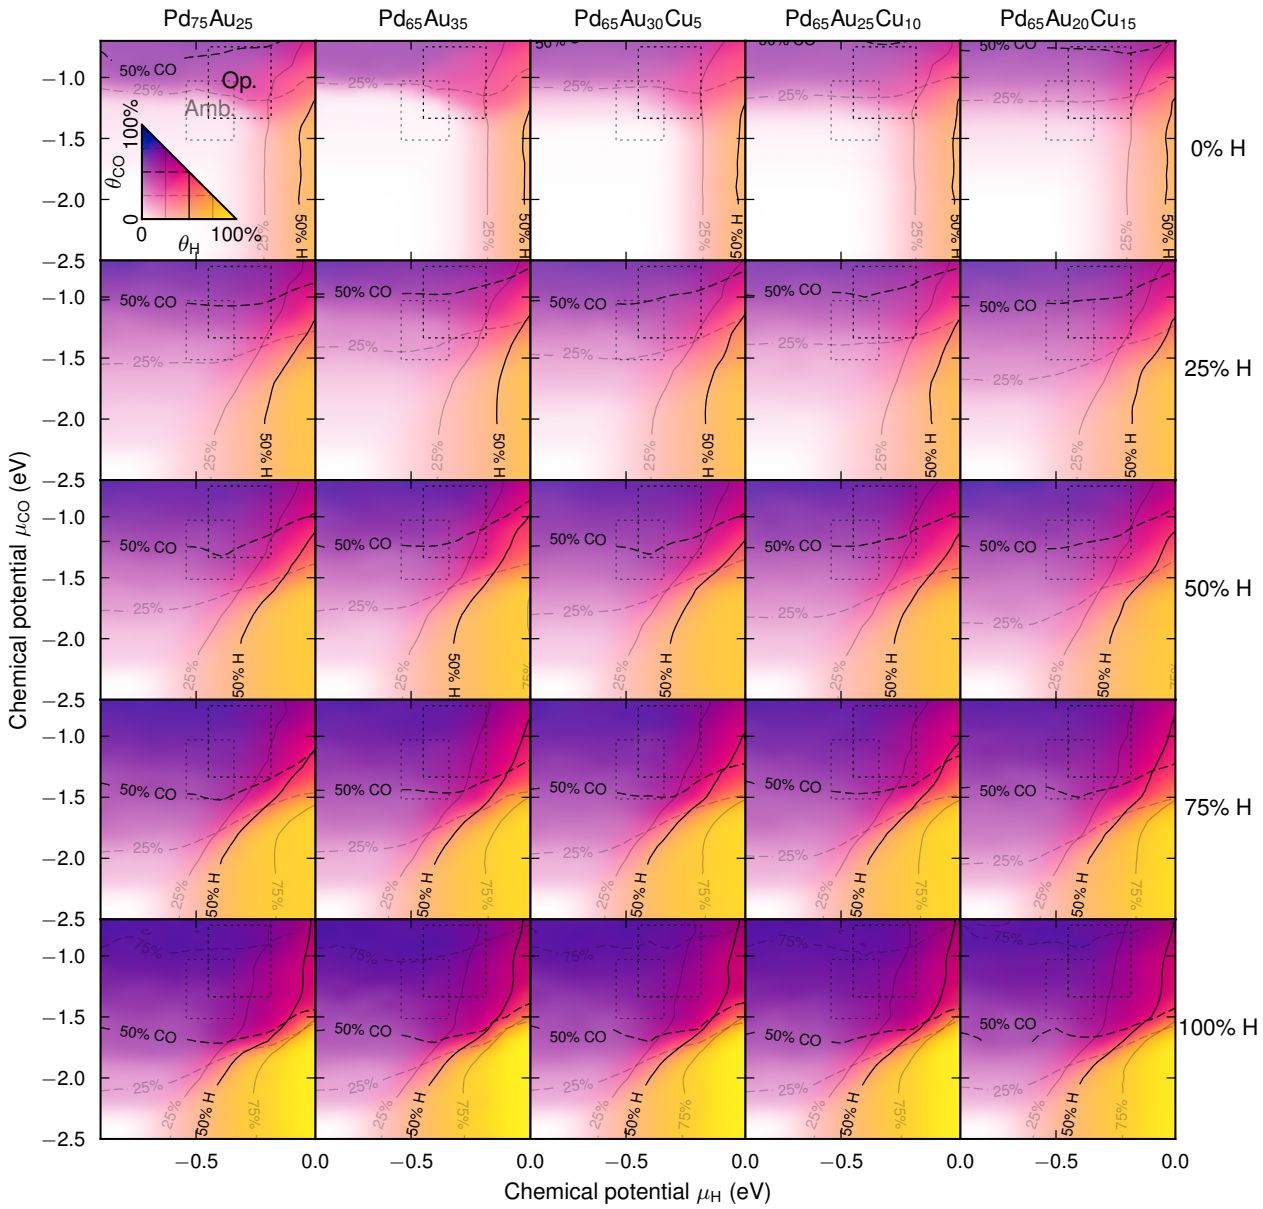

Figure S17: Surface phase diagrams for PdAuCu {110} surface slabs at 300 K. The slab configuration is obtained from MC simulations at 773 K with varying H coverage as indicated by the labels to the right.

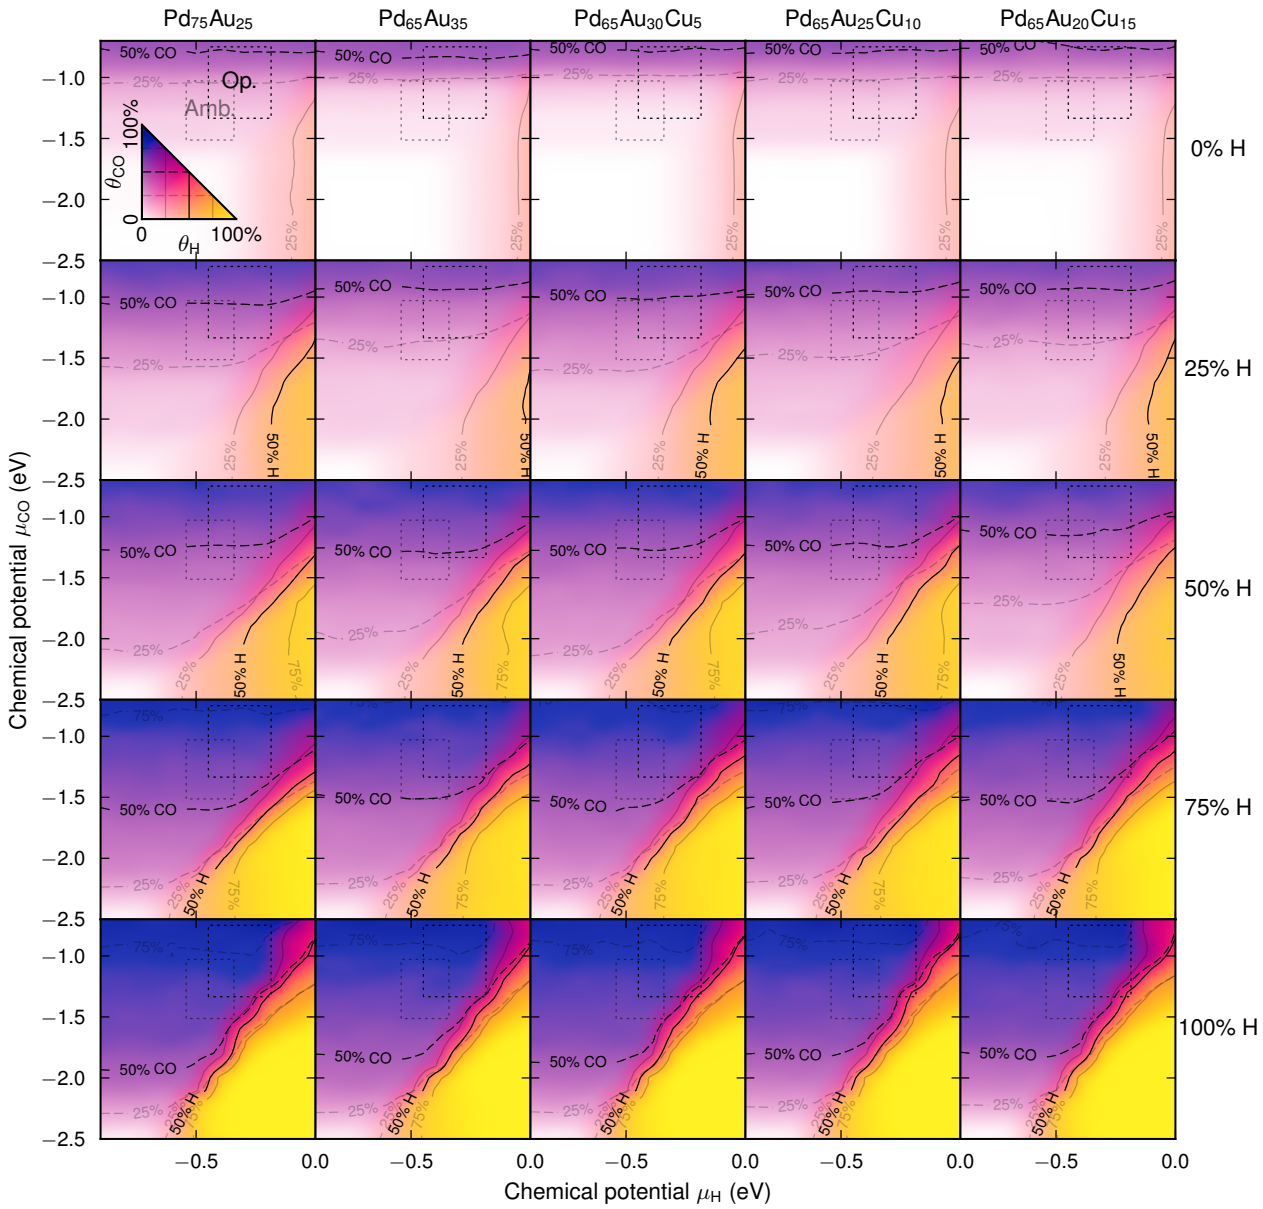

Figure S18: Surface phase diagrams for PdAuCu {100} surface slabs at 300 K. The slab configuration is obtained from MC simulations at 773 K with varying H coverage as indicated by the labels to the right.

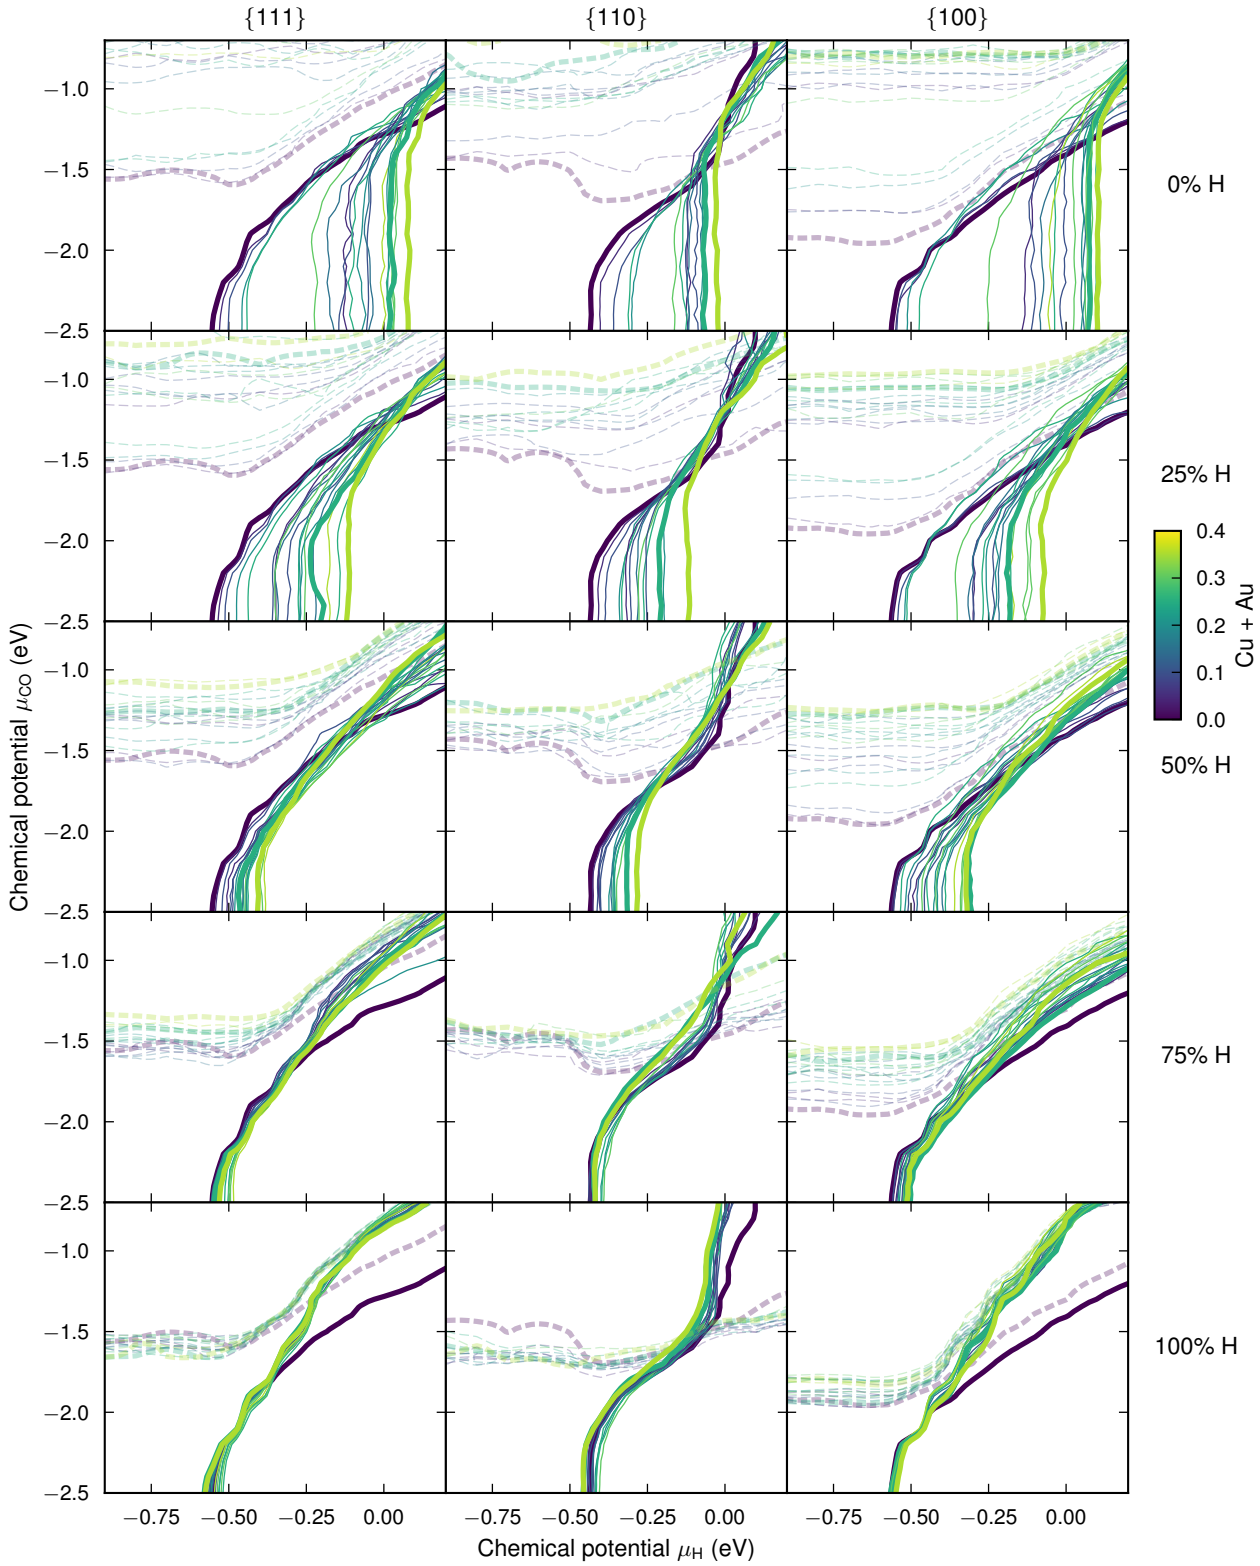

Figure S19: 50% coverage contour lines for Pd and various The slab configuration is obtained from MC simulations at 773 K with varying H coverage as indicated by the labels to the right.

Supplementary References

[1] Z. Fan, Y. Wang, P. Ying, K. Song, J. Wang, Y. Wang, Z. Zeng, K. Xu, E. Lindgren, J. M. Rahm, A. J. Gabourie, J. Liu, H. Dong, J. Wu, Y. Chen, Z. Zhong, J. Sun, P. Erhart, Y. Su, and T. Ala-Nissila, *GPUMD: A package for constructing accurate machine-learned potentials and performing highly efficient atomistic simulations,* Journal of Chemical Physics **157**, 114801 (2022). doi:10.1063/5.0106617.

[2] K. Song, R. Zhao, J. Liu, Y. Wang, E. Lindgren, Y. Wang, S. Chen, K. Xu, T. Liang, P. Ying, N. Xu, Z. Zhao, J. Shi, J. Wang, S. Lyu, Z. Zeng, S. Liang, H. Dong, L. Sun, Y. Chen, Z. Zhang, W. Guo, P. Qian, J. Sun, P. Erhart, T. Ala-Nissila, Y. Su, and Z. Fan, *General-purpose machine-learned potential for 16 elemental metals and their alloys,* Nature Communications **15**, 10208 (2024). doi:10.1038/s41467-024-54554-x.

[3] E. Fransson, J. Wiktor, and P. Erhart, *Phase transitions in inorganic halide perovskites from machine-learned potentials,* The Journal of Physical Chemistry C **127**, 13773 (2023). doi:10.1021/acs.jpcc.3c01542.

[4] K. Xu, H. Bu, S. Pan, E. Lindgren, Y. Wu, Y. Wang, J. Liu, K. Song, B. Xu, Y. Li, T. Hainer, L. Svensson, J. Wiktor, R. Zhao, H. Huang, C. Qian, S. Zhang, Z. Zeng, B. Zhang, B. Tang, Y. Xiao, Z. Yan, J. Shi, Z. Liang, J. Wang, T. Liang, S. Cao, Y. Wang, P. Ying, N. Xu, C. Chen, Y. Zhang, Z. Chen, X. Wu, W. Jiang, E. Berger, Y. Li, S. Chen, A. J. Gabourie, H. Dong, S. Xiong, N. Wei, Y. Chen, J. Xu, F. Ding, Z. Sun, T. Ala-Nissila, A. Harju, J. Zheng, P. Guan, P. Erhart, J. Sun, W. Ouyang, Y. Su, and Z. Fan, *GPUMD 4.0: A high-performance molecular dynamics package for versatile materials simulations with machine-learned potentials,* Materials Genome Engineering Advances **3**, e70028 (2025). doi:10.1002/mgea.70028.

[5] E. Lindgren, J. M. Rahm, E. Fransson, F. Eriksson, N. Österbacka, Z. Fan, and P. Erhart, *calorine: A Python package for constructing and sampling neuroevolution potential models,* Journal of Open Source Software **9**, 6264 (2024). doi:10.21105/joss.06264.

[6] D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters, and J. Schmidhuber, *Natural Evolution Strategies,* Journal of Machine Learning Research **15**, 949 (2014). http://jmlr.org/papers/v15/wierstra14a.html.

[7] I. Batatia, D. P. Kovacs, G. N. C. Simm, C. Ortner, and G. Csanyi, *MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields,* in *Advances in Neural Information Processing Systems*, edited by A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Curran Associates, Inc., 2022. arXiv:2206.07697. doi:10.48550/arXiv.2206.07697.

[8] I. Batatia, S. Batzner, D. P. Kovács, A. Musaelian, G. N. C. Simm, R. Drautz, C. Ortner, B. Kozinsky, and G. Csányi, *The Design Space of E(3)-Equivariant Atom-Centered Interatomic Potentials,* 2022. doi:10.48550/arXiv.2205.06643.

[9] A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. Bjerre Jensen, J. Kermode, J. R. Kitchin, E. Leonhard Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. Bergmann Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, and K. W. Jacobsen, *The atomic simulation environment—a Python library for working with atoms,* Journal of Physics: Condensed Matter **29**, 273002 (2017). doi:10.1088/1361-648X/aa680e.

[10] F. Eriksson, E. Fransson, and P. Erhart, *The hiphive package for the extraction of high-order force constants by machine learning,* Advanced Theory and Simulations **2**, 1800184 (2019). doi:10.1002/adts.201800184.

[11] G. L. W. Hart and R. W. Forcade, *Algorithm for generating derivative structures,* Physical Review B **77**, 224115 (2008). doi:10.1103/PhysRevB.77.224115.

[12] G. L. Hart and R. W. Forcade, *Generating derivative structures from multilattices: Algorithm and application to hcp alloys,* Physical Review B **80**, 014120 (2009). doi:10.1103/PhysRevB.80.014120.

[13] M. Ångqvist, W. A. Muñoz, J. M. Rahm, E. Fransson, C. Durniak, P. Rozyczko, T. H. Rod, and P. Erhart, *ICET – A Python Library for Constructing and Sampling Alloy Cluster Expansions,* Advanced Theory and Simulations **0**, 1900015 (2019). doi:10.1002/adts.201900015.

[14] G. Henkelman and H. Jonsson, *Improved Tangent Estimate in the NEB Method for Finding Minimum Energy Paths and Saddle Points,* Journal of Chemical Physics **113**, 9978 (2000). doi:10.1063/1.1323224.

[15] G. Henkelman, B. P. Uberuaga, and H. Jonsson, *A Climbing-Image NEB Method for Finding Saddle Points and Minimum Energy Paths,* Journal of Chemical Physics **113**, 9901 (2000). doi:10.1063/1.1329672.

[16] P. Lindgren, G. Kastlunger, and A. A. Peterson, *Scaled and Dynamic Optimizations of Nudged Elastic Bands,* Journal of Chemical Theory and Computation **15**, 5787 (2019). doi:10.1021/acs.jctc.9b00633.

[17] T. Mitsui, M. K. Rose, E. Fomin, D. F. Ogletree, and M. Salmeron, *Dissociative hydrogen adsorption on palladium requires aggregates of three or more vacancies,* Nature **422**, 705 (2003). doi:10.1038/nature01557.

[18] T. Engel and H. Kuipers, *A molecular-beam investigation of the scattering, adsorption and absorption of H_2 and D_2 from/on/in Pd(111),* Surface Science **90**, 162 (1979). doi:10.1016/0039-6028(79)90017-7.

[19] M. Wilde, M. Matsumoto, K. Fukutani, and T. Aruga, *Depth-resolved analysis of subsurface hydrogen absorbed by Pd(100),* Surface Science **482-485**, 346 (2001). doi:10.1016/S0039-6028(01)00727-0.

[20] H. Unterhalt, G. Rupprechter, and H.-J. Freund, *Vibrational Sum Frequency Spectroscopy on Pd(111) and Supported Pd Nanoparticles: CO Adsorption from Ultrahigh Vacuum to Atmospheric Pressure,* The Journal of Physical Chemistry B **106**, 356 (2002). doi:10.1021/jp013000+.

[21] W. K. Kuhn, J. Szanyi, and D. W. Goodman, *CO adsorption on Pd(111): the effects of temperature and pressure,* Surface Science **274**, L611 (1992). doi:10.1016/0039-6028(92)90834-S.

[22] V. V. Kaichev, I. P. Prosvirin, V. I. Bukhtiyarov, H. Unterhalt, G. Rupprechter, and H.-J. Freund, *High-Pressure Studies of CO Adsorption on Pd(111) by X-ray Photoelectron Spectroscopy and Sum-Frequency Generation,* The Journal of Physical Chemistry B **107**, 3522 (2003). doi:10.1021/jp021992t.

[23] J. He and P. R. Norton, *Interaction of CO with a Pd(110) surface, studied by low energy electron diffraction, thermal desorption spectroscopy, and $\Delta\Phi$,* The Journal of Chemical Physics **89**, 1170 (1988). doi:10.1063/1.455225.

[24] J. W. He, U. Memmert, K. Griffiths, W. N. Lennard, and P. R. Norton, *Absolute coverages for the various surface phases of CO and O adsorbed on Pd(110),* Surface Science **202**, L555 (1988). doi:10.1016/0039-6028(88)90031-3.

[25] R. J. Behm, K. Christmann, G. Ertl, and M. A. Van Hove, *Adsorption of CO on Pd(100),* The Journal of Chemical Physics **73**, 2984 (1980). doi:10.1063/1.440430.