

Article

Chemically Interpretable Machine Learning for Predicting HER Activity in Au-Based Alloys


Maja Kubik, Shiqi Wang *, and Pedro H. C. Camargo *

Department of Chemistry, University of Helsinki, A.I. Virtasen aukio 1, FIN-0014 Helsinki, Finland

* Correspondence: shiqi.z.wang@helsinki.fi (S.W.); pedro.camargo@helsinki.fi (P.H.C.C.)

Received: 10 October 2025; Revised: 13 November 2025; Accepted: 17 November 2025; Published: 24 November 2025

Abstract: Hydrogen evolution reaction (HER) catalysts are essential for renewable hydrogen production, but conventional screening through density functional theory (DFT) remains computationally demanding. Here, we present a proof-of-concept machine learning (ML) framework that predicts hydrogen adsorption energies on Au-based alloys using a curated dataset and nine chemically interpretable descriptors. Tree-based ensemble models achieve quantitative accuracy (mean absolute error ≈ 0.13 eV), sufficient to resolve weak, optimal, and strong adsorption regimes. Feature attribution highlights generalized coordination number and site-specific electronic descriptors as dominant factors, consistent with established chemisorption principles. Guided by these insights, the trained model was applied to a library of Au–M alloys, with targeted DFT validation identifying Au–Mo, Au–W, Au–Ta, Au–Hf, and Au–Tc as promising candidates that couple favorable adsorption with alloy stability. Beyond specific predictions, this work introduces an accessible, reproducible workflow that connects ML predictions to catalytic intuition. The framework serves as a reference for experimental chemists, lowering the barrier for integrating data-driven approaches into catalyst discovery.



Keywords: hydrogen evolution reaction (HER); machine learning in catalysis; Au-based alloys; adsorption energy prediction; interpretable descriptors

1. Introduction

The discovery of efficient catalysts is central to advancing sustainable energy technologies, particularly in the context of green hydrogen production from water electrolysis [1–3]. The hydrogen evolution reaction (HER) plays an important role in this process, as its activity depends sensitively on the adsorption energy of hydrogen intermediates [4,5]. Conventional catalyst discovery has relied on experimental trial-and-error or density functional theory (DFT) calculations, but both approaches are resource-intensive and scale poorly to large chemical spaces [6]. Machine learning (ML) offers a promising alternative by recognizing patterns in large datasets and predicting catalytic descriptors with far lower cost, thereby accelerating the design of new materials [7–10].

Despite this potential, ML has not been widely adopted in experimental catalysis laboratories. A major reason is that many existing studies focus on technical aspects of model development rather than on providing workflows that are accessible to chemists without a background in data science [11,12]. Data representations such as one-hot encodings or high-dimensional fingerprints often lack direct chemical interpretability [13,14]. For example, they encode ‘which element’ without reflecting chemical similarity or catalytic trends, which can make it difficult for experimentalists to connect model outputs to catalytic principles [15]. In addition, the implementation of ML pipelines typically requires extensive computational expertise, further limiting their uptake by the broader catalysis community [16–18]. These barriers have created a gap between the capabilities of ML and its practical use in guiding experimental catalyst discovery.

In this work, we present an ML workflow for predicting hydrogen adsorption on Au-based alloy surfaces. By combining a publicly available DFT dataset with a set of curated, physically meaningful descriptors, such as generalized coordination number, weighted electronegativity, and site-specific electronic parameter, we demonstrate an interpretable workflow that can be readily understood and reproduced by chemists. Using tree-



Copyright: © 2025 by the authors. This is an open access article under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Publisher’s Note: Scilight stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

based regression models, we achieve reliable prediction accuracy (mean absolute error ~ 0.13 eV for Au-containing systems) while also providing feature-importance analyses that connect directly to catalytic intuition. The aim of this study is not to introduce a new machine learning methodology, but rather to illustrate how chemists can integrate accessible ML tools into their research. The same descriptor set and ML workflow can be applied to Pt-M and Pd-M alloys once suitable training data are available, making the approach transferable to other HER catalyst families. By presenting a transparent and reproducible workflow, we seek to lower the barrier to entry for ML in catalysis and to provide a reference framework for future studies of HER and beyond.

2. Methodology

2.1. Machine Learning Workflow

Our machine learning framework follows a structured pipeline from dataset curation to catalyst screening (Figure 1). Adsorption data were first extracted and preprocessed, then transformed into chemically interpretable descriptors capturing both geometric and electronic factors that govern hydrogen binding. Tree-based ensemble algorithms were trained on these features, with iterative feedback guiding hyperparameter tuning and feature selection. The optimized model was then applied to screen Au-M alloy candidates, translating predictive insights into experimental guidance. This workflow explicitly links catalytic intuition with data-driven predictions, ensuring that each stage remains transparent and interpretable for chemists.

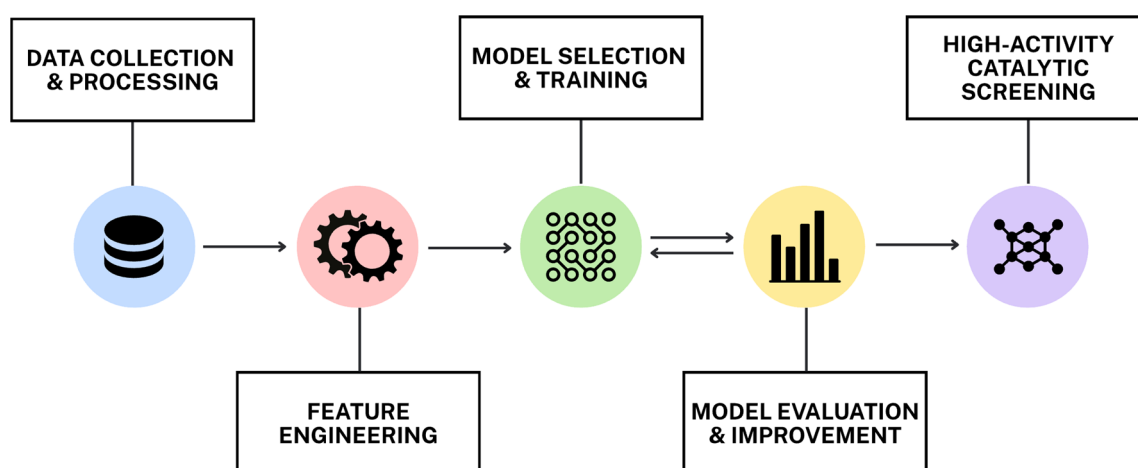


Figure 1. The machine learning workflow begins with data collection and preprocessing, followed by feature engineering to derive chemically interpretable descriptors. Models are then trained and optimized, evaluated for accuracy and interpretability, and finally applied to high-throughput screening of candidate catalytic compositions.

2.2. Data Extraction and Processing

Robust machine learning predictions require reliable and well-curated input data. For this purpose, we employed the dataset of Mamun et al. [19], which reports approximately 90,000 DFT adsorption and reaction energies across 2035 alloy surfaces for diverse adsorbates. Here, the dataset was restricted to dilute hydrogen adsorption, ensuring that the target variable corresponds directly to the HER. Only entries with a single H atom per surface unit cell were retained, while fourfold sites were discarded due to their scarcity and limited relevance under realistic catalytic conditions.

The hydrogen adsorption geometries used for model training were obtained directly from this dataset [19], in which all unique high-symmetry adsorption sites (ontop, bridge, *fcc* hollow, and *hcp* hollow) were systematically enumerated using CatKit. The slab models used in the ML part were obtained directly from the same dataset [19]. Alloy surfaces were constructed as three-layer 2×2 *fcc*-based slabs ($A1$, $L1_2$, $L1_0$), cleaved along the (111) facet for $A1/L1_2$ and the equivalent (101)/(111) facet for $L1_0$. Slabs were generated with CatKit, with the bottom two layers fixed during relaxation and only non-reconstructed surfaces retained. These structures form the basis for the adsorption data used in this study. The optimized slab geometries are publicly available through Catalysis-Hub.org and Materials Cloud.

Feature construction considered only the identity and arrangement of first-layer atoms surrounding the adsorption site. The Mamun dataset often reports multiple rows that differ only in second-layer composition, particularly for bridge motifs. Because hydrogen binding is primarily governed by the first coordination shell

[20,21], these variants were treated as replicate measurements and replaced by their mean adsorption energy. This consolidation reduced label noise and prevented over-representation of near-duplicate records. After filtering and averaging, the HER subset was reduced from 8856 to 6749 unique entries, which formed the foundation for model training and validation. This curated dataset provides a reliable platform for developing compact yet accurate ML models for catalyst discovery.

2.3. Feature Engineering

Our chosen descriptor set of nine features was designed to encode both composition and local structure at the adsorption site. Two descriptors captured the valence electron counts of the constituent elements A (S_A) and B (S_B) [22]. Three geometric descriptors were included: generalized coordination number (GCN, Equation (S1)) [23], surface facet, and volume per atom (\AA^3 , Equation (S2)) [24]. Three additional descriptors quantified composition-weighted elemental properties: average ionization energy (WIE, Equation (S3)) [25], average electronegativity (WEN, Equation (S4)) [26], and average atomic radius (WAR, Equation (S5)) [27]. The final input was a site-specific electronic descriptor (ψ , Equation (S6)) [28], which integrates local electronegativity with the valence environment of nearest neighbors. All inputs were obtained from the Mamun dataset or from standard elemental property databases. In this design, elemental type and proportion are captured numerically, in contrast to one-hot representations that only indicate the presence or absence of an element [29], providing only categorical identity but no chemical information.

The geometric descriptors establish the structural context for adsorption [30]. GCN provides a continuous measure of local coordination and serves as a compact identifier of adsorption motif [31]. Within the HER subset, which is derived from 3-layer 2×2 *fcc*-based slabs including single *fcc* metals (A_1), A_3B *fcc* alloys ($L1_2$, Cu_3Au type), and ordered AB *fcc* alloys ($L1_0$, CuAu type) cleaved along (111) or (101), this mapping yields four characteristic values corresponding to ontop, bridge, and hollow sites (*fcc* and *hcp*). All the geometries of hydrogen adsorption on various models are fully relaxed with the top layer and adsorbate. The facet descriptor, given by the Miller index, accounts for crystallographic orientation, while the atomic volume provides a proxy for bond length trends. Since many alloy compositions in the Mamun dataset are not tabulated in standard repositories, volumes were estimated by Vegard-type mixing of pure elements.

Electronic descriptors capture how alloy composition modulates charge donation and withdrawal at the adsorption site. S_A and S_B summarize shell occupancy near the Fermi level, linking directly to the familiar d-band filling picture. WEN reflects the ability of a site to attract charge, WIE describes how tightly valence electrons are bound, and WAR provides a coarse measure of size that correlates with orbital overlap. The site-specific ψ descriptor combines local electronegativity with valence counts of nearest neighbors, providing a compact representation of the immediate electronic environment and a strong predictor of hydrogen binding energy [32].

Together, our descriptor set provide a minimal yet transparent representation of adsorption environments, built entirely from basic quantities available a priori. They capture both elemental identity and local arrangement without heavy featurization or system-specific metrics, ensuring transferability across compositions, surfaces, and motifs. Importantly, each input can be directly rationalized in terms of chemisorption principles and ensemble effects in alloys, thereby enabling mechanistic interpretation alongside predictive accuracy.

3. Results and Discussion

3.1. Exploratory Data Analysis

After constructing the HER subset and defining the descriptor set, an exploratory data analysis was performed prior to model training. Elemental coverage, adsorption site labels, and the ranges of key descriptors (WAR, WEN, WIE, GCN, and Ψ) were examined to establish baseline expectations for model uncertainty and to guide data partitioning (Table S1). This ensured that subsequent learning was based on a consistently structured dataset.

The elemental distribution is shown in Figure 2A, with Au-containing entries highlighted. Broad coverage across late transition metals is evident, and the Au subset contributes with ~405 entries, sufficient for assessing Au-based systems independently. The distribution of adsorption site types is summarized in Figure 2B, where hollow motifs (*fcc* and *hcp*) dominate, while bridge and ontop sites are comparatively underrepresented. The imbalance suggests higher predictive confidence for hollow sites and motivates the use of site-resolved diagnostics in subsequent model evaluation. A schematic representation of the four characteristic adsorption site types (*fcc*, *hcp*, bridge, and ontop) considered in this study is provided in Figure S1.

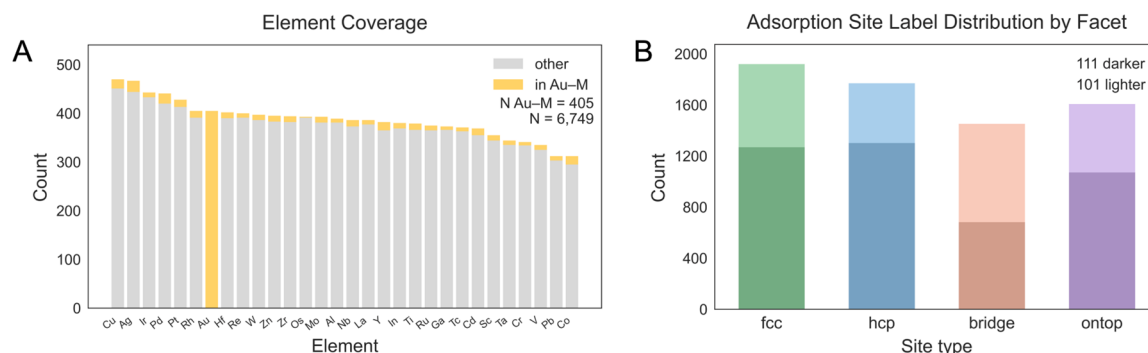


Figure 2. (A) Elemental distribution across the dataset [19] with Au-containing entries highlighted in yellow (405 rows). (B) Frequency of adsorption motifs (ontop, bridge, *fcc* hollow, and *hcp* hollow), showing dominance of close-packed hollow sites relative to bridge and ontop motifs. These distributions illustrate the chemical diversity of the dataset and the relative scarcity of undercoordinated sites, which might contribute to site-dependent variability in model performance.

Descriptor distributions were further analyzed to confirm that the feature set captures distinct structural and electronic environments. Histograms of the engineered descriptors reveal broad variation across WAR, WEN, WIE, Ψ , and GCN values, reflecting the chemical diversity of the curated dataset (Figure S2). These checks confirmed that the dataset is chemically rich and well balanced for Au systems, while also highlighting motif-dependent sampling biases that influence predictive uncertainty.

The distribution of hydrogen adsorption energies (E_{ads}) in the processed HER subset is shown in Figure 3A. Adsorption values span from approximately -2.0 to $+3.0$ eV, reflecting the diversity of local atomic environments sampled in the dataset. The histogram is skewed toward negative energies, with a broad maximum near the practical target of -0.24 eV. This skewness arises primarily from relaxation behavior during DFT optimization, where unstable high-energy configurations frequently relax into more stable hollow sites. In addition, the prevalence of threefold hollow motifs on *fcc*-derived surfaces contributes to the overrepresentation of strongly binding configurations in the dataset.

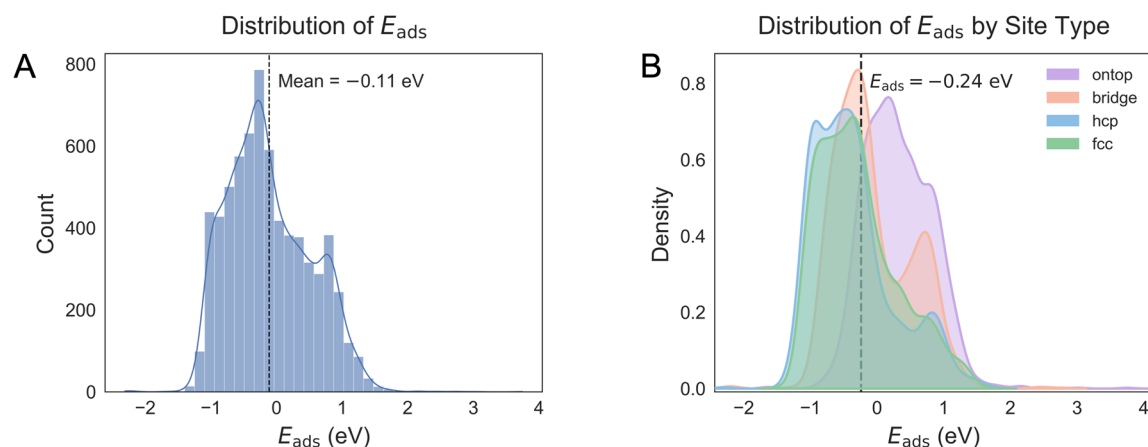


Figure 3. (A) Distribution of hydrogen adsorption energies (E_{ads}) in the curated HER dataset, spanning from approximately -2.0 to $+3.0$ eV with a peak density near the Sabatier optimum (-0.24 eV). (B) Site-resolved distributions, showing that hollow motifs (*fcc* and *hcp*) cluster near optimal binding, while ontop sites are shifted toward weaker binding with broader positive tails. These trends highlight the strong dependence of adsorption energetics on local coordination environment.

Site dependence of adsorption energetics is further illustrated in Figure 3B, which separates the distributions by adsorption motif. *fcc* and *hcp* hollow sites cluster on the negative side of the distribution and display the narrowest spreads, consistent with their thermodynamic stability [33]. Bridge sites shift toward weaker binding and exhibit a secondary lobe at positive energies, while ontop sites extend even further into positive values with the broadest tails. These long positive tails reflect the reduced stability of undercoordinated motifs and their high sensitivity to subtle changes in the local atomic environment. The magnitude of these site-dependent shifts, reaching up to ~ 1 eV between motifs, is consistent with prior reports and underscores the need for encoding of site

identity in the feature set (Table S1) [33,34]. Together, these observations justify the inclusion of coordination-sensitive descriptors and stratified model validation to ensure that predictive accuracy is assessed separately for each adsorption motif rather than in aggregate.

3.2. Correlation Analysis

After defining the feature set, relationships among descriptors were examined prior to model training. Pairwise Pearson correlation coefficients were calculated for all descriptors and for the target variable, as summarized in Figure 4. Strong positive correlations were observed among WIE, WEN, and S_A , reflecting their shared origin in periodic trends. Similarly, the unit cell volume per atom (\AA^3) tracked closely with the WAR, with a correlation coefficient of 0.80, as expected for two descriptors that encode atomic size. By contrast, the GCN and facet label exhibited negligible correlation with other descriptors (Table S2), indicating that geometric terms provide largely independent information relative to element-based quantities. Adsorption energy itself showed only weak linear correlations with any single descriptor, reinforcing the need for a multivariate ML model rather than a screening approach based on individual variables.

This correlation analysis served two main purposes. It first confirmed that geometric and electronic descriptors provide complementary information for modeling adsorption behavior. It also revealed pairs of features whose collinearity could undermine the reliability of importance estimates, as correlated variables may share or distort their relative contributions. These findings informed our modeling strategy, where we represented site geometry through the GCN and explicitly retained the facet label to preserve essential structural distinctions while avoiding one-hot encoding. This approach follows recent HER-focused ML studies in which compact descriptor sets are systematically analyzed for collinearity, and tree ensemble methods are favored for their ability to capture non-linear interactions while remaining robust to between-feature correlation. Complete distributions of all descriptors and additional statistics are provided in Figure S2. Together, these checks ensured that the final feature set is compact, interpretable, and physically meaningful, thereby establishing a robust foundation for predictive modeling.




Figure 4. Pearson correlation matrix of the engineered descriptors and hydrogen adsorption energy (E_{ads}). Strong correlations are observed between weighted electronegativity (WEN), weighted ionization energy (WIE), and valence electron count (S_A), while geometric descriptors such as generalized coordination number (GCN) and facet remain largely independent. The weak correlation between any single descriptor and E_{ads} emphasizes the need for multivariate ML models to capture adsorption trends.

3.3. Model Development

The prediction of hydrogen adsorption energies was framed as a supervised regression problem [35]. To avoid data leakage, data were partitioned using a group-aware strategy in which surface composition–facet–site identifiers served as grouping keys. This ensured that chemically identical motifs did not appear in both training

and test sets. An 80/20 data split was used for model evaluation, and StratifiedGroupKFold cross-validation was applied within the training set to ensure both balanced representation of Au-containing entries and group integrity. As summarized in Table S3, the mean absolute error (MAE) across folds remains consistent, indicating stable model behavior and confirming the effectiveness of the grouping strategy in preventing data leakage.

Three tree-based ensemble algorithms were evaluated: Random Forest (RF) [36], Extra Trees (ET) [37], and Extreme Gradient Boosting (XGBoost) [38]. The chosen models all rely on ensembles of decision trees, with RF and ET using randomized averaging strategies to reduce variance, while XGBoost builds trees sequentially via gradient boosting to minimize prediction error. These methods were selected for their ability to capture non-linear relationships among features while providing interpretable measures of feature importance. Hyperparameters for each model were optimized using grid search with cross-validation, and the final settings are listed in Table S4. Model performance was primarily assessed MAE, with root-mean-square error (RMSE) and coefficient of determination (R^2) reported as secondary metrics.

Training on the full alloy dataset yielded quantitative accuracy sufficient to resolve motif-dependent differences in adsorption behavior. The Extra Trees algorithm exhibited the best overall performance, achieving a mean average error of 0.131 eV ($R^2 = 0.86$) on the full test set and 0.136 eV ($R^2 = 0.82$) for the Au subset. This level of accuracy is sufficient to distinguish between weak, optimal, and strong hydrogen binding regimes, which is the relevant resolution for catalytic screening. RF and XGBoost showed slightly higher errors, with ~0.146 and ~0.149 eV, respectively. Site-resolved diagnostics revealed strong heterogeneity in predictive accuracy (Table S5): hollow sites (*fcc*, *hcp*) displayed the lowest errors (0.07 eV), while ontop sites exhibited errors exceeding 0.25 eV. Attempts to improve performance using site-specific ET submodels, each trained and tuned separately while retaining the same descriptor set, did not outperform the unified model (Table S6). This confirms that a global framework incorporating site descriptors remains preferable for high-throughput screening.

These results demonstrate that a compact set of physically interpretable descriptors, when combined with tree-based ensemble models, can achieve prediction errors on the order of 0.1 eV for stable adsorption motifs. The higher errors observed for ontop sites likely arise from their intrinsic instability in the underlying DFT dataset. As shown by the plateau of the learning curves around 3500 training points (Figure S3), the limited sampling of undercoordinated motifs appears to have only a minor impact on overall accuracy. Instead, the model performance seems constrained by the informational capacity of the current descriptor set, highlighting the need for coordination-sensitive features to better capture adsorption trends.

3.4. Model Evaluation

The performance of the Extra Trees (ET) model was examined in detail through residual and feature attribution analyses. Residual plots (Figure 5A) show that errors are symmetrically distributed around zero with no significant systematic bias. A modest widening of the residual spread at higher adsorption energies reflects heteroscedasticity in the underlying dataset. Site-resolved residuals (Figure 5B) confirm the motif-dependent accuracy: *fcc* and *hcp* hollow sites exhibit narrow, centered distributions, while ontop sites display the largest variance, consistent with their intrinsic instability and sparse representation in the training data.

Feature attribution was investigated using both Gini importance and SHAP (SHapley Additive exPlanations) values (Figure 6). Gini importance captures the average decrease in model error from splits on each feature, while SHAP values estimate the marginal contribution of each feature to individual predictions. These methods exhibit which chemical descriptors the model relies on most, and in which direction they influence adsorption strength. The Gini ranking identified the ψ as the most influential input, followed by GCN, WEN, and S_A . Secondary contributions were observed from WAR, volume per atom, and WIE, while facets contributed modestly. Collinearity among certain descriptors tends to redistribute feature importance within correlated groups, which can mask the apparent contribution of secondary variables. Nevertheless, the overall ranking of descriptors and the chemical interpretability of the model remain largely unaffected. SHAP analysis provided additional mechanistic insight: higher GCN and WEN values shifted predictions toward stronger binding (more negative E_{ads}), whereas larger S_A and volume per atom were associated with weaker binding. Notably, ψ exhibited a broad SHAP distribution with both positive and negative contributions, highlighting its ability to capture non-linear interactions between local electronic environment and coordination.




Figure 5. (A) Residuals of the Extra Trees (ET) model plotted against predicted adsorption energies (\hat{E}_{ads}), showing symmetric distribution around zero and modest heteroscedasticity at higher energies. (B) Site-resolved residuals for *fcc*, *hcp*, *bridge*, and *ontop* motifs. Narrow, centered distributions for *fcc* and *hcp* sites indicate reliable predictions, while *ontop* sites display the largest variance, reflecting their intrinsic instability.




Figure 6. Feature attribution analysis of the Extra Trees (ET) model. (A) Gini importance ranking, highlighting Ψ , generalized coordination number (GCN), and weighted electronegativity (WEN) as dominant descriptors. (B) SHAP summary plot, which provides directional insight: higher GCN and WEN values drive predictions toward stronger binding (more negative E_{ads}), while larger S_A and volume per atom shift predictions toward weaker binding. The broad distribution of Ψ underscores its role in capturing non-linear site–composition interactions. Agreement between the two approaches confirms that physically motivated descriptors govern model predictions.

The strong agreement between Gini and SHAP analyses underscores that the model relies on chemically interpretable features to reproduce adsorption trends [39]. In particular, the prominence of Ψ and GCN is consistent with established catalytic intuition, as adsorption strength is governed jointly by coordination environment and electronic character. These findings confirm that the ML framework not only achieves quantitative accuracy but also produces insights that can be rationalized within well-known chemisorption principles.

3.5. Discussion and Implications

The results demonstrate that hydrogen adsorption on Au-based alloys can be predicted with quantitative accuracy using compact, chemical interpretable descriptors combined with tree-based ensemble models. The achieved mean absolute error of ~ 0.1 eV is sufficient to differentiate weak, optimal, and strong binding regimes, providing a practical level of resolution for catalytic prescreening. Site-resolved diagnostics confirm that the most stable motifs (*fcc* and *hcp* hollow sites) are predicted with highest fidelity, while undercoordinated *ontop* sites remain a source of uncertainty. This performance profile mirrors the intrinsic thermodynamic variability of adsorption motifs in the underlying DFT dataset, highlighting both the strengths and the limitations of data-driven modeling in catalysis.

A key outcome of this work is the demonstration that predictive accuracy does not require large descriptor sets or complex neural architectures. Instead, the most influential features, coordination number and site-specific electronic descriptors, are directly tied to well-established chemisorption principles. This interpretability allows

experimental chemists to rationalize predictions within familiar frameworks, thereby lowering the barrier to adopting ML tools in catalyst discovery. Such predictions can help experimentalists prioritize alloy families before investing in DFT validation or synthesis. The workflow presented here thus complements intuition and conventional high-throughput screening by offering rapid, reproducible insights into alloy composition–structure–property relationships.

At the same time, limitations must be acknowledged. The training data originate from static DFT calculations and do not capture solvation, applied potential, or long-term surface restructuring under operating conditions [40]. The reduced accuracy for ontop sites reflects both their intrinsic instability and their underrepresentation in the dataset, suggesting that further improvements will require expanded sampling of undercoordinated environments. Broader validation beyond Au-based alloys will also be necessary to establish the generality of this approach.

Despite these caveats, the present study provides a proof of concept for integrating ML into catalytic discovery pipelines in a manner that is accessible to chemists. By focusing on chemically meaningful features and transparent workflows, the approach serves as a reference model for bridging computational screening and experimental exploration of electrocatalysts for the HER and related transformations.

Although the present study focuses on Au-based alloys as a proof of concept, the proposed ML framework is composition-agnostic and readily transferable to other metallic systems. The nine engineered descriptors capture general geometric and electronic factors, such as coordination environment, electronegativity, valence electron count, and local site potential, that are not specific to Au chemistry. Consequently, the same workflow can be directly retrained on DFT datasets for Pt-, Pd-, or other transition-metal catalysts without modification of the feature set or model architecture. Because these descriptors encode fundamental chemisorption principles, their physical interpretability ensures consistent predictive trends across different catalyst families. In this sense, the present work provides a transferable template for integrating interpretable machine learning into the discovery of diverse HER catalysts beyond Au-based systems.

While the present framework is trained on DFT-calculated adsorption energies, its purpose is to decouple predictive screening from the need for additional first-principles calculations. Once trained, the model can rapidly estimate adsorption trends across uncalculated alloy compositions at negligible computational cost, providing scalability that would be unattainable through DFT alone. The use of descriptors restricted to the first coordination shell intentionally balances accuracy and generalizability: local atomic and electronic environments dominate hydrogen binding, as reflected by the quantitative accuracy achieved here (MAE \approx 0.13 eV). This design allows the framework to retain physical interpretability while remaining computationally lightweight and transferable to broader catalyst discovery tasks.

Although the training data comprise ordered *fcc*-based alloys, the feature design is inherently local and therefore independent of long-range periodicity. Each descriptor is constructed from first-shell coordination and elemental properties, quantities that can be equally obtained for amorphous or defect-rich structures from atomic coordinates or short-range order parameters. As such, the same ML workflow can, in principle, be applied to amorphous catalysts without modification, offering a generalizable route for modeling hydrogen adsorption in both crystalline and non-crystalline materials.

3.6. Catalyst Screening

Guided by the feature attribution analyses and the site-resolved error diagnostics, the final stage applied the trained Extra Trees model to a targeted library of Au–M alloys. Partner elements were drawn from the 3d, 4d, and 5d transition series, together with representative p-block and rare-earth elements, to ensure broad chemical coverage while remaining within the interpolation regime of the training set. The chosen set included Sc, Ti, V, Cr, Mn, Fe, Co, Ni, Cu, Zn, Y, Zr, Nb, Mo, Tc, Ru, Rh, Pd, Ag, Cd, Hf, Ta, W, Re, Os, Ir, Pt, Hg, Al, Bi, Ga, In, La, Pb, Sn, and Tl. For each Au–M pair, compositions in AB, AB₃, and A₃B stoichiometries were generated, preserving local coordination environments consistent with those in the curated dataset. Adsorption motifs were expanded to include ontop, bridge, *fcc* hollow, and *hcp* hollow sites, with further enumeration into first-layer variants (e.g., *fcc*AAA, *fcc*AAB, *fcc*ABB, *fcc*BBB), as summarized in Table S7. In total, 1404 candidate rows were constructed, of which 1003 were absent from the training data and thus represent a genuine screening space.

To evaluate thermodynamic feasibility, DFT calculations were performed on representative screened alloys. Formation energies are presented in Figure 7A, where early 4d and 5d partners (Zr, Nb, Mo, Tc, Hf, Ta, W) yield strongly negative values (as low as −0.45 eV), consistent with robust heteronuclear bonding and a high likelihood of forming ordered intermetallics [41]. Chromium also exhibits negative formation energies near equiatomic compositions. In contrast, late transition partners (Rh, Pd, Ag) and light elements (Ti, V) display near-zero or slightly positive formation energies, reflecting reduced alloy stability. A clear compositional trend is observed:

many systems show maximum stability near 50% M (AB-type ordering), with only mild destabilization at Au-rich or M-rich extremes.

Cohesive energies, shown in Figure 7B, further reinforce these trends. Early transition partners (Mo, W, Ta, Hf, Tc) exhibit strong cohesion (-6 to -7 eV), suggesting resistance to dealloying and long-term structural integrity. In contrast, Au–Pd and Au–Ag exhibit weaker cohesion (-4 to -5 eV), indicating reduced stability under corrosive conditions [42]. Integrating adsorption and stability maps identifies Au–Mo, Au–W, Au–Ta, Au–Hf, and Au–Tc as particularly promising candidates that combine near-optimal hydrogen adsorption with robust alloy stability.

Together, these results provide a chemically grounded roadmap for Au-based HER catalyst design. Au-rich stoichiometries (A_3B , AB_3) yield surfaces resistant to corrosion while maintaining favorable adsorption energies at hollow sites, while equiatomic AB systems provide bulk stabilization. Reference alloys such as Au–Pd and Au–Ag remain useful benchmarking systems but are less likely to withstand aggressive electrochemical environments. This integrated ML–DFT workflow demonstrates how interpretable machine learning, coupled with targeted validation, can accelerate the identification of practical electrocatalyst candidates.




Figure 7. Thermodynamic stability of Au–M alloys across three stoichiometries (A_3B , AB , AB_3). (A) Formation energies, showing negative values for early 4d and 5d partners (e.g., Mo, W, Ta, Hf, Tc, Nb), consistent with strong heteronuclear bonding and ordered intermetallic formation. Late transition partners (Pd, Ag, Rh) and light elements (Ti, V) yield values near zero or positive, reflecting reduced alloy stability. (B) Cohesive energies, which strengthen progressively from late to early transition series. Au–Mo, Au–W, Au–Ta, and Au–Hf exhibit high cohesive energies (-6 to -7 eV), while Au–Pd and Au–Ag fall closer to -4 to -5 eV. Combined, the formation and cohesive energy maps identify Au–Mo, Au–W, Au–Ta, Au–Hf, and Au–Tc as promising candidates that couple robust alloy stability with favorable hydrogen adsorption.

4. Conclusions

This study presents a proof-of-concept machine learning framework for predicting hydrogen adsorption energies on Au-based alloys, built on a curated dataset and nine chemically interpretable descriptors. Using tree-based models, we show that catalytic trends across different compositions and adsorption motifs can be captured in a transparent and reproducible manner. Feature attribution highlights coordination environment and local electronic structure as the dominant factors, in agreement with established chemisorption principles. Applying the trained model to a targeted Au–M alloy library, followed by DFT-based stability validation, identified promising candidates such as Au–Mo, Au–W, Au–Ta, Au–Hf, and Au–Tc. Beyond these specific predictions, the workflow

provides an accessible reference for experimental chemists, offering a practical way to prioritize alloy families for further DFT calculations or synthesis. More broadly, this work illustrates how interpretable ML can accelerate the discovery of HER catalysts and related materials.

Supplementary Materials: The following supporting information can be downloaded at: <https://media.sciltp.com/articles/others/2511241531482990/MI-25100039-SI.pdf>, Figure S1: Schematic illustration of the four characteristic adsorption site types, ontop, bridge, *fcc* hollow, and *hcp* hollow, on an *fcc* (111) surface; Figure S2: Histograms of engineered descriptors in the HER-only training set ($N=1394$); Figure S3: Learning curve analysis for Extra Trees, Random Forest, and XGBoost models; Table S1: Mapping of coefficients used to calculate the site-specific electronic descriptor (Ψ); Table S2: Values of generalized coordination number (GCN) for canonical adsorption sites (*fcc*, *hcp*, bridge, ontop) on *fcc*(111) and *fcc*(101) surfaces, adapted from Martínez-Alonso et al. (2024); Table S3: Per-fold mean absolute error (MAE, eV) for Extra Trees (ET), Random Forest (RF), and XGBoost models under the 5-fold StratifiedGroupKFold split; Table S4: Optimized hyperparameters for Extra Trees, Random Forest, and XGBoost models, together with the ranges explored during grid search; Table S5: Site-resolved test metrics for the unified Extra Trees model; Table S6: Performance of site-specific models trained separately on ontop, bridge, *fcc*, and *hcp* subsets; Table S7: Canonical adsorption site types used in this study (ontop, bridge, *fcc*, *hcp*) and the corresponding normalized site-label variants grouped under each type. Reference [43] is cited in Supplementary Materials.

Author Contributions: M.K. and S.W.: conceptualization, methodology, visualization, software, data curation, writing—original draft preparation; P.H.C.C.: supervision; conceptualization, writing—reviewing and editing. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Jane and Aatos Erkko Foundation (grant no. 240052), and the Academy of Finland (decision no. 334826).

Data Availability Statement: All scripts and notebooks used for data preprocessing, feature engineering, model training, and alloy screening are openly available at <https://github.com/majakubik/her-au-ml-code> (accessed on 10 November 2025). The repository contains step-by-step Jupyter notebooks reproducing the main results and figures of this study, together with example input data and trained models. This ensures full reproducibility of our workflow and provides an accessible reference for researchers seeking to apply chemically interpretable machine-learning frameworks in catalysis.

Conflicts of Interest: The authors declare no conflict of interest.

Use of AI and AI-Assisted Technologies: AI tools were employed for text editing. After using this, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

- Gao, Y.; Xu, Y.; Guo, H.; Li, J.; Ding, L.; Wang, T.; He, J.; Chang, K.; Wu, Z. A 17.73% Solar-To-Hydrogen Efficiency with Durably Active Catalyst in Stable Photovoltaic-Electrolysis Seawater System. *Angew. Chem. Int. Ed.* **2025**, *64*, e202420814. <https://doi.org/10.1002/anie.202420814>.
- Nsanzimana, J.M.V.; Cai, L.; Djire, A.; Pagot, G.; Mattana, P.; Vezzù, K.; Negro, E.; Xia, B.Y.; Di Noto, V. Tailoring Chemical Microenvironment of Iron-Triad Electrocatalysts for Hydrogen Production by Water Electrolysis. *Adv. Energy Mater.* **2025**, *15*, 2501686. <https://doi.org/10.1002/aenm.202501686>.
- Tang, J.; Guo, K.; Guan, D.; Hao, Y.; Shao, Z. A Semi-Vapor Electrolysis Technology for Hydrogen Generation from Wide Water Resources. *Energy Environ. Sci.* **2024**, *17*, 7394–7402. <https://doi.org/10.1039/D4EE02722A>.
- Zhang, X.; Cao, C.; Ling, T.; Ye, C.; Lu, J.; Shan, J. Developing Practical Catalysts for High-Current-Density Water Electrolysis. *Adv. Energy Mater.* **2024**, *14*, 2402633. <https://doi.org/10.1002/aenm.202402633>.
- Liu, Z.; Zhang, L.; Zheng, C.J.; Zhang, Y.; Chen, B.; Shao, Z.; Ge, J. Advanced Electrode Materials for Efficient Hydrogen Production in Protonic Ceramic Electrolysis Cells. *Adv. Mater.* **2025**, 2503609. <https://doi.org/10.1002/adma.202503609>.
- Brunin, G.; Ricci, F.; Ha, V.-A.; Rignanese, G.-M.; Hautier, G. Transparent Conducting Materials Discovery Using High-Throughput Computing. *NPJ Comput. Mater.* **2019**, *5*, 63. <https://doi.org/10.1038/s41524-019-0200-5>.
- Wang, C.; Wang, B.; Wang, C.; Chang, Z.; Yang, M.; Wang, R. Efficient Machine Learning Model Focusing on Active Sites for the Discovery of Bifunctional Oxygen Electrocatalysts in Binary Alloys. *ACS Appl. Mater. Interfaces* **2024**, *16*, 16050–16061. <https://doi.org/10.1021/acsami.3c17377>.
- Mai, H.; Le, T.C.; Chen, D.; Winkler, D.A.; Caruso, R.A. Machine Learning for Electrocatalyst and Photocatalyst Design and Discovery. *Chem. Rev.* **2022**, *122*, 13478–13515. <https://doi.org/10.1021/acs.chemrev.2c00061>.
- Jeong, I.; Shim, Y.; Oh, S.; Yuk, J.M.; Roh, K.; Lee, C.; Lee, K.T. A Machine Learning-Enhanced Framework for the Accelerated Development of Spinel Oxide Electrocatalysts. *Adv. Energy Mater.* **2024**, *14*, 2402342. <https://doi.org/10.1002/aenm.202402342>.
- Chen, L.; Tian, Y.; Hu, X.; Yao, S.; Lu, Z.; Chen, S.; Zhang, X.; Zhou, Z. A Universal Machine Learning Framework for Electrocatalyst Innovation: A Case Study of Discovering Alloys for Hydrogen Evolution Reaction. *Adv. Funct. Mater.* **2022**, *32*, 2208418. <https://doi.org/10.1002/adfm.202208418>.
- Liu, J.; Luo, W.; Wang, L.; Zhang, J.; Fu, X.; Luo, J. Toward Excellence of Electrocatalyst Design by Emerging Descriptor-Oriented Machine Learning. *Adv. Funct. Mater.* **2022**, *32*, 2110748. <https://doi.org/10.1002/adfm.202110748>.

12. Jung, W.; Choi, J.; An, S.; Yun, S.; Chung, D.S.; Cha, H.; Lim, J.; Park, T. Photocatalytic Hydrogen Evolution with Conjugated Polymers: Structure–Property Insights and Design Strategies. *Adv. Energy Mater.* **2025**, *15*, 2501600. <https://doi.org/10.1002/aenm.202501600>.
13. Lee, W.J.; Kwak, H.S.; Lee, D.; Oh, C.; Yum, E.K.; An, Y.; Halls, M.D.; Lee, C.-W. Design and Synthesis of Novel Oxime Ester Photoinitiators Augmented by Automated Machine Learning. *Chem. Mater.* **2022**, *34*, 116–127. <https://doi.org/10.1021/acs.chemmater.1c02871>.
14. Nuñez, M. Exploring Materials Band Structure Space with Unsupervised Machine Learning. *Comput. Mater. Sci.* **2019**, *158*, 117–123. <https://doi.org/10.1016/j.commatsci.2018.11.002>.
15. Wang, C.; Wang, B.; Wang, C.; Li, A.; Chang, Z.; Wang, R. A Machine Learning Model with Minimize Feature Parameters for Multi-Type Hydrogen Evolution Catalyst Prediction. *NPJ Comput. Mater.* **2025**, *11*, 111. <https://doi.org/10.1038/s41524-025-01607-4>.
16. Wang, Z.; Yang, M.; Xie, X.; Yu, C.; Jiang, Q.; Huang, M.; Algadi, H.; Guo, Z.; Zhang, H. Applications of Machine Learning in Perovskite Materials. *Adv. Compos. Hybrid. Mater.* **2022**, *5*, 2700–2720. <https://doi.org/10.1007/s42114-022-00560-w>.
17. Mou, L.; Han, T.; Smith, P.E.S.; Sharman, E.; Jiang, J. Machine Learning Descriptors for Data-Driven Catalysis Study. *Adv. Sci.* **2023**, *10*, 2301020. <https://doi.org/10.1002/advs.202301020>.
18. Zhang, N.; Yang, B.; Liu, K.; Li, H.; Chen, G.; Qiu, X.; Li, W.; Hu, J.; Fu, J.; Jiang, Y.; et al. Machine Learning in Screening High Performance Electrocatalysts for CO₂ Reduction. *Small Methods* **2021**, *5*, 2100987. <https://doi.org/10.1002/smt.202100987>.
19. Mamun, O.; Winther, K.T.; Boes, J.R.; Bligaard, T. High-Throughput Calculations of Catalytic Properties of Bimetallic Alloy Surfaces. *Sci. Data* **2019**, *6*, 76. <https://doi.org/10.1038/s41597-019-0080-z>.
20. Chen, X.; Zhou, T.; He, T.; Liu, Q. Vacancy Engineering in the First Coordination Shell of Single-Atom Catalysts for Enhanced Hydrogen and Oxygen Evolution Reactions. *Small* **2025**, *21*, 2412000. <https://doi.org/10.1002/sml.202412000>.
21. Li, Y.; Zheng, S.; He, Y.; Yang, S.; Huang, W.-H.; Pao, C.-W.; Hu, Z.; Huang, X. Masked Second-Shell Sulfur Coordinating Atomically Dispersed Pd on Tin Oxide Boosts the Direct Synthesis of Hydrogen Peroxide. *Chem. Eng. J.* **2024**, *500*, 157297. <https://doi.org/10.1016/j.cej.2024.157297>.
22. Lv, Y.; Sun, W.; Luo, Q.; Gao, J.; Gu, G.; Ma, F. Fast Screening of High Anti-Corrosion Ta Ternary Alloys by Machine Learning and Electron-Level Descriptors. *Mater. Chem. Phys.* **2025**, *339*, 130820. <https://doi.org/10.1016/j.matchemphys.2025.130820>.
23. Tereshchenko, A.; Pashkov, D.; Guda, A.; Guda, S.; Rusalev, Y.; Soldatov, A. Adsorption Sites on Pd Nanoparticles Unraveled by Machine-Learning Potential with Adaptive Sampling. *Molecules* **2022**, *27*, 357. <https://doi.org/10.3390/molecules27020357>.
24. Shapera, E.P.; Bučar, D.-K.; Prasankumar, R.P.; Heil, C. Machine Learning Assisted Prediction of Organic Salt Structure Properties. *NPJ Comput. Mater.* **2024**, *10*, 176. <https://doi.org/10.1038/s41524-024-01355-x>.
25. Chen, K.; Kunkel, C.; Cheng, B.; Reuter, K.; Margraf, J.T. Physics-Inspired Machine Learning of Localized Intensive Properties. *Chem. Sci.* **2023**, *14*, 4913–4922. <https://doi.org/10.1039/D3SC00841J>.
26. Zou, Y.; Qian, J.; Wang, X.; Li, S.; Li, Y. Machine Learning-Assisted Prediction and Interpretation of Electrochemical Corrosion Behavior in High-Entropy Alloys. *Comput. Mater. Sci.* **2024**, *244*, 113259. <https://doi.org/10.1016/j.commatsci.2024.113259>.
27. Mishra, A.; Kompella, L.; Sanagavarapu, L.M.; Varam, S. Ensemble-Based Machine Learning Models for Phase Prediction in High Entropy Alloys. *Comput. Mater. Sci.* **2022**, *210*, 111025. <https://doi.org/10.1016/j.commatsci.2021.111025>.
28. Kapse, S.; Janwari, S.; Waghmare, U.V.; Thapa, R. Energy Parameter and Electronic Descriptor for Carbon Based Catalyst Predicted Using QM/ML. *Appl. Catal. B: Environ.* **2021**, *286*, 119866. <https://doi.org/10.1016/j.apcatb.2020.119866>.
29. Cerda, P.; Varoquaux, G.; Kégl, B. Similarity Encoding for Learning with Dirty Categorical Variables. *Mach. Learn.* **2018**, *107*, 1477–1494. <https://doi.org/10.1007/s10994-018-5724-2>.
30. Esterhuizen, J.A.; Goldsmith, B.R.; Linic, S. Uncovering Electronic and Geometric Descriptors of Chemical Activity for Metal Alloys and Oxides Using Unsupervised Machine Learning. *Chem. Catal.* **2021**, *1*, 923–940. <https://doi.org/10.1016/j.checat.2021.07.014>.
31. Lansford, J.L.; Vlachos, D.G. Spectroscopic Probe Molecule Selection Using Quantum Theory, First-Principles Calculations, and Machine Learning. *ACS Nano* **2020**, *14*, 17295–17307. <https://doi.org/10.1021/acsnano.0c07408>.
32. Chen, M.S.; Zuehlsdorff, T.J.; Morawietz, T.; Isborn, C.M.; Markland, T.E. Exploiting Machine Learning to Efficiently Predict Multidimensional Optical Spectra in Complex Environments. *J. Phys. Chem. Lett.* **2020**, *11*, 7559–7568. <https://doi.org/10.1021/acs.jpclett.0c02168>.
33. Wexler, R.B.; Martinez, J.M.P.; Rappe, A.M. Chemical Pressure-Driven Enhancement of the Hydrogen Evolving Activity of Ni₂P from Nonmetal Surface Doping Interpreted via Machine Learning. *J. Am. Chem. Soc.* **2018**, *140*, 4678–4683. <https://doi.org/10.1021/jacs.8b00947>.

34. Padama, A.A.B.; Palmero, M.A.; Shimizu, K.; Chookajorn, T.; Watanabe, S. Machine Learning and Density Functional Theory-Based Analysis of the Surface Reactivity of High Entropy Alloys: The Case of H Atom Adsorption on CoCuFeMnNi. *Comput. Mater. Sci.* **2025**, *247*, 113480. <https://doi.org/10.1016/j.commatsci.2024.113480>.
35. Pandit, N.K.; Roy, D.; Mandal, S.C.; Pathak, B. Rational Designing of Bimetallic/Trimetallic Hydrogen Evolution Reaction Catalysts Using Supervised Machine Learning. *J. Phys. Chem. Lett.* **2022**, *13*, 7583–7593. <https://doi.org/10.1021/acs.jpclett.2c01401>.
36. Dong, C.; Zhu, Y.; Qu, C.; Chen, Y.; Ma, Y.; Yu, Y.; Li, C. Coordination Chemistry-Driven Oxygen Vacancy Strategy for Rational Design of High-Performance Catalysts in BTX Oxidation. *Coord. Chem. Rev.* **2025**, *545*, 217007. <https://doi.org/10.1016/j.ccr.2025.217007>.
37. Chen, R.; Yang, J.; Yu, Y.; Liu, Z.; Wang, C.; Wen, Y.; Gao, Y.; Zhao, Y.; Sampara, C.S.; Li, W.; et al. Extra Trees Regression Assisted 1D Monolith Reactor Simulations Based on Microkinetic Analysis and Rate Transformation. *Chem. Eng. Sci.* **2025**, *302*, 120721. <https://doi.org/10.1016/j.ces.2024.120721>.
38. Da Costa, M.L.; Oviedo, L.R.; Franco, D.S.P.; Da Silva, W.L.; De Oliveira, J.S. Catalytic Ozonation for the Efficient Degradation of Tetracycline Using CoFe₂O₄@TiO₂ Ceramic Nanocomposite: Kinetic, Thermodynamic and Machine Learning Study. *Ceram. Int.* **2025**, *51*, 7143–7158. <https://doi.org/10.1016/j.ceramint.2024.12.149>.
39. Naqvi, S.K.H.; Chong, K.T.; Tayara, H. Machine Learning-Enhanced Analysis of Catalyst Particle Size Effects and Performance Prediction of Platinum on Carbon Electrocatalysts. *Comput. Mater. Sci.* **2025**, *259*, 114105. <https://doi.org/10.1016/j.commatsci.2025.114105>.
40. Panapitiya, G.; Avendaño-Franco, G.; Ren, P.; Wen, X.; Li, Y.; Lewis, J.P. Machine-Learning Prediction of CO Adsorption in Thiolated, Ag-Alloyed Au Nanoclusters. *J. Am. Chem. Soc.* **2018**, *140*, 17508–17514. <https://doi.org/10.1021/jacs.8b08800>.
41. He, X.; Liu, J.; Yang, C.; Jiang, G. Predicting Thermodynamic Stability of Magnesium Alloys in Machine Learning. *Comput. Mater. Sci.* **2023**, *223*, 112111. <https://doi.org/10.1016/j.commatsci.2023.112111>.
42. Cheng, G.; Gong, X.-G.; Yin, W.-J. An Approach for Full Space Inverse Materials Design by Combining Universal Machine Learning Potential, Universal Property Model, and Optimization Algorithm. *Sci. Bull.* **2024**, *69*, 3066–3074. <https://doi.org/10.1016/j.scib.2024.07.015>.
43. Martínez-Alonso, C.; Vassilev-Galindo, V.; Comer, B.M.; Abild-Pedersen, F.; Winther, K.T.; LLorca, J. Application of machine learning to discover new intermetallic catalysts for the hydrogen evolution and the oxygen reduction reactions. *Catal. Sci. Technol.* **2024**, *14*, 3784–3799.