

<https://doi.org/10.1038/s41524-025-01649-8>

Inverse design of catalytic active sites via interpretable topology-based deep generative models

Bingxu Wang¹, Shisheng Zheng²✉, Jie Wu^{3,4}, Jingyan Li³✉ & Feng Pan¹✉

The rational design of catalyst structures tailored to target performance is an ambitious and profoundly impactful goal. Key challenges include achieving refined representations of the three-dimensional structure of active sites and imbuing models with robust physical interpretability. Herein, we developed a topology-based variational autoencoder framework (PGH-VAEs) to enable the interpretable inverse design of catalytic active sites. Leveraging high-entropy alloys as a case, we demonstrate that persistent GLMY homology, an advanced topological algebraic analysis tool, enables the quantification of three-dimensional structural sensitivity and establishes correlations with adsorption properties. The multi-channel PGH-VAEs illustrate how coordination and ligand effects shape the latent space and influence the adsorption energies. Building on the inverse design results from PGH-VAEs, the strategies to optimize the composition and facet structures to maximize the proportion of optimal active sites are proposed. This interpretable inverse design framework can be extended to diverse systems, paving the way for AI-driven catalyst design.

Heterogeneous catalysis plays a central role in modern industry, driving advancements in energy conversion and environmental sustainability^{1,2}. The understanding and design of catalytic active sites, which are specific surface regions or atom groups that directly affect molecular adsorption³, are crucial, as they directly determine the efficiency, selectivity, and stability of catalytic processes. The intricate microstructural characteristics of active sites, arising from the combination of components, dynamic behavior, and spatial positioning, render the analysis of molecular adsorption states significant challenges, even for the simplest yet most critical adsorption energies⁴⁻⁷. High-throughput density functional theory (DFT) calculations and machine learning (ML) methods, commonly referred to as forward design, have been extensively utilized to establish structure-property relationships and predict adsorption properties⁸⁻¹⁰. These efforts have deepened our fundamental insights into heterogeneous catalysis, exemplified by the establishment of linear scaling relationships and volcano plots in various catalytic reactions¹¹⁻¹⁴. Once the optimal adsorption state for a specific catalytic reaction is understood, the focus naturally shifts to identifying which types of active sites can achieve this adsorption state beyond the benchmark catalysts. This process, referred to as inverse design, entails a transition from property to structure¹⁵. While the concept of inverse design has been around for many years, it has only recently begun to reveal its full

potential with the emergence of deep generative models. The variational autoencoder (VAE) and the generative adversarial network, two prominent generative models, have made significant strides in the fields of molecules and materials, facilitating the realization of automated “closed-loop” design processes aimed at achieving targeted performances¹⁶⁻²⁰. Yet, applications in heterogeneous catalysis remain scarce²¹.

The primary challenge is to accurately represent catalytic active sites, with two key factors contributing to their uncertainty and complexity: the variations in facets, defects, and size, referred to as the coordination effect²², and the random spatial distribution of different elements, known as the ligand effect²³. In a real catalyst, the coordination and ligand effects intertwine to create a complex and diverse distribution of catalytic active sites. An ideal representation would encode both coordination and ligand effects, be compatible with gradient-based optimization, and allow active sites to be reconstructed and decoded from the generative model. Popular expressions of catalytic active sites based on cheminformatics or graph have achieved significant progress but still face challenges in capturing the distant atomic effect and the overall three-dimensional structural complexity^{8,24-27}. Another challenge is that the “black box” nature of deep learning often makes generative models lacking in interpretability²⁸. In catalyst design, we seek not only to identify effective active sites but also to understand the reasons

¹School of Advanced Materials, Peking University Shenzhen Graduate School, Shenzhen, PR China. ²College of Energy, Xiamen University, Xiamen, PR China.

³Beijing Key Laboratory of Topological Statistics and Applications for Complex Systems, Beijing Institute of Mathematical Sciences and Applications, Beijing, PR China. ⁴School of Mathematical Sciences, Hebei Normal University, Shijiazhuang, PR China. ✉e-mail: zhengss@xmu.edu.cn; jingyanli@bimsa.cn;

panfeng@pkusz.edu.cn

behind their effectiveness. This raises a fundamental question in heterogeneous catalysis: which physical properties of a catalyst surface dictate the chemisorption strength of adsorbates^{29–34}. Although generative models can reverse-engineer properties-structure relationships to identify new potential candidate materials, the use of high-dimensional latent spaces to represent data structures often means that the generation process lacks intuitive physical intuition, making it difficult to understand the specific mechanisms behind the model's output^{35,36}. Extracting hidden patterns from the latent space into interpretable formats can lead to testable theories and hypotheses, further advancing scientific understanding²⁸.

In this work, we developed a novel topology-based VAE framework (PGH-VAEs) to enable the high-resolution representation of catalytic active sites and interpretable inverse design, using *OH adsorption—a key step in the oxygen reduction reaction (ORR)—on IrPdPtRhRu high-entropy alloys (HEAs) as an example. HEAs have recently emerged as a promising approach for fine-tuning catalytic properties due to their extensive variety of active site types, stemming from the high variability in local structural composition and coordination^{37–41}. GLMY homology, proposed by Grigor'yan et al., is a generalized homology theory based on path complexes that extends classical homology to directed and non-symmetric structures^{42–44}. It enables the topological analysis of complex systems with directionality or asymmetry, making it particularly useful for capturing subtle structural features and sensitivity in crystalline structures. We employ persistent GLMY homology (PGH) to achieve a refined characterization of the three-dimensional spatial features of catalytic active sites. A multi-channel VAE with modules dedicated to encoding and decoding the coordination and ligand features is developed, enabling the latent design space of active sites to possess substantial physical meaning. Leveraging a semi-supervised learning framework, we achieved a high-precision VAE

model using only around 1100 DFT data points, attaining a remarkably low mean absolute error (MAE) of 0.045 eV in *OH adsorption energy predictions. We further elucidate how coordination and ligand effects, especially the distant atoms that do not directly contact with adsorbate, shape *OH adsorption states, offering targeted optimization strategies for HEA catalysts through inverse design. This proof-of-concept protocol establishes a solid foundation for interpretable inverse design of catalytic active sites and can be extended to any other catalytic processes and systems. The integration of topology-based descriptors with interpretable property-structure relationships renders the ML “black box” more transparent, bringing us a step closer to on-demand catalyst design rather than relying on traditional trial-and-error approaches.

Results

Workflow overview

The overall workflow of the PGH-VAEs is illustrated in Fig. 1. To characterize the intricate microstructure of catalytic active sites, we developed a topology-based descriptor enriched with chemical information, enabling a unified representation of coordination and ligand effects (Fig. 1a). Subsequently, we generated structures for *OH adsorption on HEAs with uniform elemental ratios across various crystal facets. As HEAs represent one of the most complex catalytic systems, employing them as a model enables an effective simulation of the multifaceted complexities found at active sites in real-world catalysts. Constrained by the substantial computational demands of DFT simulations, we are limited to conducting first-principles calculations on only a small selection of structures. To overcome this, we developed a semi-supervised ML model to enhance our dataset. Specifically, we start with a labeled database of adsorption sites with known energies obtained from DFT calculations. A lightweight and efficient ML model is trained on

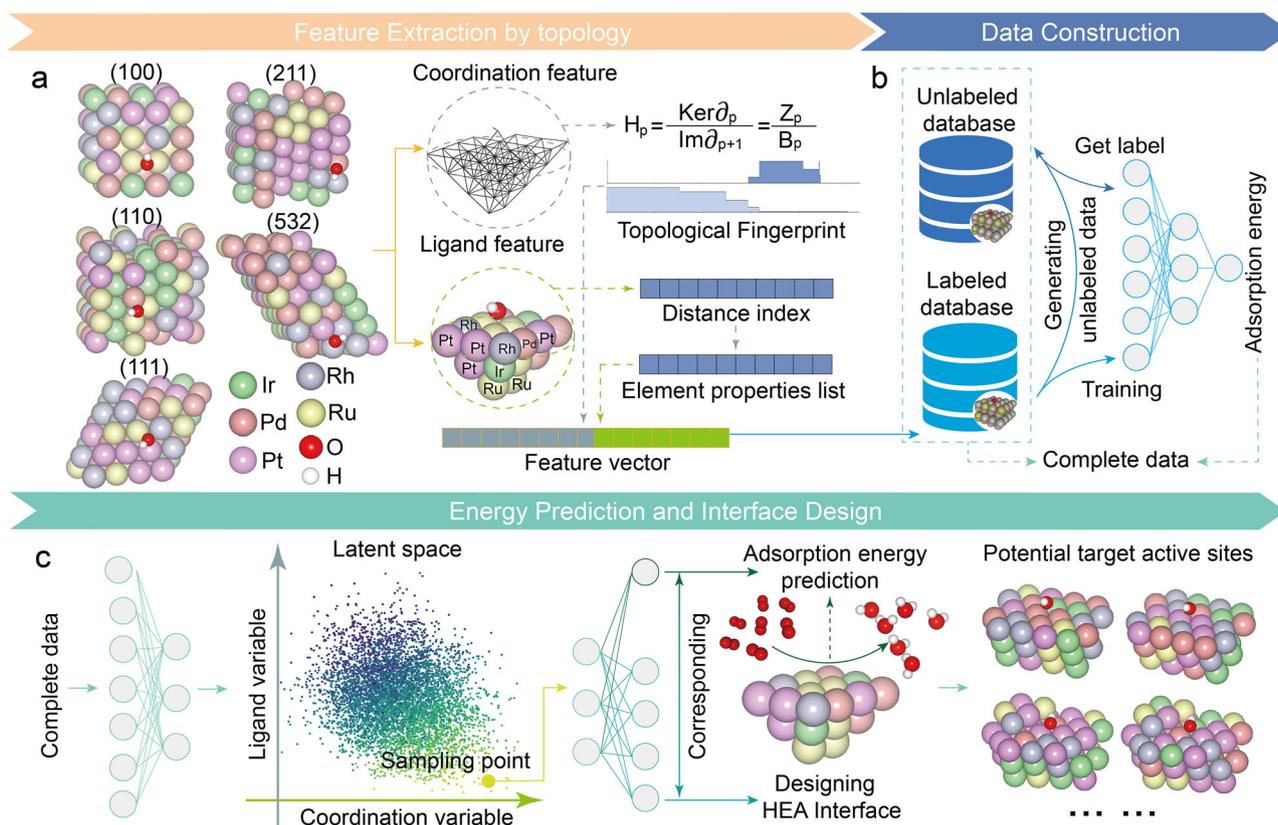


Fig. 1 | Overview of feature extraction, dataset construction, and workflow for energy prediction and interface design. **a** Schematic illustration of feature construction. Coordination features are extracted using the PGH method, while ligand features are represented based on element properties. **b** Schematic of dataset construction using DFT and semi-supervised learning. A GBR model was initially

trained on DFT-calculated adsorption energies and subsequently employed to predict adsorption energies for additional simulated active sites, enabling the construction of an expanded, pseudo-labeled dataset for model training. **c** Framework of PGH-VAEs, including modules for encoding, latent space visualization, latent space sampling, and decoding to generate potential active sites.

this dataset and subsequently used to predict the adsorption energies of newly generated structures that form the unlabeled database. By randomly generating unlabeled data and using the ML model trained on DFT data to label these samples—specifically by predicting their adsorption energies—we can effectively augment the dataset used for VAE training (Fig. 1b). This process transforms the original labeled data, the predicted adsorption energies, and the newly generated structures into a complete dataset, which is then used to train the VAE. Subsequently, a multi-channel VAE framework was introduced, with its latent space offering strong interpretability, allowing structure-performance relationships to be understood in terms of coordination and ligand effects. To avoid introducing bias into model evaluation, the unlabeled database was only used during training, and all testing was conducted exclusively on original DFT-calculated data. Further, the VAEs can generate novel active site structures tailored to specific adsorption energy criteria, offering insightful guidance for catalyst optimization and advancing the design of high-performance catalytic materials (Fig. 1c).

Active sites identification and representation

Catalytic active sites typically refer to specific regions or groups of atoms on the catalyst surface that directly influence the molecular adsorption state³. To maximize the diversity of active sites, we sampled on various Miller index surfaces, including (111), (100), (110), (211), and (532) of IrPdPtRhRu HEAs (Fig. 1a). The (111), (100), (110), (211), and (532) facets were selected because they are among the most commonly studied surfaces in ORR research^{45–47}, representing a diverse set of low-index and high-index surfaces that capture a range of atomic coordination environments commonly observed in transition metal catalysts. IrPdPtRhRu HEAs were selected for their stability, ease of synthesis, and demonstrated potential in ORR applications^{37,48,49}. According to Previous studies, the bridge site is considered to place the *OH adsorbate^{40,50}. Figure S1 shows all 13 unique bridge sites on these Miller index surfaces with their designation. The bridge site and the first and second-nearest neighbors of the bridge atoms are considered to constitute the primary chemical environment (active sites) influencing molecular adsorption^{51,52}.

Active sites are primarily influenced by coordination and ligand effects. Coordination effects refer to the spatial arrangement of atoms within the active site, encompassing structural features such as crystal facets, defects, and corner sites. We introduced PGH as a novel mathematical tool for capturing nuanced structural variations. Ker and Im represent the kernel and image of the boundary operators in PGH, capturing the cycles and boundaries in the topological structure. We include this formulation to illustrate the mathematical foundation of our topological fingerprinting approach. The detailed mathematical principles of PGH are presented in the Methods section. Here, we briefly introduce its logic and applications, using *OH adsorption on the (211) surface, denoted as (211) summit, as an example (Fig. 2a). First, the active site atoms are represented as a colored point cloud, with paths established based on bonding and element properties (group, period and atom radius) differences between points. Once paths are established, the atomic structure is converted into a path complex. The geometric characteristics of this path complex can be captured across various spatial scales through a process called filtration. As the filtration parameter (distance) increases, the number of visible paths expands, provided that the path lengths remain below the filtration threshold. Insets of Fig. 2a show adjacency matrices for path complexes at different filtration parameters. Simultaneously, the filtration process will generate the distance-based persistent GLMY homology (DPGH) fingerprint. Since each topological invariant (i.e., each Betti number) has its own persistence range, the number of recorded features can vary across different structures, leading to fingerprints of inconsistent dimensions. To resolve this issue, instead of recording each individual barcode, we discretize the continuous filtration parameter with a fixed step size of 0.1 Å and count the number of Betti numbers present at each discrete filtration value. These counts are then plotted as a line chart against the filtration parameter (the blue figure in Fig. 2a, b), and the resulting line chart is represented as a feature vector.

Additionally, we set the maximum filtration distance to 4 Å, which ensures that the filtration process is complete for all samples. As a result, the input feature vectors for all structures have a consistent dimensionality of 40. Therefore, by using this statistical summarization approach over persistence diagrams, we ensure that the resulting fingerprint vectors have a uniform length, making them compatible with ML models. In algebraic topology, Betti numbers are topological invariants that describe the number of independent features in different dimensions of a space. For a given dimension n , the Betti number is defined as the rank of the n th homology group H_n , and it is typically denoted by β_n . In our work, the Betti numbers have the following interpretations: (1) β_0 corresponds to the number of independent components (atoms), representing 0-dimensional features. (2) β_1 reflects the number of directed cycles composed of 1-paths (directed edges), capturing 1-dimensional topological structures. (3) β_2 denotes the number of voids in terms of paths composed of 2-paths, representing 2-dimensional features. This notation offers an interpretable representation of topological characteristics relevant to catalytic active sites across multiple geometric scales. The blue figure, representing the DPGH fingerprint, is the line plot of Betti numbers as a function of filtration parameters, which captures topological invariants of the catalytic site using GLMY homology. The horizontal axis denotes the filtration parameter (in Å), and the vertical axis shows the Betti numbers in dimensions 0, 1, and 2. The more detailed analysis of the DPGH is presented in Supporting Note 1. This filtration process highlights the structural intricacies and allows us to capture the topological features essential for active sites.

The application of PGH reveals distinctions between catalytic sites with similar coordination numbers but different structural features. Figure 2b presents another *OH adsorption configuration, denoted as (211) valley, on the (211) facet with the same neighbor coordination as the (211) summit in Fig. 2a. The representation based on coordination number miss such differences with the bridge atoms both having one 9-coordinated atom and one 10-coordinated atom (Fig. 2c) while the evolution of Betti numbers in Fig. 2b effectively differentiates this structure from that in Fig. 2a. Additional results for the other 11 active sites are presented in Fig. S2, offering a comprehensive analysis of structural features by PGH. Ultimately, incorporating the approximate treatment, mean values of β_0 , β_1 , and β_2 in the DPGH fingerprints are computed across the filtering process. These values are amalgamated into a 1×3 vector to serve as the coordinate feature for the active sites. For the ligand effect, we represent it directly based on the element properties of atoms (group, period, and atom radius), ordered by their spatial distance from the adsorbate (Fig. 2d). Ultimately, the vector representation of the active site is constructed by concatenating the coordination effect vector with the ligand effect vector, thereby achieving a dual description of the active site's topological and chemical properties.

The *OH adsorption dataset

Although there are no established guidelines on the optimal size of training datasets for deep generative models, empirical evidence suggests that larger datasets generally lead to higher-quality generated data¹⁶. Calculating the adsorption energy of *OH relies on DFT, which is highly computationally demanding. To address this, we introduce a semi-supervised learning framework to efficiently construct and expand the dataset. We constructed a dataset of 1159 DFT data points of *OH adsorption, randomly sampled to cover each unique bridge site across various Miller index surfaces. Notably, the adsorption energy here refers to the difference between the *OH adsorption energy on HEA active sites and that on Pt(111). It is widely recognized that an adsorption energy within 0.1 eV above that of Pt(111) can facilitate *OH desorption, thereby enhancing ORR performance^{37,53}. Using this dataset, we trained a Gradient Boosted Regression (GBR) model with the combined feature representation, achieving a test set MAE of 0.078 eV for *OH adsorption energy prediction (Table S1). The labeled dataset for semi-supervised VAE training was constructed by randomly sampling 1159 DFT-labeled structures across diverse coordination environments, ensuring a broad representation of

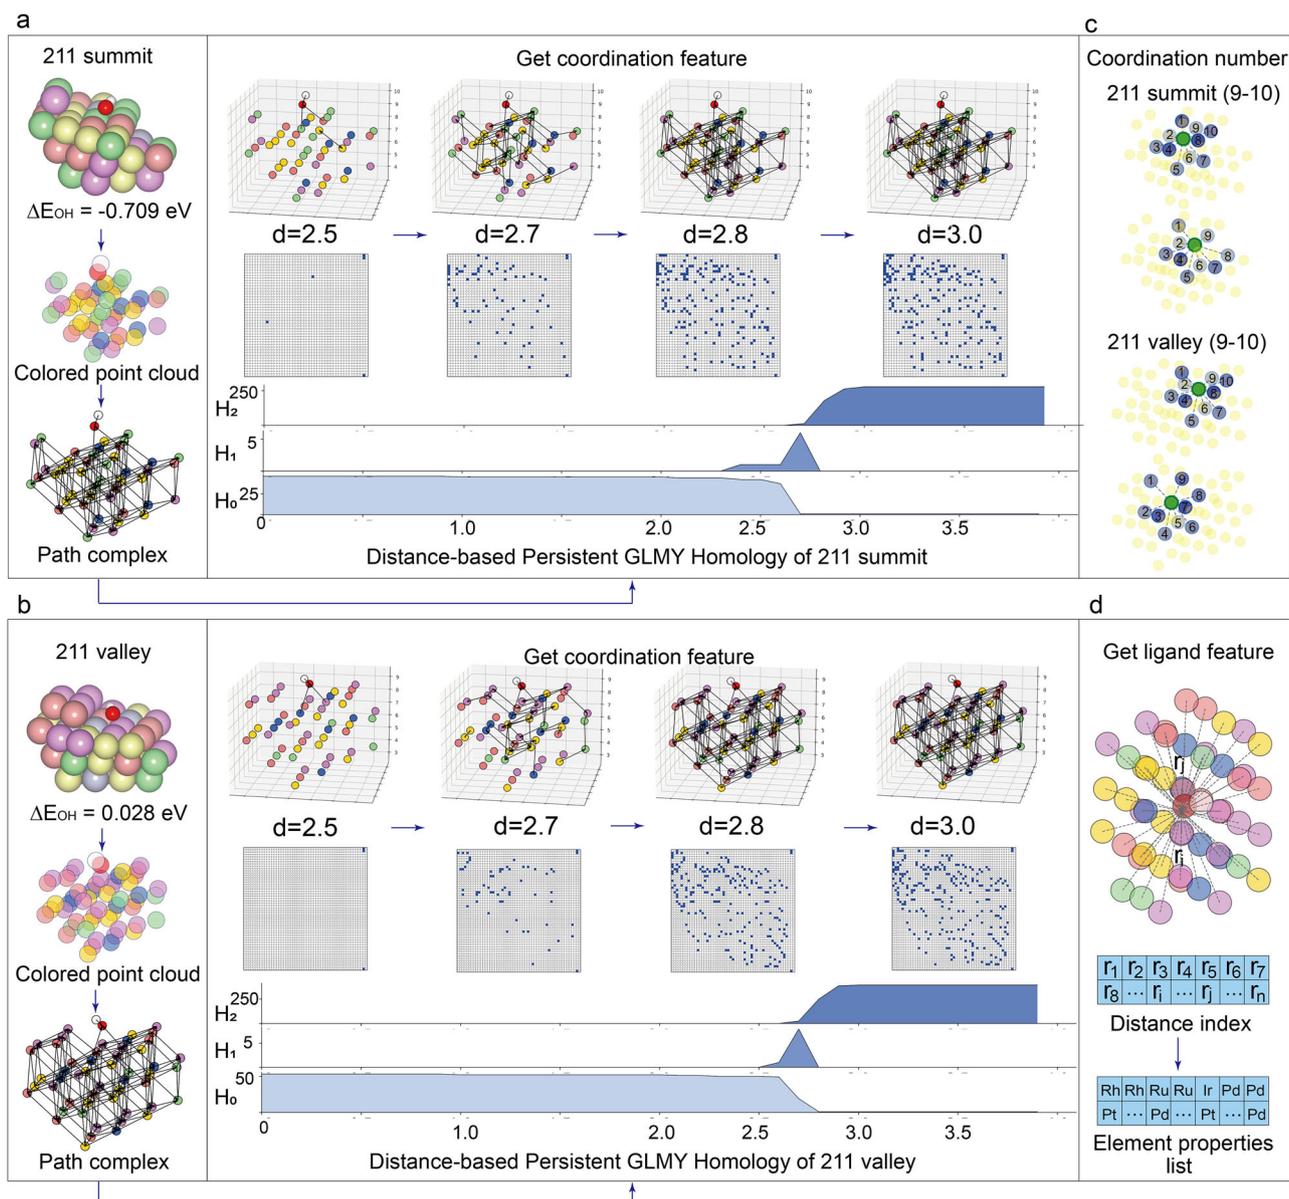


Fig. 2 | Workflow for coordination and ligand feature extraction of active sites.

a A schematic workflow for extracting the coordination features of the (211) summit bridge active site using PGH. Active site atoms are initially represented as a colored point cloud, with paths direction defined by bonding and electronegativity differences. These paths form a path complex, whose geometric features are analyzed across scales using filtration. As the filtration parameter increases, visible paths expand and are recorded, then converted into vectors for machine learning. The blue bar chart is the DPGH fingerprint. The DPGH fingerprint encapsulates the GLMY homology in dimensions 0 (H_0), 1 (H_1), and 2 (H_2). The vertical axis denotes the

value of the Betti number, while the horizontal axis represents the filtering parameter in angstroms (Å). **b** The workflow for extracting the coordination features of the (211) valley bridge active site. **c** Coordination numbers of neighboring atoms for the (211) summit and (211) valley bridge sites are shown. It can be observed that they share the same coordination environment, with the bridge atoms both having one 9-coordinated atom and one 10-coordinated atom. **d** Illustration of the process for obtaining the ligand features of active sites. The element properties of atoms (group, period, and atom radius) are arranged as a vector according to their spatial distance from the adsorbate.

active site types. To expand the dataset, an additional 3477 structures were pseudo-labeled using a GBRT model trained on PGH-derived features. The simulated data were generated by modifying the atomic species within the averaged atomic Cartesian coordinates derived from the DFT-optimized structures. This approach is justified by the observation that *OH adsorption perturbs surface metal atoms by less than 0.1 Å in the DFT dataset, allowing the use of averaged coordinates to represent new structures. Notably, although the newly generated structures share the same spatial geometry, the directionality of the connections is altered, enabling the PGH to distinguish between different samples within the same local environment based on their topological features. This expanded our labeled dataset to 4636 points. Although 75% of these labels were not

obtained via direct DFT calculations, they enhance VAE model convergence without sacrificing accuracy, as detailed below.

The training of PGH-VAEs

The PGH-VAEs are designed to predict and generate catalytic active sites, with the core innovations in their multi-channel architecture and high-precision adsorption energy predictions, which establish interpretability in the latent space and accuracy in inverse design.

In the training process, the model receives feature vectors representing coordination and ligand characteristics of HEA interfaces. These are separated into two distinct encoders (Fig. 3a), each generating Gaussian distributions for the coordination and ligand variables (mean and variance).

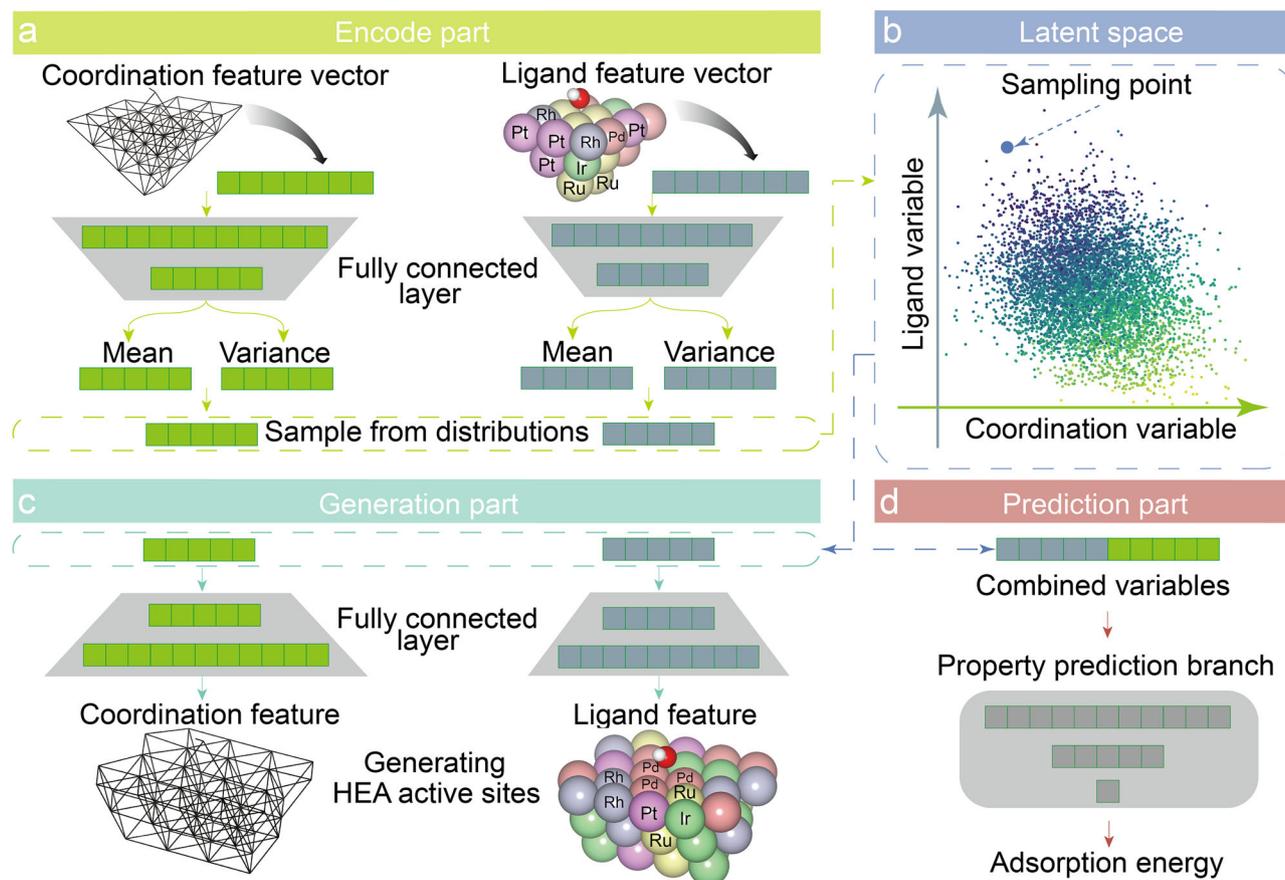


Fig. 3 | PGH-VAEs: encoding, generation, and prediction workflow. **a** Encode Part: Two encoders process coordination and ligand feature vectors from the HEA interface, outputting mean and variance parameters as Gaussian variables. Ligand and coordination variables (green dashed lines) are randomly sampled. **b** Latent Space: Ligand and coordination variables are independently clustered and visualized

on a 2D scatter plot, with data point colors representing adsorption energy.

c Generation Part: Coordination and ligand variables from new sampling points are decoded by two decoders into feature vectors for new HEA interfaces. **d** Prediction Part: The framework predicts *OH adsorption energy by combining sampled ligand and coordination variables in the property prediction branch.

Random sampling within these distributions provides coordination and ligand variables, which are then input into corresponding decoders (Fig. 3b). This process constructs a structured latent space where each point encodes specific coordination and ligand information (Fig. 3c). In the Generation section of Fig. 3d, the decoders reconstruct the original feature vectors, while the prediction branch combines coordination and ligand features to forecast adsorption energy. PGH-VAEs are trained by minimizing four losses: ligand feature reconstruction error (L_{ligand}), coordination feature reconstruction error ($L_{coordination}$), property (the adsorption energy) prediction error ($L_{property}$), and Kullback-Leibler divergence (KL divergence) (L_{KL}), yielding a total loss function:

$$L_{total} = L_{ligand} + L_{coordination} + L_{property} + L_{KL}. \quad (1)$$

Initially, we attempted to use only the 1159 DFT-calculated OH adsorption energies as the model dataset, with a training-to-test set ratio of 7:3. Both the MAEs and RMSEs of the property predictions show a sharp initial decrease (Fig. S4), followed by a gradual decline. In testing, however, errors initially decrease but then rise, indicating overfitting—a common issue in computational materials science due to limited data. To address this, we incorporated the expanded dataset of 4636 points generated through semi-supervised learning. In training, only the semi-supervised data were added to the training set, while the test set remained as 30% of the original DFT-calculated data, ensuring robust and reliable test accuracy. PGH-VAEs were retrained with the updated dataset; architecture and parameters are provided in Tables S3 and S4. We observed that during the training of PGH-

VAEs, L_{KL} typically has a larger magnitude (~ 10) compared to L_{ligand} (~ 1), $L_{coordination}$ (~ 1), and $L_{property}$ (~ 0.1). This behavior is expected and arises from differences in task complexity and output dimensionality. Specifically, L_{KL} is computed over latent distributions across the batch and all latent dimensions, making it inherently larger. In contrast, the prediction branch optimizes a single scalar property, which is easier to fit and naturally yields a smaller loss. To ensure balanced training, we carefully weighted the ligand and coordination reconstruction losses according to their dimensionality^{54,55}. Figure S5 illustrates the MAE and RMSE trends over epochs in training and testing, showing a sharp initial drop followed by a gradual decline, indicating stable training and effective mitigation of the prior overfitting issue. The final results show that PGH-VAEs achieve an MAE of 0.045 eV for adsorption energy prediction, reducing the error by an impressive 50% compared to the coordination number-based neural network model for HEA catalysts (Table S5)⁴⁰. These results underscore the exceptional performance of PGH-VAEs and highlight the transformative impact of the semi-supervised learning approach in significantly enhancing model performance.

Demonstration of PGH-VAEs on active sites design and optimization

Imbuing deep generative models with interpretability facilitates the understanding of the physical principles underlying materials design. During the prediction and generation process, PGH-VAEs employ a two-dimensional latent space, where each data point represents an input instance. The coordinates of these points in the latent space are derived from the downscaling of coordination and ligand encoders via principal

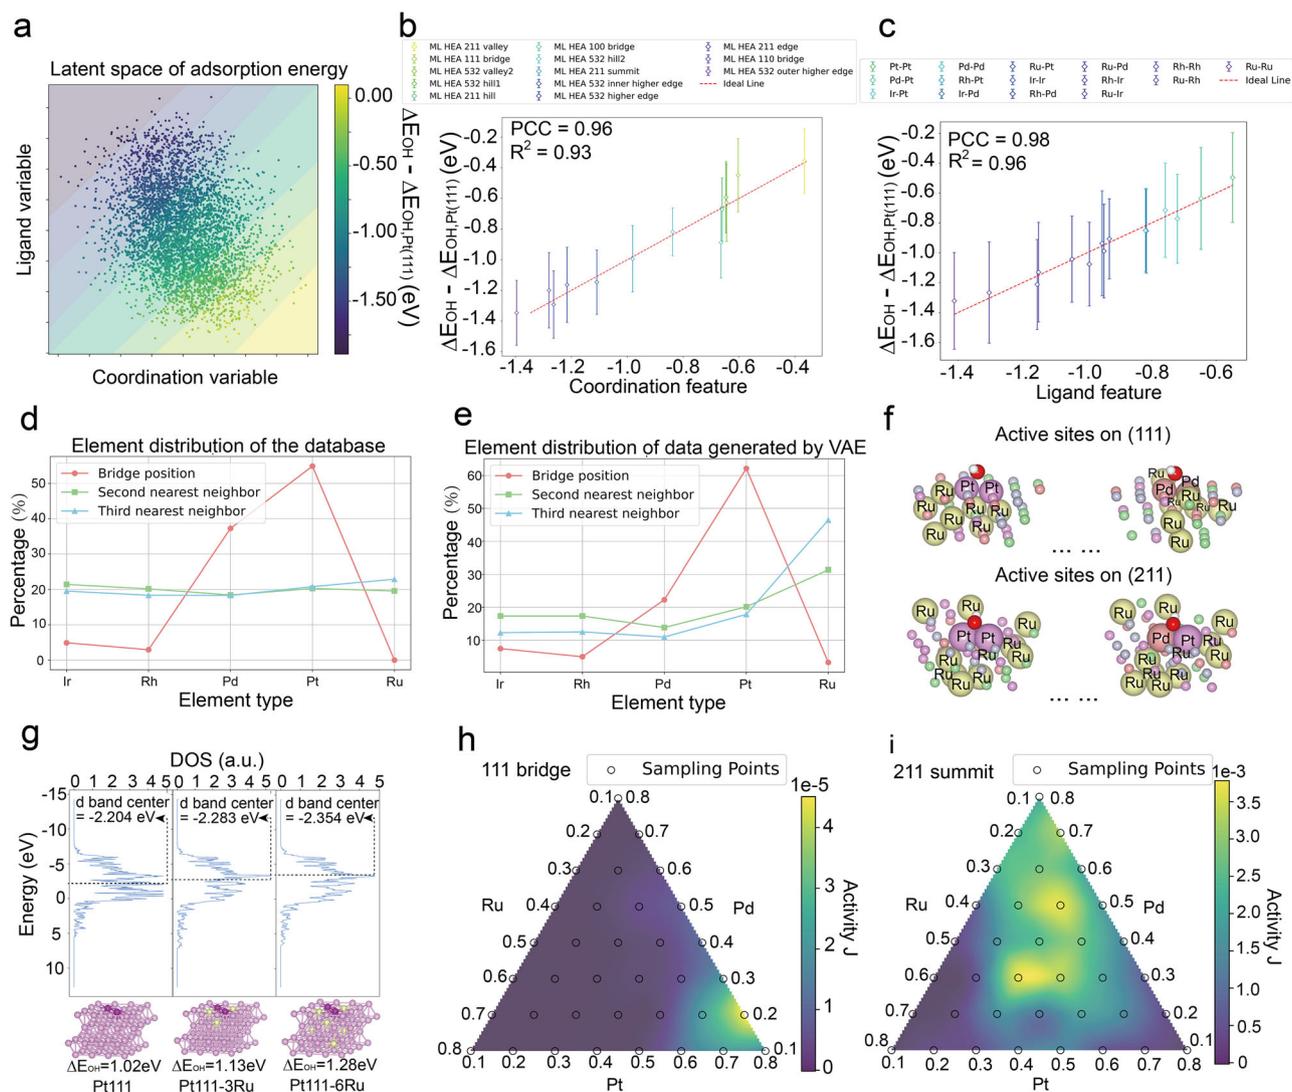


Fig. 4 | Demonstration of PGH-VAEs on active sites design and optimization.

a Latent space of adsorption energy plotted against coordination and ligand variables derived via PCA. Data points are color-coded based on their adsorption energy values. **b** Correlation between numerical representation and mean adsorption energy for 13 unique bridge-site structures. The numerical representation of structures is obtained by linear combination of their Betti numbers (β_0 , β_1 , and β_2). **c** Correlation between numerical representation and mean adsorption energy for 15 unique bridge-site elemental combinations. The numerical representation of elemental combinations is derived by linear combination of the element properties of bridge atoms. Standard deviations are shown for each point to visualize the variability in adsorption energy within specific structures or elemental combinations. **d** Percentage distribution of elemental composition for bridge-site atoms, first-nearest neighbors, and second-nearest

neighbors in the latent space associated with optimal *OH adsorption. **e** Percentage distribution of elemental composition for bridge-site atoms, first-nearest neighbors, and second-nearest neighbors in active sites generated via inverse design. **f** Illustration of active sites generated through inverse design, highlighting bridge-site atoms and distal Ru atoms. **g** Density of states (DOS) analysis of Ru-doped Pt (111) panels from left to right show the d-orbital states, d-band centers, and *OH adsorption energies for Pt adsorption sites without Ru doping, with three Ru atoms doped in the first-nearest neighbors, and with three additional Ru atoms doped in the second-nearest neighbors of the bridge site. **h, i** Heatmaps of the relationship between PtPdRu alloy composition ratios and catalytic activity on the (111) and (211) facets, respectively. Color codes: Ir: green, Pd: pink, Pt: purple, Rh: blue, Ru: yellow, O: red, H: white. In (g), the bridge atoms that adsorb *OH are highlighted by dark purple.

component analysis (PCA). Both ligand variables and coordination variables explain over 95% of the total variance, ensuring effective dimensionality reduction. The results (Fig. 4a) reveal a significant correlation between the relative adsorption energy of samples and the coordination and ligand variables in the latent space. Along the diagonal from the top left to the bottom right, the adsorption energy transitions from negative to positive, indicating that coordination and ligand effects can be utilized to explain variations in adsorption performance, with their influence exhibiting a coupled nature.

To further distinguish how the two effects regulate adsorption energy, we analyzed their individual correlations with changes in adsorption energy. For the ligand effect, leveraging the use of PGH to describe the active sites,

we can directly quantify structural features into specific numerical values by linearly combining the average values of β_0 , β_1 , and β_2 , detailed in Supporting Note 2. The 13 unique active sites show distinct numerical differences, even among bridge sites on the same Miller index surface, showcasing the strength of topological data analysis in capturing the three-dimensional structure sensitivity and reducing them to simple numerical representations (Tables S6 and S7). We mapped the numerical representations of these 13 active sites against their corresponding average adsorption energies, revealing a perfect linear correlation that underscores the strong influence of structure sensitivity on adsorption energy (Fig. 4b). The substantial standard deviations observed for each point reveal the considerable tuning space for adsorption energy by varying atomic species within identical structural

environments. Furthermore, the weights of the average Betti numbers reveal that β_0 (connected components) and β_2 (three-dimensional cavities) dominate in the structural representation, underscoring their pivotal roles in determining adsorption energy, whereas the influence of β_1 (ring structures) is comparatively minor. Chemically, both β_0 and β_2 reflect structural density from complementary perspectives: β_0 represents the number of independent atomic components, with higher values indicating increased local atomic density, while β_2 captures the presence of three-dimensional cavities. A greater number of such cavities implies a higher density of path combinations (chemical bonds) within the local structure, corresponding to a more compact atomic arrangement and a higher degree of bond saturation, thereby facilitating the *OH desorption. Consequently, the 211-valley bridge site and 111-bridge site (Fig. S1), which all exhibit the characteristic close-packed structure of FCC metals, exhibit more positive *OH adsorption energies, favoring ORR performance. The (532) surface displays substantial variations in *OH adsorption energy across its various bridge sites, elegantly captured through the numerical representations provided by PGH. A comparative analysis with the Generalized Coordination Number (GCN)^{56,57} descriptor shows that the GCN model achieved a correlation coefficient of 0.43 and an *R*-squared value of 0.18 (Fig. S8), substantially lower than the VAE-based model's values of 0.96 and 0.93. This highlights PGH's remarkable sensitivity in characterizing subtle microscopic structural features.

We applied a similar approach to analyze the ligand effect. The 15 unique atomic combinations at the bridge sites were quantified into specific numerical values based on their linear combination of element properties (Supporting Note 3, Tables S8 and S9). Plotting these numerical representations against the corresponding average adsorption energies also revealed a perfect linear correlation, emphasizing that the bridge-site atoms, being in direct contact with the adsorbate, play another pivotal role in determining adsorption energy (Fig. 4c). Despite identical bridge-site atoms, significant standard deviations in adsorption energy suggest that adjustments to the active site structure or distal atoms can further modulate *OH adsorption. Among the bridge-site combinations, Pt–Pt, Pd–Pt, and Pd–Pd exhibit the most positive adsorption energies, suggesting they are the most favorable active sites. This is intuitive, as both Pt and Pd in their metallic states exhibit excellent ORR performance. Based on these findings from the latent space, active sites on the (111) or (211) facets with Pd and Pt forming the bridge sites are identified as the most advantageous for achieving the optimal *OH adsorption state.

Discovering active sites with targeted properties is the core objective of inverse design. Therefore, we conducted ~240 structural samplings in the lower-right region of the latent space, where the optimal *OH adsorption state is observed. The corresponding coordination and ligand environments were decoded to identify the associated active sites. All these active sites exhibited characteristics of the (111) and (211) facets (Fig. S6), in accordance with the best-performing samples in the latent space and validating the reliability of our generative approach. Furthermore, we analyzed the elemental composition of these structures by averaging across all samples (Fig. 4e). Data points in the latent space with adsorption energies greater than that on Pt (111) were also plotted for comparison (Fig. 4d). Both the inversely generated data and the latent space data consistently exhibited a preference for Pt or Pd as the bridge-site atoms. Interestingly, in the inversely generated data, the proportion of Ru atoms in the second and third nearest neighbors is notably high compared to other elements (Fig. 4e, f). This suggests that incorporating Ru within the second and third nearest neighbors of the bridge site favors achieving the optimal *OH adsorption state. Such effects may stem from Ru's high electron affinity than Ir and Rh, which can strongly influence the electronic structure of neighboring Pt and Pd atoms, thereby modulating the *OH adsorption state. To investigate this, we constructed simplified models of Pt (111) surfaces doped with Ru at various positions. The results (Fig. 4g) reveal that as the Ru content increases in the first and second-nearest neighbors of the bridge site, the d-band center of the Pt–Pt site directly interacting with *OH shifts significantly downward, facilitating *OH desorption at this site. Incorporating Ir or Rh while

maintaining the same configuration results in significantly smaller effects on the d-band center and adsorption energy (Fig. S7). These findings demonstrate that distal atoms, even without direct adsorbate interaction, can modulate adsorption energies via local coordination and electronic effects. This may be attributed to the contribution of void structures (β_2) (Fig. 4b and Table S7), as Ru in the second-nearest neighbors affects the bridge site via remote spatial interaction. This insight not only reinforces the robustness of our generative model but also provides a mechanistic basis for tailoring active sites in HEAs.

Based on the results of inverse design, an optimization strategy for IrPdPtRhRu HEAs can be proposed. On one hand, the surface should predominantly expose Pt and Pd sites as the direct *OH adsorption centers; on the other hand, the incorporation of Ru in the neighbors can enhance the activity than the Ir and Rh elements. To further quantify this strategy and provide quantitative guidance for experimental synthesis, we systematically explored the compositional ratios of PtPdRu ternary alloys. Alloy compositions were varied from 0.1 to 0.9, yielding 36 distinct concentration gradients. For each gradient, 1000 potential active sites were enumerated, and their *OH adsorption energies were predicted using the VAE's prediction branch. The relative activity (relative to the catalytic performance of the pure Pt(111) surface) of each site was calculated via the Arrhenius equation, and the cumulative activity for each specific composition was determined (details in methods). Our findings (Fig. 4h, i) reveal that the highest activity is achieved with a Pt: Pd: Ru ratio of 0.7:0.2:0.1 on the (111) facet and 3:3:4 on the (211) facet, with activities 6646 times and 11 times greater than those of the PtPdIrRuRh (111) and (211) facets, respectively. Experimentally, such tailored alloy compositions can be synthesized through advanced techniques, including vapor phase spark discharge, sputtering, and acute chemical reduction, among other approaches³⁸.

Discussion

GLMY topology leverages the relationships between vertices (atoms) and paths (chemical bonds) to represent structures, making it naturally suited for the description of atomic configurations. It retains critical three-dimensional spatial information across multiple scales, overcoming the dimensionality reduction issues often associated with graph and cheminformatics-based representations. Traditional graph-based descriptors, built on vertices and edges, are limited by their inability to capture directionality and high-dimensional interactions, making them insufficient to distinguish topologically distinct environments with similar coordination numbers. Furthermore, it translates structural features into quantitative values by extracting topological invariants, such as Betti numbers. These invariants—capturing attributes like connectivity, loops, and voids—provide a clear numerical representation of different structures, effectively quantifying structural differences and laying a foundation for the interpretability of deep generative models. The GLMY topology is inherently adaptable to systems with dissimilar atomic radii, as both vertex weighting and interatomic distances are naturally incorporated into the path complex construction. While the atomic radii of elements studied in this work are relatively similar, the proposed method remains applicable to more diverse systems, where variations in atomic size could play a more pronounced role in shaping topological features. Additionally, when combined with a multi-channel encoding structure, the VAE model organizes the latent space into a two-dimensional landscape of structural and chemical information. This arrangement visualizes regions with optimal *OH adsorption properties, guiding further exploration of catalytic active sites through targeted sampling in the latent space. Unlike traditional forward-learning methods that require extensive screening over predefined latent spaces, this inverse design approach selectively focuses on areas of interest, strengthening the insights derived from existing data and revealing new features that can drive catalyst optimization. We employed IrPdPtRhRu HEAs as a case example to demonstrate the framework's capability. HEAs capture the multi-element interactions, diverse local structures, and dynamic coordination environments commonly seen in real-world catalysts, making them an ideal system for testing the framework's generalizability and multi-scale structural

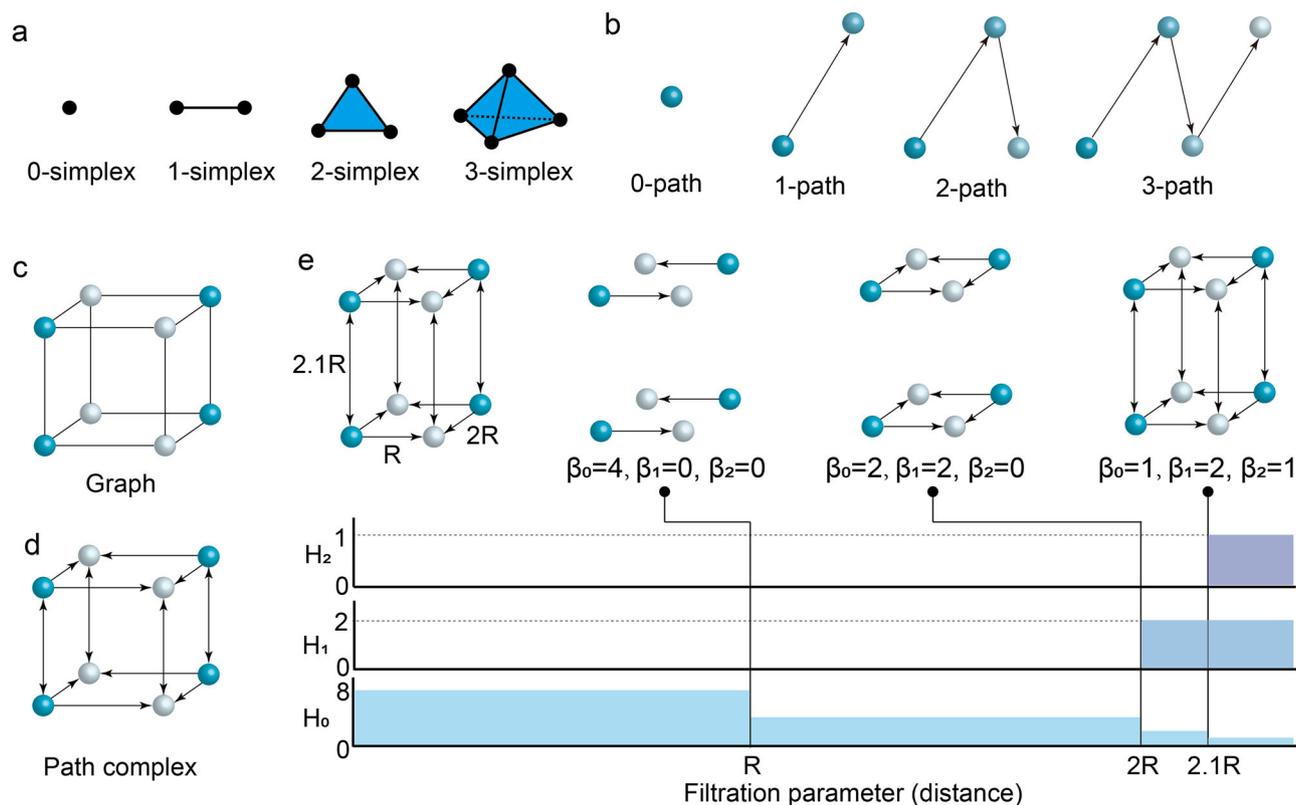


Fig. 5 | Building blocks and topological analysis of a cubic system using simplicial and path complexes. **a** Basic building blocks of a simplicial complex. **b** Basic building blocks of a path complex. **c** Graph representation of a cubic system with

distinct vertices. **d** Path complex representation of a cubic system with distinct vertices. **e** Zero-, one-, and two-dimensional Betti numbers of the cubic system with distinct vertices under different connection schemes.

characterization abilities. Notably, the PGH-VAEs are not limited to this specific elemental combination. Leveraging the multi-scale insights of GLYM topology and the flexibility of its multi-channel VAE, this framework is highly adaptable, with broad potential for application across catalytic systems.

The PGH framework demonstrates significant data and computational efficiency advantages over traditional machine learning potentials (MLPs). While MLP-based approaches often require training on thousands to tens of thousands of DFT-relaxed configurations—as exemplified by recent work on HEA catalysts using over 135,000 adsorption sites⁵⁸—our method achieves accurate structure-property characterization with only ~1000 descriptors and without iterative optimization steps. Notably, a compelling vision for future development of PGH-VAEs involves integrating pretrained MLPs that capture structure-energy relationships over a broader elemental landscape. Pretrained MLPs, such as DPA-1 and DeePMD-kit, are typically trained on millions of atomic configurations spanning dozens of elemental species. These models offer broad coverage of atomic interactions, although their predictive accuracy for specific systems often remains limited without further fine-tuning. When extending the VAE framework to a broader chemical space, it is essential to construct an additional representative and chemically diverse training set, including new elements. To streamline this process, pretrained MLPs can be leveraged to perform large-scale pre-optimization of newly sampled structures. From these, structurally informative configurations can be selectively refined using high-accuracy DFT calculations. This hybrid strategy not only preserves the precision of DFT where it matters most but also substantially reduces the computational burden associated with dataset generation, thereby enabling broader generalization of the model across chemical spaces without compromising predictive fidelity. The feasibility of this approach has been demonstrated by prior studies, where similar strategies have been applied to efficiently search for low-energy adsorption configurations of adsorbates on diverse surface types⁵⁹. The most challenging aspect of this workflow may lie in identifying

representative and meaningful structures from the pre-optimized configurations. Low-energy states, high-uncertainty points, and geometries that are difficult to converge are potential candidates of interest; however, a systematic and reliable selection strategy remains to be fully developed. Such advancement would enhance the model's design capabilities, paving the way for the rational creation and evaluation of complex, multi-element catalytic materials and ushering in a new paradigm in catalyst design.

Methods

Path complexes and their homology

PGH, an extension of GLMY homology theory, offers a method to quantify the persistence of homology within path complexes across various scales^{42–44}.

GLMY homology is a generalized homology theory developed to extend classical topological tools to settings where directionality and asymmetry play essential roles. In classical homology theory, the fundamental building blocks of the studied space are simplices—vertices, edges, triangles, and higher-dimensional analogs—sembled into simplicial complexes, as shown in Fig. 5a. The goal is to analyze the connectivity and higher-dimensional structures within such complexes, often built from undirected graphs or point cloud data.

In contrast, as shown in Fig. 5b, GLMY homology is defined on path complexes, where the basic units are directed paths rather than simplices. This framework allows the encoding of directional relationships between vertices, which are especially important in systems where vertex identity matters, such as in crystals or molecules with heterogeneous atomic types.

As shown in Fig. 5c, d, while homology captures the topological features of undirected structures, GLMY homology provides a richer representation by accounting for ordered, vertex-sensitive connections. As a result, it is particularly advantageous for studying crystalline and molecular systems, where the asymmetry and directional nature of local configurations carry essential chemical or physical meaning. By preserving more

information about connectivity and composition, GLMY homology offers a more expressive and comprehensive topological descriptor in such contexts.

In algebraic topology, Betti numbers are a sequence of integers that describe the number of independent k -dimensional holes in a topological space. Specifically, the k th Betti number β_k is the rank of the k th homology group, which counts the number of k -dimensional cycles that are not boundaries of $(k + 1)$ -dimensional objects. For example, β_0 counts connected components (0-dimensional holes), β_1 counts independent loops or cycles (1-dimensional holes), and β_2 counts voids enclosed by surfaces (2-dimensional holes). As shown in Fig. 5e, the Betti numbers in each dimension of a cube vary with changes in connectivity, capturing the topological features of the path complex. When there are two holes and two disconnected components, $\beta_0 = 2$ and $\beta_1 = 2$. When a cavity forms, $\beta_1 = 2$ and $\beta_2 = 1$.

Let V be a nonzero finite set. For a given integer $p \geq 0$, the elementary p -path on V is a sequence $i_0 i_1 i_2 \dots i_p$ of elements in V . Let $e_{i_0 i_1 i_2 \dots i_p}$ be the generator corresponding to the elementary p -paths, then a \mathbb{K} -linear space can be generated by all the elementary p -path, which is denoted as $\Lambda_p = \Lambda_p(V)$. Specifically, we make the convention that $\Lambda_{-1} = 0$. An element v in Λ_p can be uniquely written as

$$v = \sum_{i_0, i_1, i_2, \dots, i_p \in V} a^{i_0 i_1 i_2 \dots i_p} e_{i_0 i_1 i_2 \dots i_p}, \quad a^{i_0 i_1 i_2 \dots i_p} \in \mathbb{K}. \tag{2}$$

For any integer $p \geq 0$, a \mathbb{K} -linear map $\partial_p: \Lambda_p \rightarrow \Lambda_{p-1}$ is defined on the generator $e_{i_0 i_1 i_2 \dots i_p}$ as

$$\partial_p e_{i_0 i_1 \dots i_p} = \sum_{k=0}^p (-1)^k e_{i_0 \dots \hat{i}_k \dots i_p}, \quad p > 0, \tag{3}$$

and $\partial_0 e_{i_0} = 0$ for $p = 0$, where \hat{i}_k indicates omission of the index i_k , and $\partial^p \circ \partial^{p+1} = 0$. Thus $\partial = (\partial_p)_p$ can be deduced as a boundary operator on $(\Lambda_p)_p$.

A path complex on V is a nonempty collection \mathcal{P} of elementary paths on V , and it satisfies that $i_0 i_1 i_2 \dots i_p \in \mathcal{P}$ implies $i_1 i_2 \dots i_{p-1} i_p, i_0 i_1 i_2 \dots i_{p-1} \in \mathcal{P}$.

A digraph $G = (V, E)$ consists of a set V of vertices and a subset $E \subset \{V_G \times V_G\}$ of ordered pairs (v, w) of vertices called arrows. The arrow (v, w) is denoted $v \rightarrow w$. The collection $\{i_0 i_1 i_2 \dots i_p | i_k \rightarrow i_{k+1} \text{ for all } 0 \leq k \leq p-1, p \geq 0\}$ of paths on G is a path complex, denoted by $\mathcal{P}(G)$. The p -paths in \mathcal{P} are called allowed p -paths, and the \mathbb{K} -linear space spanned by the allowed p -paths is denoted as

$$\mathcal{A}_p = \mathcal{A}_p(\mathcal{P}) = \left\{ \sum_{i_0, i_1, \dots, i_p \in V} a^{i_0 i_1 \dots i_p} e_{i_0 i_1 \dots i_p} \mid i_0 i_1 \dots i_p \in \mathcal{P}, a^{i_0 i_1 \dots i_p} \in \mathbb{K} \right\}. \tag{4}$$

Here, as convention, let $\mathcal{A}_{-1} = 0$ be the null space. The space of ∂ -invariant p -paths can be deduced by

$$\Omega_{-1} = 0, \quad \Omega_p = \Omega_p(\mathcal{P}) = \{x \in \mathcal{A}_p \mid \partial x \in \mathcal{A}_{p-1}\}, \quad p \geq 0. \tag{5}$$

Then $\partial_p|_{\Omega_p}: \Omega_p \rightarrow \Omega_{p-1}$ satisfies $\partial_p|_{\Omega_p} \partial_{p+1}|_{\Omega_{p+1}} = 0$ and $(\Omega_p)_p$ with the boundary operator $\partial| = (\partial_p|_{\Omega_p})_p$ is a sub-chain complex of $(\Lambda_p(V))_p$. The GLMY homology of a path complex \mathcal{P} is defined by

$$H_p(\mathcal{P}; \mathbb{K}) := \frac{\ker \partial_p|_{\Omega_p}}{\text{im} \partial_{p+1}|_{\Omega_{p+1}}}, \quad p \geq 0. \tag{6}$$

The GLMY homology of a digraph G is that of the path complex $\mathcal{P}(G)$. The p th Betti number of the digraph G is the rank of the homology $H_p(G; \mathbb{K}) = H_p(\mathcal{P}(G); \mathbb{K})$, denoted as $\beta_p(G)$.

Let (S, \leq) be an order set and (S, \leq) can be regarded as a category with elements in S as objects and all the binary orders as morphisms. A filtration of path complexes means a covariant functor $\mathcal{F}: (S, \leq) \rightarrow \mathbf{Path}$ from the category (S, \leq) to the category of path complexes. For each element $a \in S$, \mathcal{F}_a is a path complex. Let $f_{a,b}: \mathcal{F}_a \rightarrow \mathcal{F}_b$ be the morphism induced by $a \rightarrow b$, then $f_{b,c} \circ f_{a,b} = f_{a,c}$ for $a \leq b \leq c$. The morphism $f_{a,b}$ induces a morphism of GLMY homology

$$\tilde{f}_{a,b}: H_p(\mathcal{F}_a; \mathbb{K}) \rightarrow H_p(\mathcal{F}_b; \mathbb{K}).$$

The p th (a, b) -persistent GLMY homology of \mathcal{F} is defined by

$$H_p^{ab}(\mathcal{F}; \mathbb{K}) = \text{im}(H_p(\mathcal{F}_a; \mathbb{K}) \rightarrow H_p(\mathcal{F}_b; \mathbb{K})), \quad p \geq 0. \tag{7}$$

The (a, b) -persistent Betti number is defined as the rank of $H_p^{ab}(\mathcal{F}; \mathbb{K})$.

In practice, the path complex is usually defined on digraphs. Let $\mathbf{Digraph}$ be the category of digraphs and digraph maps. A filtration of digraphs is a covariant functor $\mathcal{D}: (S, \leq) \rightarrow \mathbf{Digraph}$ from the category (S, \leq) to the category $\mathbf{Digraph}$. A filtration of digraphs can induce a filtration of path complexes, which results in the PGH of digraphs. Different filtration can result in different persistence.

Let $G = (V, E)$ be a digraph and V represents the set of data points in a metric space $(X, \|\cdot\|)$. Then, there is a weight function $d: E \rightarrow \mathbb{R}$ on the edge set E deduced by

$$d(x_1, x_2) = \|x_1 - x_2\|, \quad (x_1, x_2) \in E \subseteq X \times X. \tag{8}$$

Here, this work specifies the metric space $(X, \|\cdot\|)$ as the Euclidean space with L_2 -norm. Then, let $E_t = \{(x, y) \in E \mid d(x, y) \leq t\}$ and $\mathcal{G}_t = (V, E_t)$. It can be deduced that $\mathcal{G}: (\mathbb{R}, \leq) \rightarrow \mathbf{Digraph}$, $t \rightarrow \mathcal{G}_t$ is a filtration of digraphs, which leads to a persistent diagram $\mathcal{D}(G)$ of G .

Variational auto-encoders

In PGH-VAEs training, four key losses shape model performance. The reconstruction losses for ligand and coordination features quantify how accurately the decoder recreates these properties from the latent space, ensuring the network captures essential structural information. Minimizing these losses through encoder-decoder optimization improves feature representation and generalization. The property prediction loss, reflecting discrepancies between predicted and actual properties (adsorption energies), is reduced via supervised learning to enhance accuracy. KL divergence ensures the latent space aligns with a smooth prior distribution, enabling meaningful interpolation and sampling. Balancing the KL divergence term weight maintains a trade-off between reconstruction fidelity and latent space regularization.

In the sampling process, PGH-VAEs construct a low-dimensional latent space, where each point represents an input data instance. The coordinates of these points are derived from the downscaled ligand and coordination variables, respectively. Within this latent space, our framework identifies regions that may contain potential active sites by analyzing the characteristics of the existing data points. A sample point is then selected from these regions. The decoders reconstruct the ligand and coordination variables into ligand and coordination feature vectors, respectively, generating potential active sites. These feature vectors are combined and input into the property prediction branch to generate the predicted adsorption energy. This process facilitates the design of active sites with optimal adsorption capacity.

DFT calculation data

This work uses the same parameter setting and data as ref. 40. Spin-polarized plane-wave DFT⁶⁰ calculations were performed for HEAs structure optimization using VASP with projector augmented wavefunctions and the PBE exchange functional. A kinetic energy cutoff of 500 eV and a Fermi smearing width of 0.1 eV were used to ensure convergence. Van der Waals interactions were tested on a subset of data using Grimme's DFT-D3

method, revealing a negligible effect. Supercells contained at least 64 atoms, with a Monkhorst-Pack (2,2,1) *k*-point grid. Geometry optimizations were terminated when atomic forces fell below 0.02 eV/Å.

Surfaces with various Miller indices were created using ASE and PyMatGen, with lateral dimensions of at least 10 Å to minimize periodic image interactions. Lattice constants were averaged according to Vegard's law. Surfaces included at least four atomic layers, with the bottom two fixed during relaxation, and a vacuum of 10 Å was added in the vertical direction.

The adsorption energy of *OH is calculated by comparing with Pt(111), which is the most referenced material.

$$\Delta E_{OH} - \Delta E_{OH,Pt(111)} = (E_{*OH} - E_*) - (E_{*OH,Pt(111)} - E_{*,Pt(111)}). \quad (9)$$

where E_{*OH} , E_* , $E_{*OH,Pt(111)}$, $E_{*,Pt(111)}$ are the electronic energies of *OH-adsorbed interface, clean interface, *OH-adsorbed interface of Pt(111), and clean interface of Pt(111), respectively.

The activity calculation of PtPdRu

For a type of catalytic active site, where the proportion of Pt is r_{Pt} , the proportion of Pd is r_{Pd} and the proportion of Ru is r_{Ru} , with a total of N such catalytic active sites, the activity j can be determined using the following expression based on the format of the Arrhenius equation^{37,61}:

$$j = \frac{1}{N} \sum_{i=1}^N (r_{Pt} \times r_{Pd} \times r_{Ru}) e^{-\frac{|\Delta E_{OH}^i - \Delta E_{target}|}{k_b T}}. \quad (10)$$

Here, ΔE_{OH}^i represents the modeled adsorption energy for the i th active site. E_{target} is the optimal *OH adsorption energy according to the Sabatier principle, which is 0.1 eV higher than that in Pt (111)^{37,53}. Based on this setting, the activity of Pt-Pd-Ir-Ru-Rh (111) is assigned a j value of 7.0×10^{-5} , and the activity of Pt-Pd-Ir-Ru-Rh (211) is assigned a j value of 3.3×10^{-4} . k_b is Boltzmann's constant, and T is the absolute temperature. Here, we assume that $T = 300$ K.

Data availability

The data is available from <https://zenodo.org/records/15266141>.

Code availability

Code is available from <https://github.com/PKUsum2023/PGH-VAEs>.

Received: 30 January 2025; Accepted: 8 May 2025;

Published online: 24 May 2025

References

- Seh, Z. W. et al. Combining theory and experiment in electrocatalysis: insights into materials design. *Science* **355**, eaad4998 (2017).
- Zheng, R., Liu, Z., Wang, Y., Xie, Z. & He, M. The future of green energy and chemicals: rational design of catalysis routes. *Joule* **6**, 1148–1159 (2022).
- Vogt, C. & Weckhuysen, B. M. The concept of active site in heterogeneous catalysis. *Nat. Rev. Chem.* **6**, 89–111 (2022).
- Qin, R., Liu, K., Wu, Q. & Zheng, N. Surface coordination chemistry of atomically dispersed metal catalysts. *Chem. Rev.* **120**, 11810–11899 (2020).
- Liu, L. & Corma, A. Metal catalysts for heterogeneous catalysis: from single atoms to nanoclusters and nanoparticles. *Chem. Rev.* **118**, 4981–5079 (2018).
- Andersen, M. & Reuter, K. Adsorption enthalpies for catalysis modeling through machine-learned descriptors. *Acc. Chem. Res.* **54**, 2741–2749 (2021).
- Wang, Q. & Yao, Y. Harnessing machine learning for high-entropy alloy catalysis: a focus on adsorption energy prediction. *npj Comput. Mater.* **11**, 91 (2025).
- Tran, K. & Ulissi, Z. W. Active learning across intermetallics to guide discovery of electrocatalysts for CO₂ reduction and H₂ evolution. *Nat. Catal.* **1**, 696–703 (2018).
- O'Connor, N. J., Jonayat, A., Janik, M. J. & Senftle, T. P. Interaction trends between single metal atoms and oxide supports identified with density functional theory and statistical learning. *Nat. Catal.* **1**, 531–539 (2018).
- Xu, H., Cheng, D., Cao, D. & Zeng, X. C. Revisiting the universal principle for the rational design of single-atom electrocatalysts. *Nat. Catal.* **7**, 207–218 (2024).
- Abild-Pedersen, F. et al. Scaling properties of adsorption energies for hydrogen-containing molecules on transition-metal surfaces. *Phys. Rev. Lett.* **99**, 016105 (2007).
- Liu, X., Jiao, Y., Zheng, Y., Jaroniec, M. & Qiao, S.-Z. Building up a picture of the electrocatalytic nitrogen reduction activity of transition metal single-atom catalysts. *J. Am. Chem. Soc.* **141**, 9664–9672 (2019).
- Li, H. et al. Machine learning big data set analysis reveals C–C electrocoupling mechanism. *J. Am. Chem. Soc.* **146**, 22850–22858 (2024).
- Zheng, S., Zuo, C., Liang, X., Li, S. & Pan, F. Valence state of transition metal center as an activity descriptor for CO₂ reduction on single atom catalysts. *J. Energy Chem.* **56**, 444–448 (2021).
- Freeze, J. G., Kelly, H. R. & Batista, V. S. Search for catalysts by inverse design: artificial intelligence, mountain climbers, and alchemists. *Chem. Rev.* **119**, 6595–6612 (2019).
- Yao, Z. et al. Inverse design of nanoporous crystalline reticular materials with deep generative models. *Nat. Mach. Intell.* **3**, 76–86 (2021).
- Kench, S. & Cooper, S. J. Generating three-dimensional structures from a two-dimensional slice with generative adversarial network-based dimensionality expansion. *Nat. Mach. Intell.* **3**, 299–305 (2021).
- Kim, B., Lee, S. & Kim, J. Inverse design of porous materials using artificial neural networks. *Sci. Adv.* **6**, eaax9324 (2020).
- Xia, W. et al. Accelerating the discovery of novel magnetic materials using machine learning-guided adaptive feedback. *Proc. Natl. Acad. Sci. USA* **119**, e2204485119 (2022).
- Noh, J. et al. Inverse design of solid-state materials via a continuous representation. *Matter* **1**, 1370–1384 (2019).
- Song, Z. et al. Inverse design of promising electrocatalysts for CO₂ reduction via generative models and bird swarm algorithm. *Nat. Commun.* **16**, 1053 (2025).
- Liu, P., Qin, R., Fu, G. & Zheng, N. Surface coordination chemistry of metal nanomaterials. *J. Am. Chem. Soc.* **139**, 2122–2131 (2017).
- Bligaard, T. & Nørskov, J. K. Ligand effects in heterogeneous catalysis and electrochemistry. *Electrochim. Acta* **52**, 5512–5516 (2007).
- Xu, W., Reuter, K. & Andersen, M. Predicting binding motifs of complex adsorbates using machine learning with a physics-inspired graph representation. *Nat. Comput. Sci.* **2**, 443–450 (2022).
- Pablo-García, S. et al. Fast evaluation of the adsorption energy of organic molecules on metals via graph neural networks. *Nat. Comput. Sci.* **3**, 433–442 (2023).
- Lin, X. et al. Machine learning-assisted dual-atom sites design with interpretable descriptors unifying electrocatalytic reactions. *Nat. Commun.* **15**, 8169 (2024).
- Zheng, S. et al. Active phase discovery in heterogeneous catalysis via topology-guided sampling and machine learning. *Nat. Commun.* **16**, 1–13 (2025).
- Esterhuizen, J. A., Goldsmith, B. R. & Linic, S. Interpretable machine learning for knowledge generation in heterogeneous catalysis. *Nat. Catal.* **5**, 175–184 (2022).
- Bunting, R. J., Wodaczek, F., Torabi, T. & Cheng, B. Reactivity of single-atom alloy nanoparticles: modeling the dehydrogenation of propane. *J. Am. Chem. Soc.* **145**, 14894–14902 (2023).
- Jiao, Y., Li, H., Jiao, Y. & Qiao, S.-Z. Activity and selectivity roadmap for C–N electro-coupling on MXenes. *J. Am. Chem. Soc.* **145**, 15572–15580 (2023).
- Lunger, J. R. et al. Towards atom-level understanding of metal oxide catalysts for the oxygen evolution reaction with machine learning. *npj Comput. Mater.* **10**, 80 (2024).

32. Saini, S., Halldin Stenlid, J. & Abild-Pedersen, F. Electronic structure factors and the importance of adsorbate effects in chemisorption on surface alloys. *npj Comput. Mater.* **8**, 163 (2022).
33. Liu, H.-S., Zhang, X.-M., Liang, G.-H., Zheng, S. & Li, J.-F. Investigation of water structure and proton transfer within confined graphene by ab initio molecule dynamics and multiscale data analysis. *Chin. J. Struct. Chem.* 100596 (2025).
34. Liang, G.-H., Liu, H.-S., Zhang, X.-M., Li, J.-F. & Zheng, S. Topology-based machine learning for predicting curvature effects in metal-nitrogen-carbon single-atom catalysts. *J. Energy Chem.* **105**, 608–616 (2025).
35. Anstine, D. M. & Isayev, O. Generative models as an emerging paradigm in the chemical sciences. *J. Am. Chem. Soc.* **145**, 8736–8750 (2023).
36. Fuhr, A. S. & Sumpter, B. G. Deep generative models for materials discovery and machine learning-accelerated innovation. *Front. Mater.* **9**, 865270 (2022).
37. Batchelor, T. A. et al. High-entropy alloys as a discovery platform for electrocatalysis. *Joule* **3**, 834–845 (2019).
38. Yao, Y. et al. High-entropy nanoparticles: synthesis-structure-property relationships and data-driven discovery. *Science* **376**, eabn3103 (2022).
39. Löffler, T., Ludwig, A., Rossmeis, J. & Schuhmann, W. What makes high-entropy alloys exceptional electrocatalysts? *Angew. Chem. Int. Ed.* **60**, 26894–26903 (2021).
40. Lu, Z., Chen, Z. W. & Singh, C. V. Neural network-assisted development of high-entropy alloy catalysts: decoupling ligand and coordination effects. *Matter* **3**, 1318–1333 (2020).
41. Pedersen, J. K. et al. Bayesian optimization of high-entropy alloy compositions for electrocatalytic oxygen reduction. *Angew. Chem.* **133**, 24346–24354 (2021).
42. Grigor'yan, A. A., Lin, Y., Muranov, Y. V. & Yau, S.-T. Path complexes and their homologies. *J. Math. Sci.* **248**, 564–599 (2020).
43. Chen, D. et al. Path topology in molecular and materials sciences. *J. Phys. Chem. Lett.* **14**, 954–964 (2023).
44. Chowdhury, S. & Mémoli, F. Persistent path homology of directed networks. In *Proc. Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms* 1152–1169 (SIAM, 2018).
45. Duan, Z. & Wang, G. Comparison of reaction energetics for oxygen reduction reactions on Pt (100), Pt (111), Pt/Ni (100), and Pt/Ni (111) surfaces: a first-principles study. *J. Phys. Chem. C* **117**, 6284–6292 (2013).
46. Yamada, Y. et al. Oxygen reduction reaction activities for Pt-enriched Co/Pt (111), Co/Pt (100), and Co/Pt (110) model catalyst surfaces prepared by molecular beam epitaxy. *Surf. Sci.* **607**, 54–60 (2013).
47. Rizo, R., Herrero, E. & Feliu, J. M. Oxygen reduction reaction on stepped platinum surfaces in alkaline media. *Phys. Chem. Chem. Phys.* **15**, 15416–15425 (2013).
48. Minamihara, H. et al. Continuous-flow reactor synthesis for homogeneous 1 nm-sized extremely small high-entropy alloy nanoparticles. *J. Am. Chem. Soc.* **144**, 11525–11529 (2022).
49. Li, Y. et al. Nanoscale design for high entropy alloy electrocatalysts. *Small* **20**, 2310006 (2024).
50. Chaves, A. S., Piotrowski, M. J., Guedes-Sobrinho, D. & Da Silva, J. L. Theoretical investigation of the adsorption properties of CO, NO, and OH on monometallic and bimetallic 13-atom clusters: the example of Cu₁₃, Pt₇Cu₆, and Pt₁₃. *J. Phys. Chem. A* **119**, 11565–11573 (2015).
51. Pedersen, J. K., Batchelor, T. A., Bagger, A. & Rossmeis, J. High-entropy alloys as catalysts for the CO₂ and CO reduction reactions. *ACS Catal.* **10**, 2169–2176 (2020).
52. Clausen, C. M., Batchelor, T. A., Pedersen, J. K. & Rossmeis, J. What atomic positions determines reactivity of a surface? long-range, directional ligand effects in metallic alloys. *Adv. Sci.* **8**, 2003357 (2021).
53. Nørskov, J. K. et al. Origin of the overpotential for oxygen reduction at a fuel-cell cathode. *J. Phys. Chem. B* **108**, 17886–17892 (2004).
54. Asperti, A. & Trentin, M. Balancing reconstruction error and Kullback-Leibler divergence in variational autoencoders. *IEEE Access* **8**, 199440–199448 (2020).
55. Groenendijk, R., Karaoglu, S., Gevers, T. & Mensink, T. Multi-loss weighting with coefficient of variations. In *Proc. IEEE/CVF Winter Conference on Applications of Computer Vision, IEEE Computer Society*, 1469–1478 (2021).
56. Cheng, D. et al. Structure Sensitivity and Catalyst Restructuring for CO₂ Electro-reduction on Copper. *Nat. Commun.* **16**, 4064 (2025).
57. Zhang, S., Tang, Q., Zhu, B. & Gao, Y. Reaction-driven varieties of active sites on Cu (100) in electrochemical CO₂ reduction reaction. *ACS Catal.* **15**, 6497–6506 (2025).
58. Sun, C., Goel, R. & Kulkarni, A. R. Developing cheap but useful machine learning-based models for investigating high-entropy alloy catalysts. *Langmuir* **40**, 3691–3701 (2024).
59. Lan, J. et al. Adsorbml: a leap in efficiency for adsorption energy calculations using generalizable machine learning potentials. *npj Comput. Mater.* **9**, 172 (2023).
60. Grimme, S., Antony, J., Ehrlich, S. & Krieg, H. A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H-Pu. *J. Chem. Phys.* **132**, 154104 (2010).
61. Zhou, Q. et al. Analyzing the active site and predicting the overall activity of alloy catalysts. *J. Am. Chem. Soc.* **146**, 15167–15175 (2024).

Acknowledgements

This work was financially supported by the Guangdong Basic and Applied Basic Research Foundation (2020A1515110843), Young S&T Talent Training Program of Guangdong Provincial Association for S&T (SKXRC202211), National Natural Science Foundation of China (22402163, 22109003), the Major Science and Technology Infrastructure Project of Material Genome Big-science Facilities Platform supported by Municipal Development and Reform Commission of Shenzhen, Soft Science Research Project of Guangdong Province (No. 2017B030301013), Natural Science Foundation of Xiamen, China (3502Z202472001), High-level Scientific Research Foundation of Hebei Province and Fundamental Research Funds for the Central Universities (20720240054).

Author contributions

B.W. derived the models and analyzed the data; S.Z. drafted the manuscript and supervised the project; J.W. proposed methods; J.L. supervised the project; F.P. conceived of the idea and drafted the manuscript with inputs from the other authors.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41524-025-01649-8>.

Correspondence and requests for materials should be addressed to Shisheng Zheng, Jingyan Li or Feng Pan.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025