

Global optimization workflow with compact graph neural networks for comprehensive transition-state mapping

Difan Zhang

Oak Ridge National Laboratory

Roger Rousseau

Oak Ridge National Laboratory

Vassiliki-Alexandra Glezakou

`glezakouva@ornl.gov`

Oak Ridge National Laboratory

Article

Keywords: Transition state prediction, Reaction pathways, Graph neural networks, Global optimization, Heterogeneous catalysis

Posted Date: October 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11175345/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Locating the transition states in them remains a central computational bottleneck, particularly in heterogeneous catalysis where reference data are scarce and the configurations that matter are not known in advance. Most machine-learning approaches to transition-state prediction pursue increasingly expressive architectures trained on abundant gas-phase data, a strategy poorly matched to data-limited surface chemistry. We present an integrated workflow that couples our global optimization code with a deliberately compact graph neural network with essential physical invariances encoded into the input representation rather than learned by the network architecture. The global search maps not only ground states but also the low-lying metastable configurations that conventional searches discard. A lightweight network, using 2–3 orders of magnitude fewer trainable parameters than competing models, predicts transition states for matching configuration pairs. Our results indicate that, in data-scarce scientific domains, matching model capacity to the available data offers a robust and transferable alternative to ever-larger models.

Introduction

Across chemistry, materials science, and biology, the rates of the processes that matter most, such as chemical bond making and breaking, diffusion, nucleation, conformational change, phase transformation, are governed not by the stable states a system spends most of its time in, but by the fleeting, high-energy configurations it must pass through to get from one stable state to another.^{1–5} The highest-energy configuration along the minimum energy pathway between two stable states is the transition state (**TS**), a first-order saddle point on the potential energy surface whose relative energy sets the kinetic barrier and therefore the timescale of the transformation. Despite its fundamental importance, the reliable identification of TS remains one of the foundational computational problems in diverse contexts, including chemical reactions^{6–8} and phase transformation⁹. TSs are extremely difficult to observe directly by experiments due to their unstable and short-lived nature, and their characterization has rested predominantly on computational approaches for decades.

Traditionally, computational TS searches fall into two families that differ in the information they require. The single-ended methods¹⁰, e.g., eigenvector-following method and single-ended growing string method, start from a single stable configuration or an approximate TS guess, and walk uphill toward the saddle point. The double-ended methods^{11, 12}, e.g., nudged elastic band (**NEB**) and double-ended string methods, exploit the interpolation between the two endpoint configurations as input. Despite their widespread use, both methods are often limited by high computational cost, limited scalability, and persistent convergence difficulties, particularly when the initial TS guess is far from the true saddle point or when the underlying energy landscape is rugged.^{13, 14}

The rapid advance of deep learning has motivated a new generation of data-driven approaches that directly predict TS configurations in chemical transformations, often in conjunction with double-ended methods^{6, 15, 16}. For example, Pattanaik et al. show a fully connected graph neural network (**GNN**) for TS

prediction of isomerization reactions.¹⁷ Jackson et al. developed TSNet, a Siamese message-passing neural network based on tensor field networks, for TS prediction in nucleophilic substitution reactions.¹⁸ Makos et al. proposed a set of generative adversarial networks (TS-GAN) for TS predictions of organic reactions.¹⁹ Choi employed pair-sequence interaction layers containing transformer encoders for TS prediction in gas-phase organic reactions.²⁰ Schreiner et al. introduced NeuralNEB, an equivariant GNN to perform NEB calculations for TS estimation.²¹ More recently, Duan et al. reported an equivariant diffusion model (OA-ReactDiff) for TS prediction,²² while Kim et al. developed TSDiff, a diffusion-based framework that generates TS configurations from two-dimensional molecular graphs.²³ These generative models expanded the accessible reaction space and explored diverse reaction pathways beyond single predefined conformations. In addition, Si et al. combined bitmap representation, convolutional neural networks, and a genetic algorithm to design a computational framework for organic TS prediction.²⁴ This trajectory of method development for TS prediction shows a steady trend toward increasingly expressive model architectures, such as the equivariant GNN, diffusion models, transformers, etc., often leading to improved predictive accuracy.

However, three fundamental questions that extend beyond chemistry and are pervasive across computation in physical sciences are also exposed. First, the benchmark successes described above rely largely on comparatively abundant datasets of gas-phase organic reactions.^{25, 26} In contrast, reference datasets containing TS for heterogeneous catalytic systems, where reactions occur at molecule/solid interfaces in most catalytic processes, remain extremely limited. The recently released OC20NEB²⁷ database contributes several hundred converged surface NEB trajectories for elementary processes such as desorption, dissociation, and mass transfer of organic species on inorganic surfaces. Yet, against the enormous combinatorial diversity of surface chemistry, topology, adsorbates and their binding geometries, the available data is far too sparse to parameterize high-capacity models for direct, single-shot TS prediction in practice.²⁸

This data scarcity creates a second closely related problem on the mismatch between model complexity and information scope of available data. Contemporary TS prediction models typically contain $> 10^6$ trainable parameters,^{17–23} whereas heterogeneous catalysis datasets typically comprise only $\sim 10^3$ examples²⁷. Although such highly expressive models can achieve impressive performance on in-distribution benchmarks, they can generalize poorly to chemically distinct systems. Improving predictive robustness with limited, high-dimensional scientific data is still a central challenge throughout computational science, and TS prediction in heterogeneous catalysis is a prime example of this difficulty.

A third question is rather methodological. Typically, TS searches almost always focus on pathways connecting two local energy minima, most-often determined by a gradient-based optimization. Under realistic, catalytic conditions, however, the kinetically dominant pathway does not necessarily proceed via ground-state structures.^{2, 3, 29, 30} Finite-temperature fluctuations can transiently populate metastable states that are slightly higher in energy yet connected by lower energy barriers. Therefore, restricting TS

search to ground-state configurations alone can potentially overlook important operative pathways and rate-determining steps, regardless of how accurately any single saddle point is computed.^{2, 31, 32}

In this work, we address these challenges through an integrated TS search framework that pairs global optimization with a deliberately lightweight GNN, providing an alternative to purely deep-learning-based approaches for TS prediction. The TS prediction in heterogeneous catalytic processes is focused on as an example as illustrated in Fig. 1a. The framework proceeds in two stages. First, reactant, product, and intermediate configurations are explored using RANGE,³³ our global optimization code that identifies not only ground-state structures, but also low-lying metastable configurations that are frequently neglected in conventional TS searches. This exploration yields a broader view of the accessible reaction landscape beyond pathways connecting only the lowest-energy states. Second, a compact edge-aware attention GNN, pretrained on the Transition1x dataset²⁵ and subsequently fine-tuned on OC20NEB²⁷ through transfer learning, predicts TS geometries for physically plausible reactant–product pairs. The predicted TS structures may be further refined using established TS optimization methods such as NEB.

Rather than increasing model complexity, our strategy emphasizes physically motivated structural representations and a model capacity commensurate with the available data. We demonstrate the framework on propane dehydrogenation over a platinum cluster supported by carbon nitride. We show that, despite employing approximately two orders of magnitude fewer trainable parameters than other deep-learning models, the combination of global configurational exploration and lightweight learning accurately reconstructs the relevant reaction landscape and identifies pathways that are inaccessible to conventional ground-state-focused searches. More broadly, our results suggest that, in data-scarce scientific domains, the integration of global optimization with compact machine-learning models can provide a practical route toward transferable and chemically meaningful TS prediction across complex energy landscapes.

Methods

Data and graph feature generation

To prepare appropriate training and validation data for the GNN, we utilized the Transition1x²⁵ and OC20NEB²⁷ databases. From the Transition1x dataset, we extracted the reactant, TS, and product structures for each reaction, resulting in 10073 three-frame trajectories. For the OC20NEB dataset, reactant, TS, and product structures were similarly extracted from the converged NEB trajectories. Although OC20NEB lists 932 converged NEB trajectories, about a third of pathways do not represent an activated process and only exhibit a monotonic variance of the energy along the reported reaction pathway. We therefore retained only trajectories featuring a distinct energy maximum, which was taken as the TS structure. This filtering yielded 699 three-frame trajectories from the OC20NEB database. To ensure consistency of atomic displacements, all trajectories were unwrapped using the reactant configuration as the reference frame, thereby eliminating artifacts arising from periodic boundary conditions.

To ensure seamless integration of gas-phase molecular systems from Transition1x and surface-based systems from OC20NEB, a unified feature representation applicable to both environments is required. Here, we construct a common representation that can be applied to gas-phase molecules as well as organic–inorganic surface interactions, as briefly shown in Fig. 1b. All three-frame trajectories from both databases are processed to generate the corresponding inputs and target outputs for GNN training. For a given three-frame trajectory, the feature construction proceeds as follows:

1. The displacement of each atom between the reactant and product configurations is computed. Atoms with displacement distances greater than 0.3 Å are recorded in a list (called active list). For systems with periodic boundary conditions, the minimum image convention is applied when computing displacements.
2. For each atom in the active list, neighboring atoms within a cutoff threshold (R_A , 5 Å by default, optimized from our following parameter screening) are identified and added to the active list if not already included. This procedure yields a subset of atoms that captures the structural changes associated with the reaction. The resulting subset defines a cluster model, independent of whether the original system is a gas-phase molecule or a surface slab and is used for subsequent feature generation.
3. The geometric center of the cluster is computed using both reactant and product configurations to define a common reference point. Atoms in the cluster are then reordered using a lexicographic scheme: first by increasing displacement magnitude, then by atomic number, and finally by distance to the common geometric center. This procedure yields a deterministic ordering of atoms, ensuring permutation invariance.
4. The displacement vector of the atom with the largest displacement is rotated to align with the positive X-axis. Subsequently, the cluster is rotated about the X-axis so that the displacement vector of the atom with the second-largest displacement lies in the XY-plane and points in the positive Y-direction. The resulting rotations are recorded in a rotation matrix that leads to rotational invariance. Then the entire cluster is translated so that its geometric center coincides with the origin. The corresponding translation vector is recorded, ensuring translational invariance.
5. A fixed dimension ($D_{\max}$, 36 by default, a number that was optimized based on tests presented below) is imposed on the active list. If the number of atoms in the cluster (i.e. the length of the active list) is smaller than $D_{\max}$, dummy atoms with zero atomic number and zero atomic positions are added as padding. If the number exceeds $D_{\max}$, atoms with the smallest displacements are removed until the final length equals $D_{\max}$.
6. In the end, for each reaction this procedure yields an ordered active atom list, a rotation matrix, and a translation vector, derived from the reactant and product configurations. During GNN training, the same transformations are applied to the corresponding TS structures, producing three-frame trajectories with cluster representations that are invariant to translation, rotation, and atomic permutation.
7. The resulting cluster representations are converted into graph structures for GNN input. Each atom is represented as a node, with node features including information of atomic type and position

changes between the reactant and product configurations, including atomic number, the interpolated midpoint position between the reactant and product configurations, and the atom's displacement vector. Edges are defined between atom pairs whose distances are less than a cutoff threshold (R_G , 5 Å by default, optimized by parameter screening as shown below). The edge feature contains the bonding distance changes between the reactant and product configurations, defined as the difference in interatomic distances between them. These node and edge features constitute the GNN graph input, while the target output is the displacement vector of each atom in the true TS configuration relative to the interpolated midpoint of reactant and product configurations.

Additional details for the node and edge features, GNN architecture and data flow is provided in the following section.

Graph neural network training and evaluation

A lightweight edge-aware attention GNN is developed to predict a TS configuration using a given pair of reactant and product configurations. Considering the limited size of the available datasets and the need for fast inference, large-scale models such as highly expressive equivariant GNNs or computationally intensive models such as diffusion-based models are not employed in this work. As shown in Fig. 1c, each retained atom in the selected representative segment becomes a node in the input graph, carrying its atomic number, the interpolated midpoint position between the reactant and product configurations, the displacement vector between them and its magnitude, and a local reaction score that measures the cumulative change in interatomic distances over all edge on that atom. Each edge carries the change and the mean of the interatomic distance between the two configurations, a categorical label indicating whether the contact is formed, broken, or unchanged, the change in the interatomic unit vector, the cosine alignment between the reactant and product bond directions, and the change in alignment between each node's displacement vector and the bond direction.

The GNN network maps this graph to a per-atom displacement from the reactant and product geometries to the predicted TS. Node features are first projected to the hidden dimension by a linear embedding, then updated by a stack of edge-aware attention, message-passing layers, and finally projected to three Cartesian components by a linear output head. Predictions are masked so that only atoms displaced between the reactant and product configurations are assigned a displacement, and the predicted TS geometry is obtained by adding this displacement to the input configuration.

Within each attention layer, illustrated in Fig. 1d, the edge features are first encoded by a multilayer perceptron. The encoded edge representation, together with the features of the two nodes it connects, is used along two parallel branches. The first computes an attention score for each edge, normalized over the neighbors of the receiving node, and an independent sigmoid gate that scales the edge contribution. The second constructs the message itself from the encoded edge representation and the sending node's features. The message is multiplied by both the attention weight and the gate and summed over all neighbors, and the aggregated result is combined with the node's own features through a further multilayer perceptron with layer normalization and a residual connection. Summation rather than

averaging is used for aggregation, as it preserves the magnitude of the localized geometric changes that distinguish a transition state from its neighboring minima.

To optimize model performance, the number of hidden layers varies from 1 to 3, and hidden layer dimensions ranging from 8 to 128 are explored, examining a wide range of GNN architectures while avoiding oversized networks.³⁴ The message-passing layers propagate residual messages to enhance training stability. Model training is performed using the Adam optimizer³⁵ with an initial learning rate of 1×10^{-3} and a weight decay of 1×10^{-4} . A dropout rate of 0.15 is applied to mitigate overfitting given the limited dataset size. The loss function consists of the root-mean-square deviation (RMSD) between the predicted and target displacement vectors, augmented by a penalty term to discourage atomic overlaps. Training is carried out for up to 300 epochs with a batch size of 128. Early stopping with a patience of 30 epochs is employed. If the validation loss does not improve over this interval, training is terminated to prevent overfitting, see tests in **Figure S1**. The model achieving the lowest validation loss is retained as the final model.

Ideally, TS predictions should be invariant to the ordering of the reactant and product configurations, thus swapping the reactant and product positions in the inputs should not alter the predicted TS structure significantly. To enforce this symmetry and augment the limited dataset, reactant and product configurations are swapped for each trajectory, and the resulting data are added back to the dataset and assigned to the same subset as the original frames.³⁶ We found this approach has no significant impact on the training results, but made the learning curve more stable.

Due to the extremely small size of the OC20NEB dataset for surface systems, a transfer learning strategy is adopted. The GNN is first trained on the Transition1x dataset, which is split into training, validation, and test sets using a 60%/20%/20% ratio. The model with the lowest validation loss is saved and used as the initialization for subsequent fine-tuning. During fine-tuning on the OC20NEB dataset, the learning rate is reduced to 1×10^{-4} and the weight decay is increased to 1×10^{-3} . The OC20NEB dataset is split using the same proportions as in the Transition1x dataset, and 5% of the Transition1x data are incorporated to mitigate catastrophic forgetting. The final model is selected based on the lowest validation loss at this stage, and is evaluated using the testing sets of Transition1x and OC20NEB, as well as an isomerization dataset by Grambow et al.²⁶ In addition to RMSD, distance mean absolute error (D-MAE) is used to compare the our model with others. D-MAE quantifies the mean absolute error of pairwise distance matrices between two structures, which focuses more on the relative positions of atoms, while RMSD focuses more on the relative orientation and position differences of the entire structure. All training and evaluation are performed on a NVIDIA A100 GPU using the Perlmutter machine at NERSC.

Transition state search

The TS prediction consists of two main steps. The first step is the generation of reactant, product, and possible intermediate configurations, which is carried out using RANGE.³³ RANGE is a Python package

for atomistic exploration of potential energy surfaces based on a global optimization algorithm that combines the artificial bee colony³⁷ and a genetic algorithm. In RANGE, the energies of candidate structures can be evaluated using first-principles calculations, semi-empirical methods, or empirical and machine-learning-based force fields. For an elementary step in chemical transformations, RANGE explores the configurations of chemical species in both the reactant and product states and drives the search toward energetically favorable configurations, while also identifying low-lying metastable structures. RANGE is agnostic with respect to either gas-phase molecular assemblies or heterogeneous systems involving surfaces or interfaces, as demonstrated in previous studies on gas-phase clusters, supported particles, porous materials, or solvents.^{33, 38, 39}

In this work, for a given reaction step, RANGE is used to identify possible configurations of reactants and products independently and to rank them according to their energies within the respective configurational spaces. All configurations with relative energies within 2 eV of the ground-state structure are retained in each space, although the energy threshold can be user-defined. Configurations from the reactant and product spaces are then paired to generate candidate reactant–product combinations. For each pair, atomic displacement lengths between the reactant and product configurations are computed. If the average of the three largest atomic displacements is below a threshold (3 Å in our study, but can be adjusted by the user), the pair is considered a physically reasonable candidate reaction pair. This procedure yields a set of plausible reactant-product configuration pairs for subsequent TS prediction.

The second step involves TS prediction using the GNN developed in this work. Each identified reactant–product pair is passed to GNN, which predicts a corresponding TS configuration, providing a rapid and computationally inexpensive estimation of the TS structure. These predicted TS structures can be subsequently refined using established TS optimization techniques, such as saddle-point optimization with Sella^{40, 41} or the NEB module via the Atomic Simulation Environment (**ASE**).⁴² Energy evaluations during TS refinement are performed using the same computational methods employed during the reactant and product configuration search in the first step by RANGE to ensure the energy consistency from the same level of theory.

Framework demonstration: propane dehydrogenation

To demonstrate the capability of the proposed framework, we apply it to a heterogeneous catalysis process to determine the propane dehydrogenation pathway on a platinum cluster (Pt₄) supported on carbon nitride (gC₃N₄), previously reported using conventional first-principles calculations with NEB⁴³. The choice of benchmark was motivated by the higher level of complexity of the system and its topical relevance to light alkane conversion to selective product precursors. We reconstruct its reaction pathway and energy landscape and identify additional reaction pathways that were not captured by the conventional approach which pre-determines the positions of reactants. Due to the limited availability of classical force fields for this modeling system, RANGE is coupled with the MACE force field⁴⁴ via the ASE interface for structure search and energy evaluation. For benchmarking the force field, several structures of propane interacting with a tetrahedral Pt₄ on gC₃N₄ were built, and their energies and

atomic forces predicted by various versions of MACE are compared to density functional theory calculations by CP2K⁴⁵, as shown in **Figure S2**. Although none of the MACE force fields gives a highly quantitative accuracy across energies and atomic forces, the MACE mh-1_omat_pbe provides the best overall accuracy of energy ranking without defective atomic force and is used in this work for an initial rapid energy ranking. Subsequent density functional theory calculations are performed on selected structures from MACE ranking to allow for improved accuracy. A gC₃N₄ single-layer model with a XY planar size of 2.14×2.47 nm² and a Z vacuum of 3.0 nm is used as the substrate. RANGE first searches 1000 structures for the location of 4 Pt atoms on gC₃N₄, and we extracted two Pt₄ configurations for further use: the global minimum with a more planar Pt₄ shape, and the tetrahedral shape reported in other work⁴³. The two substrates are geometrically refined by CP2K and kept frozen in the following simulation steps. This constraint is required by the displacement-based graph representation used for TS prediction. Atoms enter the reaction graph when they displace significantly between reactant and product configurations. As a result, independent relaxation of the extended support across configurations would populate the graph with substrate atoms unrelated to the reaction and, given the fixed graph size, crowd out the atoms that are involved in the reaction. This constraint also removes the rigid-body drift of the free-floating substrate atoms. The Pt₄ geometry is sampled independently, so its effects enter as a controlled comparison rather than continuous relaxation in our investigation.

For each substrate, four reaction states are considered: (1) propane adsorption on the substrate, (2) first H abstraction, (3) second H abstraction, and (4) H₂ gas formation and desorption. In the first state, RANGE explores the location of a propane molecule on the Pt₄/gC₃N₄ surface. In the second state, it explores the location of an individual hydrogen atom and a C₃H₇ species on the surface. In the third state, it explores the location of two individual hydrogen atoms and a propene molecule. In the fourth state, it explores the location of a H₂ gas molecule and a propene molecule. For each state, RANGE searches for 4000 potential configurations and ranks these configurations by their energies in each state. All structures from the four states are characterized by the following features: the distance of carbon and hydrogen atoms from the propane molecule to the geometric center (centroid) of the Pt₄ cluster, the distance between carbon and hydrogen in propane that have chemical bond breaking/formation in these reactions, and the height of carbon and hydrogen in propane with respect to the gC₃N₄ surface (average of surface carbon and nitrogen atoms in gC₃N₄). These features are used to perform a principal component analysis (**PCA**) to extract two primary components that differentiate four states on the same map. Then, in each elementary step from one state to the next, all structures from these two states are paired as a reactant–product pair if their maximum atomic displacement is within 3 Å. The trained GNN reads each successful pair and predicts their corresponding TS geometry. If a generated TS structure shows an energy barrier lower than a threshold (2 eV in this case), it is further refined by 9-image NEBs within ASE until the maximum force converges below 0.02 eV/Å. This removes TS structures with unrealistic energy barriers regardless of geometric accessibility. The resulting TS structures and their corresponding reaction state structures are grouped by atom coordinates and structure energies, where the structures showing an average coordination difference below 0.02 Å and

energy difference below 0.02 eV are treated as the same geometry. Based on these grouped geometries, the reaction pathways are generated and re-assembled onto a common x-grid and compared using Euclidean distance between the resulting curves. Pathways were grouped by average-linkage hierarchical clustering, merging curves whose average distance to other cluster members fell below a fixed threshold (15% of the maximum pairwise distance observed within the group). The actual curve closest on average to all others in its cluster is selected as the representative curve. The final selected pathways are re-optimized by the NEB module in CP2K to achieve higher level of accuracy. All CP2K calculations use PBE functional^{46–48} with DZVP-MOLOPT-GTH basis set⁴⁹ and D3 correction⁵⁰.

Results and Discussion

Graph representation and neural networks parameterization

The number of atoms of the structures in the Transition1x database (T1x) ranges from 4 to 23, whereas those in the OC20NEB dataset (OCN) span 19 to 102 atoms due to the presence of substrate surface. However, not all atoms in the OCN are involved in the reactions. Consequently, a single graph representation with appropriate size must serve both gas-phase molecular systems and surface-interaction models, covering the structural features that matter in each without either truncating the larger surface systems or burying the smaller molecular ones in padding noise. We examined this trade-off across the atomic-neighbor cutoff (R_A), the graph-connectivity cutoff (R_G), and the maximum graph dimension ($D_{\max}$) defined in the methods section.

Although the modeling sizes of the two datasets differ, the number of atoms that move between the reactant and product structures is under 28 in both cases, as shown in Fig. 2a. Therefore, any $D_{\max}$ value except 25 is sufficient to capture essential moving atoms involved in the reactions. For the R_A value, which determines the radius used to search for neighboring atoms around the moving atoms, Fig. 2a also reports the resulting graph representation sizes (i.e., the total number of moving atoms plus their neighbors within the R_A cutoff). In T1x dataset (orange), the gas-phase molecules are not constrained by a surface, and most atoms exceed the 0.3 Å displacement threshold we defined in the Methods section. As a result, nearly every atom is included in the graph, regardless of the R_A cutoff. In contrast, for the OCN dataset, the graph size depends strongly on both R_A and $D_{\max}$ values. The graph grows as either R_A (from 3 to 5 Å) or $D_{\max}$ (from 25 to 100) value increases. When $D_{\max}$ exceeds 49, the graph contains many empty nodes relative to the number of actual atoms, introducing padding noise that hinders joint learning across the two datasets. Conversely, the sharp peak in the OCN distribution at $D_{\max}$ of 25 indicates substantial truncation of the surface models, suggesting a potential loss of important structural information. In summary, a $D_{\max}$ value of 36 provides the best tradeoff.

The impact of R_G and the width and depth of GNN hidden layer is further assessed by screening GNNs trained separately on each dataset. The full results of all screened hyperparameters are provided in the supporting materials **Figure S1**, and representative results are illustrated in Fig. 2b. For T1x dataset, the

predictive error is largely insensitive to the cutoff and graph parameters (R_A , $D_{\max}$, R_G), consistent with their minor impacts observed in Fig. 2a. Instead, the GNN architecture matters more. Wider hidden layers from 8 to 32 significantly reduce error, while further widening from 32 to 128 helps only marginally. The increase of layer depth also effectively reduces predictive error, where 2 and 3 layers of the networks show similar level of accuracy. On the contrary, the OCN dataset shows an inverse pattern that deepening or widening the network yields little gain on predictive accuracy, because with a limited number of data, the noise of neural network has dominated over genuine feature variation. Although simply tuning R_G at fixed graph size has little effect, enlarging the graph itself ($D_{\max}$ from 25 to 36 and R_A from 3 to 5) sharply reduces predictive error, underscoring the minimum need to capture the longer-range interaction between surface and reaction species. Further increasing $D_{\max}$ from 36 to 64 does not provide significant gain in the GNN accuracy for this dataset.

The interplay between complexity and accuracy in the model is summarized in Fig. 2c. On T1x dataset with around 10^4 reaction systems, the predictive error falls as trainable parameters increase until the parameter count approaches roughly 2–9 times the number of structures in the database, beyond which added model capacity brings no gain. On OCN dataset, with only < 700 valid reaction systems, the predictive accuracy is essentially flat with respect to GNN model size, which is a clear signature of a data-limited region that motivates transfer learning strategy. Weighing graph representation size against the capacity supported by data, a selection of $R_A=5$, $D_{\max}=36$, $R_G=5$ with a GNN of depth 3 and width 32 makes a suitable set of hyperparameters to pursue a minimal GNN model size and predictive error at the same time (Fig. 2b). This GNN model contains 30089 trainable parameters.

Graph neural networks training and transfer learning

Using the representation and architecture selected above, a GNN model was first trained on T1x dataset exclusively. Small variations in learning rate and weight decay changed accuracy negligibly. Since the labels “reactant” and “product” are interchangeable in practice, an ideal TS predictor should be invariant to their ordering. We first confirmed that swapping the two inputs left predictive accuracy essentially unchanged for the same GNN setting. In addition, a data-augmentation strategy, where these reactant-product swapped pairs are combined with the original dataset, doubles the number of training samples and yields slightly improved predictive accuracy by reducing RMSD by only about 0.05 Å. We refer to this T1x-trained model as GNN_1 in the following discussions.

To extend this model to surfaces despite the scarcity of surface TS data, we fine-tuned GNN_1 on the OCN dataset with a reduced learning rate and increased weight decay as indicated in the method section, yielding another model GNN_2 , and compared the two GNNs and the published models for TS prediction in Fig. 3a. Before fine-tuning, it is not surprising that GNN_1 performs very poorly on OCN since gas-phase training in T1x alone transfers weakly to surface systems, which is structurally absent from T1x. After fine-tuning, GNN_2 improves substantially on OCN dataset (RMSD from 0.445 to 0.235) while

losing very minor accuracy on original T1x dataset (RMSD from 0.216 to 0.221). It is evident that a single compact model can span both data regimes with transfer learning strategies.

Furthermore, these two GNNs are benchmarked against other models for TS prediction on T1x dataset and another independent isomerization data (Iso)²⁶. On the T1x data, our GNNs show 20%–60% higher predictive error than the other single model predictors such as OA-ReactDiff²² and NeuralNEB²¹. On the Iso dataset, which contributed no training data and therefore measures strictly out-of-distribution generalization, our predictive error is comparable to TS-GAN¹⁹, higher than TSDiff²³ and Choi's model²⁰, but lower than the TSNet¹⁸ and Pattanaik's model¹⁷.

Although we do not claim state-of-the-art predictive accuracy of our GNNs, we highlight the outstanding parameter efficiency of our approach when the model accuracy is plotted against the model complexity in Fig. 3b. Those competing models in previous work carry more than two orders of magnitude more parameters than our GNNs, and many of their reliance on SE(3)-equivariant, diffusion, or transformer architectures naturally pushes them to 10^6 – 10^7 parameters, yet they are trained on datasets with only 10^4 structures. Admittedly, whether such models are "too large" is not settled by the raw parameter-to-sample ratio alone. Strong architectural priors, such as equivariance, can reduce the effective degrees of freedom by one to two orders of magnitude below the nominal parameter count, and each TS label supplies full three-dimensional geometric information rather than a single scalar, potentially adding one to two orders of magnitude of effective supervision. These effects synergistically narrow the gap between model capacity and data information content, although achieving the same level of accuracy with a less complex model would still be preferable with limited data.

Crucially, the success of such synergy of model architecture and data depends on dataset coverage and quality, as well as on minimizing information leakage in the transformation from raw atomic structures to network inputs. These conditions are often met with massive datasets in many modern data science applications, but harder to meet in chemical reaction modeling where much less data is available.²⁸ For TS prediction in particular, the available data size on the order of 10^4 cannot hope to represent the full breadth of the reaction space. In this regime, large models risk low in-distribution error alongside degraded transfer to chemically distinct systems. We therefore argue that the model size deserves to be treated as a first-class design consideration alongside accuracy, especially in data-scarce domains such as TS prediction. Our developed GNN models contain only parameters on the order of 10^4 , yet obtain errors comparable or within roughly a factor of two, to other models that are much larger (10^6 – 10^7 parameters). These light-weighted GNNs show a favorable accuracy-complexity trade-off and build the required physical invariances into the representation rather than paying for them in model parameters. Finally, we note the corresponding limitation candidly: for surface systems, generalization remains constrained by the scarcity of high-quality TS data, and until larger and more diverse surface-relevant datasets exist, disciplined control of model capacity, not merely added complexity, is likely the more dependable path to transferable, chemically meaningful predictions.^{15, 28, 51, 52}

Transition-state mapping for propane dehydrogenation

To examine the capability of our TS mapping framework using the trained GNN₂ model, we explore the reaction pathways of propane dehydrogenation (**PDH**) on a gC₃N₄-supported Pt₄. Propane dehydrogenation is a vital industrial process for production of propylene, a primary building block for plastics and chemicals. Studying this pathway on supported transition metals is key to unlocking the precise structural and electronic mechanisms that drive catalyst selectivity and longevity.^{53, 54}

We first constructed the supported cluster by placing four individual Pt atoms on gC₃N₄ substrate and started a search for stable Pt₄ geometries with RANGE as described in the method section. The lowest-energy configuration is a near-planar Pt₄ cluster as shown in Fig. 4a, consistent with a previous study that adopted the same structure to examine Pt, Pd, and Co cluster binding on gC₃N₄⁵⁵, although that work did not establish whether this geometry is in fact the ground state of Pt₄ cluster. RANGE additionally locates two low-lying isomers shown in Fig. 4a, one of which is the tetrahedral Pt₄ geometry employed in an earlier PDH study⁴³ and lies 0.99 eV (95 kJ/mol) above the ground state. Since Pan et al.⁴³ conducted a thorough study using the tetrahedral Pt₄ structure for PDH reaction pathways, we investigate the tetrahedral cluster for direct comparison with Pan's work, but we also examine and compare the pathways discovered on the planar cluster below.

For each cluster, we examined four interaction states that define the dehydrogenation sequence as described in the method section: the reactant state (label R), in which propane adsorbs on the Pt₄/gC₃N₄ surface; the first intermediate (label I1), in which the first C–H bond has broken to give co-adsorbed C₃H₇ and H; the second intermediate (label I2), in which a second C–H scission yields co-adsorbed C₃H₆ and two H atoms; and the final product state (label P), comprising H₂ and propene (C₃H₆) molecules. All energies are referenced to the sum of an isolated gas-phase propane molecule and a bare gC₃N₄ surface, and are reported in kJ/mol to enable direct comparison with the earlier work⁴³. To characterize where each species resides, we track its distance from the Pt₄ center and its height above the gC₃N₄ surface and use these two coordinates to project each state onto a two-dimensional map as illustrated in **Figure S3** and **S4**.

For the R state, RANGE samples a wide range of propane adsorption geometries by probing a propane molecule at different positions on the Pt₄/gC₃N₄ surface, with periodic boundary condition considered. Figure 4b maps the resulting configurations for the tetrahedral Pt₄ system, with the propane–Pt₄ distance on the horizontal axis and the propane height above gC₃N₄ on the vertical axis. In the upper-right region (label A in Fig. 4b), propane sits far from the surface and its energy approaches that of the gas-phase reference (about 0 kJ/mol). As both coordinates decrease and the molecule nears the surface, the energy falls, signifying favorable adsorption (from A to B in Fig. 4b). At a propane–Pt₄ distance of ~ 5 Å and a height of ~ 6 Å (with respect to gC₃N₄ surface), propane binds atop the Pt₄ cluster through a single H atom (label B in Fig. 4b); shorter distances reflect additional contact between

a second propane H and the cluster. Beyond this point the propane–Pt₄ distance only slightly increases while the height continues to fall significantly, tracing the migration of propane from the top of the cluster toward its base, adjacent to the gC₃N₄ substrate (from B to C in Fig. 4b). Finally, the propane height stabilizes while the propane–Pt₄ distance grows substantially, corresponding to propane detaching from Pt₄ and adsorbing on the gC₃N₄ surface (from C to D in Fig. 4b). It is noted that the binding of propane is stronger at the gC₃N₄ support or at the interface of the support and Pt₄ cluster, rather than the Pt₄ cluster itself. This interaction map resolves the full complexity of propane adsorption on Pt₄/gC₃N₄ surface, and shows how binding strength varies with the binding site not only on Pt₄ cluster but also on gC₃N₄ substrate, suggesting the binding is strongest on the cluster/substrate interface rather than the cluster alone. This observation also demonstrates the practical advantages of RANGE sampling with minimal human bias since enumerating such configurations by hand is laborious and prone to oversight. In fact, the identified strongest interaction geometry with propane bound at the base of Pt₄ and close to gC₃N₄ has not been considered in earlier studies.^{43, 55}

The remaining states (I1, I2, and P) are less amenable to a two-coordinate description. Whereas a single propane molecule in the R state is well described by its distance and height, the I1, I2, and P states each contain two or three surface species, so a faithful description must account for the positions of all of them simultaneously. The corresponding maps are therefore more intricate and do not project cleanly onto two dimensions in Fig. 4c–4e, but they still reveal the dominant binding motifs. In state I1, the most stable configuration has H bound between the upper and lower layers of the tetrahedral cluster and C₃H₇ bound to its apex Pt atom. Alternative low-lying arrangements place both C₃H₇ and H on the lower Pt layer, both on the same apex atom, or H on the apex while C₃H₇ occupies a basal Pt atom.

Configurations in which C₃H₇ binds to the gC₃N₄ substrate rather than to Pt₄ are consistently higher in energy. In state I2, the most stable configuration has two H atoms bridging the upper and lower cluster layers with C₃H₆ on the apex Pt. Shifting one H to the apex atom raises the energy modestly, while relocating C₃H₆ to the gC₃N₄ substrate, or placing both H atoms on the apex, raises it further. For state P, the ground-state configuration has H₂ adsorbed on the apex Pt and C₃H₆ binding with two Pt atoms. The interaction of either gas-phase product with the substrate, or simultaneous binding of both to the apex Pt, again yields higher-energy states. These structures cover the geometries considered in Pan et al.⁴³ work, but are generated without binding site assumption and manual building work.

The planar Pt₄ cluster exhibits a comparably rich landscape but similar to the tetrahedral Pt₄ cluster, detailed in **Figure S4** for the binding landscape map and representative structure snapshots. In the R state, the propane can bind to the atop site of Pt₄ cluster, bottom of cluster at the interface with gC₃N₄ substrate, and the gC₃N₄ substrate alone. In the I1 state, the first detached H atom can bind to either the top or middle Pt atom of the cluster, where the top site is preferred. The C₃H₇ binds preferably on the top Pt site. Similarly, in the I2 state, the two H atoms interact with both top and middle Pt atoms, while C₃H₆ prefers the top Pt site. In the P state, H₂ is only weakly absorbed on the Pt cluster and the gC₃N₄

substrate. Overall, the planar Pt₄ cluster shows fewer interaction configurations with the reactive species than the tetrahedral Pt₄ cluster does, because the position of bottom Pt atom in the planar cluster is lower in the gC₃N₄ substrate and this impedes interaction with the reactive species. As a result, fewer Pt atoms are available than the tetrahedral cluster.

For each Pt₄ cluster type, we used PCA as described in the Methods to place the structures from all four states on a common map. The first two principal components capture roughly 60% of the total variance in both cases and separate the four states on the single map shown in **Figure S5**. PCA is used here solely to visualize all structures in a common low-dimensional frame; the analysis below rests on the structures and their TS connections themselves, not on the projection, so the fraction of variance captured does not affect our conclusions. In both systems the heights of the propane carbon and hydrogen atoms above the support emerge as the variables that most strongly discriminate the reaction states. We then paired structures both within and across states according to the criteria in the Methods, applied the GNN₂ model to predict TS for each pair, and refined these predictions with climbing-image NEB (See Methods for details). The validated pairs and their TS structures are projected onto the PCA maps in Fig. 5a and 5b for the tetrahedral and planar clusters, respectively, tracing the reaction progression from R to I1, I2 and P. For both surfaces, the four reaction states occupy distinguishable regions of the PCA map, connected by the corresponding TS structures. The tetrahedral cluster separates the states more sharply, whereas the planar cluster shows greater overlap between them. This difference reflects that the three-dimensional tetrahedral cluster imposes a distinct adsorption configuration at each reaction state, while the more open, quasi-two-dimensional planar cluster permits similar binding arrangements across several states. The same panels also present the connectivity as a graph, in which each structure is a node and an edge links any two structures joined by a predicted TS. Purple edges denote intra-state changes (pure displacement of adsorbed species) and green edges denote inter-state changes (chemical reaction). Both networks show a great number of paths, and we observed much more inter-state edges in the elementary step from I2 to P on the tetrahedral cluster.

To avoid the scaling of edge amounts with the node amounts, we normalized the edge numbers by dividing the number of realized connections by the total number of possible pairs between the relevant states. From the state configurations in Fig. 5c, the node counts differ modestly between the two clusters: the two surfaces show similar R-state configurations, while the tetrahedral cluster yields 17% fewer I1 and 27% fewer I2 configurations but 45% more P-state configurations than the planar cluster. This pattern suggests that the tetrahedral cluster has fewer distinct interactive sites for chemisorbed species, but more possible binding configurations for weakly bound closed-shell products. For the intra-state transition, the tetrahedral cluster is more densely connected within R and I1 states (1.3–1.4 times higher) and shows marginally less connectivity (< 12%) within I2 and P states than the planar cluster, indicating more diverse displacements of adsorbed propane and the first detached H atom on its surface. For the inter-state chemical steps, the tetrahedral cluster has slightly higher normalized connectivity (1.3–1.4 times) for R-to-I1 and I1-to-I2 steps, and a much larger difference (8 times) for I2-to-P step. Since this comparison is normalized by the number of available state-pairs, the difference is

not attributed to unequal sampling. Instead, it suggests that the three-dimensional tetrahedral shape provides more viable paths for two adsorbed H atoms to recombine to a H₂ gas, whereas the two-dimensional planar shape imposes more geometric restrictions for H recombination. Projecting these elementary steps onto energy landscape maps in **Figure S6** reinforces the thesis that the metastable tetrahedral Pt₄ cluster supports a richer pathway connectivity, albeit their accessibility may be hindered by overall higher energetics. The tetrahedral Pt₄ cluster allows for more readily rearrangement of adsorbed species on its surface, particularly for displacements within the I1 and I2 states.

We then enumerated all complete routes connecting the initial state R to the final state P, as shown in **Figure S6**. Not every elementary step joins with others to form an intact pathway. The minimal complete route is a three-step sequence from R to I1, I1 to I2 and I2 to P. We additionally considered four-step routes that augment this sequence with one intra-state displacement (for example, from one I1 configuration to another I1 configuration), which can bridge segments that would otherwise be disconnected. These routes are grouped by hierarchical clustering to group marginally different routes and the medoid of each group are selected as representative pathways (See Methods for details). This gives 8 three-step and 13 four-step representatives for the tetrahedral cluster and 2 and 5 for the planar cluster, consistent with the normalized connectivity analysis above, as shown in Fig. 5d. Note that in Fig. 5d all pathways use one common reference (individual ground-state planar Pt cluster plus a gas-phase propane), so all tetrahedral cluster pathways are displaced upward by 95 kJ/mol compared to the planar cluster pathways. Within each set, the two routes with lowest barriers are also highlighted in Fig. 5d. On both clusters, the first C-H activation shows a higher barrier than the second C-H activation, and the two adsorbed H atoms bind strongly with the Pt clusters, giving a deep energy valley on the reaction landscape, in good agreement with reported trends of Pt.^{43, 56, 57} The recombination of adsorbed H into H₂ gas requires crossing the highest single-step barriers on the pathways to escape the deep energy valley. Overall, the four-step routes show slightly lower energy barriers than the three-step routes, but the barrier reduction is limited (< 11 kJ/mol).

Although the tetrahedral cluster results in denser path connections and lower barriers in general, the tetrahedral Pt₄ geometry lies 95 kJ/mol above the ground-state planar geometry, which cannot be compensated by its lower reaction energy barriers. Therefore, under practical conditions the planar geometry still governs the observed kinetics. Whether a metastable catalyst geometry contributes to turnover depends on the balance between the barriers it offers and the penalty for forming it, and this balance is system-specific and cannot be anticipated without computing both.⁵⁸ In a system with a smaller isomer gap, or under conditions where cluster geometries are kinetically trapped rather than equilibrated, the balance could resolve the other way. This bears directly on the earlier study of this reaction⁴³ that obtains its mechanism on manually built tetrahedral Pt₄ cluster without establishing its position relative to the ground state. Our framework of comprehensive TS mapping indicates that this tetrahedral geometry, while chemically reasonable and reactive, is not likely to dominate the observed kinetics.

To obtain a more quantitative energy landscape, we refined the lowest-barrier pathways at the DFT level (see Methods for details). Both state configurations and TS geometries undergo geometric changes upon DFT-level optimization, and their energies shift accordingly, most notably for the I2-state structures, as shown in the DFT-level energy landscapes in Fig. 6a. The tetrahedral pathways reproduce the previously reported pathways using manually built structures.⁴³ Across the R, I1 and P states on both clusters, together with their TS, the force field and DFT results agree to within roughly 15 kJ/mol. On the other hand, the I2 state, in which two H atoms are chemisorbed on the Pt cluster, shows a discrepancy between the two methods by 20–40 kJ/mol. This is not unexpected, since the MACE force field was neither trained on a broad range of TS structures and energies nor tuned for this particular reaction. Until more robust force fields become available, DFT refinement of the generated structures, both stable states and TS, remains necessary for quantitative barriers. Nonetheless, such a discrepancy in this work is only limited to numerical values of the energy, and the relative ordering of two clusters and enumerated pathway connectivity are all reproduced by DFT. The force field is therefore reliable for exploring and ranking the reaction landscape, while DFT supplies quantitative energetics.

The elementary-step barriers of these DFT-level pathways, illustrated in Fig. 6b, confirm that the tetrahedral Pt cluster has lower energy barriers in the activation of first C-H bond (R to I1, 15–20 kJ/mol lower) and second C-H bond (I1 to I2, 20–60 kJ/mol lower), compared to the planar Pt cluster. However, the last step for H recombination and H₂ release has higher barrier (around 15–20 kJ/mol) on the tetrahedral cluster than on the planar cluster. This is because the tetrahedral cluster stabilizes I2-state intermediate (two co-adsorbed H atoms) more than the planar cluster does, by around 50 kJ/mol. As the representative binding configurations in Fig. 6c show, the undercoordinated Pt atoms make C-H activation facile but also bind the co-adsorbed H and propylene fragment strongly, creating a deep potential well for products to climb out. The observed high barrier for H₂ formation from two co-adsorbed H atoms on Pt cluster is consistent with reports of surface hydrogen accumulation at low-coordination Pt sites in PDH,^{59–61} where strongly bound surface H impedes product release and promotes catalyst deactivation.

In addition, the overall energy barrier along the pathways is only slightly lower (< 20 kJ/mol) on the tetrahedral cluster than on the planar cluster. This falls far short of compensating the 95 kJ/mol required to form the tetrahedral isomer from the planar ground state. Taken together, these results confirm that the planar Pt₄ cluster governs the observed kinetics in practice, and that the intrinsic kinetic advantage of the tetrahedral geometry is realized only if that geometry can be stabilized by additional geometric or chemical constraints.^{58, 62–64} In summary, these findings validate the capability of the framework for automated configurational sampling and TS mapping, and it recovers the established mechanism while also uncovering comparable and competing pathways that manual construction did not reveal. The mechanistic picture emerging from our finding also points to a broader principle: A metastable catalyst geometry could be intrinsically more active than the ground state, but this activity is realized only if the geometry can be stabilized sufficiently to offset its formation penalty. However, whether that balance is favorable cannot be known in advance, because it depends on quantities, such as intrinsic barrier,

binding strength, and formation energy, that must be computed rather than assumed. Tools capable of mapping TS landscapes comprehensively are therefore not merely descriptive but predictive, providing the first-principles basis on which targeted experimental synthesis of stabilized metastable catalysts can be designed.

Discussion

The results above show that coupling global configurational optimization with a compact, physically constrained TS predictor yields more than a faster route to individual saddle points. It produces a connected map of the reaction landscape, including the metastable configurations and transport steps that conventional, ground-state-anchored searches do not visit. Three aspects of that map deserve emphasis. First, unbiased sampling of the adsorbate configurations determines which pathways are visible at all. In the case of PDH, the strongest propane binding occurs not on the Pt cluster but at its interface with gC_3N_4 substrate, a configuration absent from earlier studies of this system that constructed adsorption geometries by hand. Because the operative route under reaction conditions need not begin from the configuration an expert would anticipate, sampling the configuration space without bias is not a convenience but a prerequisite for a correct mechanistic picture.

Second, and more consequential, the catalyst structure itself must be sampled rather than assumed, but sampling it leads to a more nuanced conclusion than is often supposed. In our case, the metastable tetrahedral Pt_4 isomer is intrinsically the better catalyst, because it supports denser pathway connectivity, rearranges adsorbed species more readily, and activates C-H bonds at lower cost than the planar ground state. Yet it lies 95 kJ/mol above the ground state, and its reaction barrier advantage of at most 20 kJ/mol cannot offset this penalty. Therefore, the planar geometry still governs the observed kinetics among the two geometries. The lesson is therefore not that metastable geometries dominate turnover, nor that they can be neglected, but that whether they contribute is a quantitative question with a system-specific answer.^{2, 29} It depends on the competition between intrinsic barriers and formation penalties, and neither quantity is available without explicitly computing both. In a system with a smaller isomer gap, or where cluster geometries are kinetically trapped rather than equilibrated, the balance could resolve the other way. This bears directly on the earlier study of this reaction,⁴³ which built its mechanism on the tetrahedral geometry without establishing its position relative to the ground state. Our results indicate that this geometry, while chemically reasonable and intrinsically reactive, is not the one that determines the rate. Restricting computation to the single lowest-energy geometry can systematically misrepresent a catalyst's behavior. This further generalizes well beyond catalysis that any activated process on a fluxional cluster, nanoparticle, or defective surface is subject to the same caveat.

Third, the explicit inclusion of intra-state displacement steps, such as adsorbate migration and reorientation involving no bond breaking, emphasizes that transport and reactivity are entangled in determining the accessible landscape. In our case of PDH, four-step routes incorporating one such displacement step consistently appear among the low-lying pathways on both clusters. Their

significance is less in the magnitude of direct barrier reduction than in joining pathway segments that would otherwise be disconnected. A pathway search restricted to chemical steps alone would misestimate these routes or even fail to identify them.

More broadly, our computational approach illustrates a strategy of general relevance to machine learning in the physical sciences. Where reference data are scarce, encoding known physical structures into the representation and matching model capacity to the information content of the data offers a more dependable route to transferable predictions than increasing architectural expressiveness. The framework itself is agnostic to whether the underlying system is molecular or extended and applies to any activated process in which the identity of the operative configuration, rather than the accuracy of any single saddle point, is the limiting uncertainty.

Several limitations define the path forward. On the methodological side, the TS predictor's accuracy on surface systems remains bounded by the scarcity of high-quality heterogeneous TS data. Broader and more diverse heterogeneous datasets, together with selective DFT refinement, would tighten quantitative reliability. Related to this, machine learning force fields enhanced by additional TS training data would also improve the efficiency and reliability of our pathway search. Finally, converting the computed landscape into experimentally comparable kinetics requires free-energy corrections, such as zero-point energy and vibrational entropy, and ultimately a microkinetic treatment that propagates the full network of competing pathways into observables such as turnover frequency, apparent activation energy, and the degree of rate control of each step. Such a treatment would also test directly whether the additional pathways uncovered here alter the predicted kinetics relative to a ground-state-only model.

Conclusions

We have introduced a transition-state search framework that combines global configurational searches with a lightweight, physically constrained graph neural network, and demonstrated it on propane dehydrogenation over a carbon-nitride-supported platinum cluster. Rather than scaling model capacity, the approach matches it to the limited transition-state data available for heterogeneous systems and encodes translational, rotational, and permutational invariance directly in the structural representation, reconstructing the reaction landscape at a small fraction of the parameter cost of contemporary deep-learning models. By sampling metastable configurations of both the adsorbed species and the catalyst itself, the framework reaches pathways that ground-state-anchored searches cannot access, and it allows their kinetic relevance to be assessed rather than assumed. In the testing example, this framework reveals that the metastable tetrahedral Pt_4 isomer, although intrinsically more reactive than the planar ground state, is too costly to form to contribute appreciably to turnover. Beyond the demonstrated catalytic system, many activated processes, such as dissociation, nucleation or phase transformation, share the feature that the configurations governing the rate are not known in advance. For these, comprehensive configurational sampling coupled with compact learning offers a practical and transferable route to mapping energy landscapes. Coupling these landscapes to free-energy and

microkinetic analyses and expanding the still-limited corpus of heterogeneous transition-state data will pave a path toward quantitative, predictive modeling of kinetics from first principles.

Declarations

Author Contribution

V.-A. G. conceived and planned the research and co-wrote the manuscript. D. Z. executed the research and prepared the first draft of the manuscript. R. R. edited the manuscript and contributed to discussions.

Acknowledgement

This research was performed at the Oak Ridge National Laboratory and used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. V.-A.G. acknowledges support from the Department of Energy, Office of Basic Energy Sciences, Catalysis Science Program, Grant ERKCC96. D.Z. acknowledges support from the Department of Energy, Energy Earthshot Research Center, Non-Equilibrium Energy Transfer for Efficient Reactions (NEETER) center. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a Department of Energy User Facility using NERSC award BES-ERCAP0036410.

Data availability

The data and code to reproduce results of this work is freely available at GitHub:
<https://github.com/RANGE-kit/RANGE-TS>.

Competing interests

The authors declare no competing interests.

References

1. Wang, Y.-G.; Cantu, D. C.; Lee, M.-S.; Li, J.; Glezakou, V.-A.; Rousseau, R. CO Oxidation on Au/TiO₂: Condition-Dependent Active Sites and Mechanistic Pathways. *Journal of the American Chemical Society* 2016, 138 (33), 10467–10476. DOI: 10.1021/jacs.6b04187.
2. Zhang, Z.; Zandkarimi, B.; Alexandrova, A. N. Ensembles of Metastable States Govern Heterogeneous Catalysis on Dynamic Interfaces. *Accounts of Chemical Research* 2020, 53 (2), 447–458. DOI: 10.1021/acs.accounts.9b00531.

3. Zhai, H.; Alexandrova, A. N. Ensemble-Average Representation of Pt Clusters in Conditions of Catalysis Accessed through GPU Accelerated Deep Neural Network Fitting Global Optimization. *Journal of Chemical Theory and Computation* 2016, 12 (12), 6213–6226. DOI: 10.1021/acs.jctc.6b00994.
4. Moon, S.; Limmer, D. T. A theory for diffusion-controlled reactions within nonequilibrium steady states. *The Journal of Chemical Physics* 2025, 163 (14). DOI: 10.1063/5.0284657.
5. Collinge, G.; Yuk, S. F.; Nguyen, M.-T.; Lee, M.-S.; Glezakou, V.-A.; Rousseau, R. Effect of Collective Dynamics and Anharmonicity on Entropy in Heterogenous Catalysis: Building the Case for Advanced Molecular Simulations. *ACS Catalysis* 2020, 10 (16), 9236–9260. DOI: 10.1021/acscatal.0c01501.
6. Beaglehole, I. W.; Pemberton, M. J.; Farrar, E. H. E.; Grayson, M. N. Machine Learning Transition State Geometries and Applications in Reaction Property Prediction. *WIREs Computational Molecular Science* 2025, 15 (3), e70025. DOI: <https://doi.org/10.1002/wcms.70025>.
7. Wang, Y.; Su, Y.-Q.; Hensen, E. J. M.; Vlachos, D. G. Finite-Temperature Structures of Supported Subnanometer Catalysts Inferred via Statistical Learning and Genetic Algorithm-Based Optimization. *ACS Nano* 2020, 14 (10), 13995–14007. DOI: 10.1021/acsnano.0c06472.
8. Yan, G.; Vlachos, D. G. Unraveling Multiscale Kinetics over Subnanometer Cluster Catalysts: H₂ Desorption from Pt₃(-H)₂/γ-Al₂O₃(110). *ACS Catalysis* 2023, 13 (16), 10602–10614. DOI: 10.1021/acscatal.3c02014.
9. Yin, K.; Gao, P.; Shao, X.; Gao, B.; Liu, H.; Lv, J.; Tse, J. S.; Wang, Y.; Ma, Y. An automated predictor for identifying transition states in solids. *npj Computational Materials* 2020, 6 (1), 16. DOI: 10.1038/s41524-020-0286-9.
10. Zimmerman, P. M. Single-ended transition state finding with the growing string method. *Journal of Computational Chemistry* 2015, 36 (9), 601–611. DOI: <https://doi.org/10.1002/jcc.23833>.
11. Koslover, E. F.; Wales, D. J. Comparison of double-ended transition state search methods. *The Journal of Chemical Physics* 2007, 127 (13). DOI: 10.1063/1.2767621.
12. Jung, H.; Colombi Manzi, E.; Goncalves, T. J.; Bukas, V. J.; De, S.; Margraf, J. T.; Reuter, K.; Heenen, H. H. From global optimization to transition state search: an automated workflow for surface reaction barriers. *npj Computational Materials* 2026, 12 (1), 300. DOI: 10.1038/s41524-026-02298-1.
13. Koda, S.-i.; Saito, S. Locating Transition States by Variational Reaction Path Optimization with an Energy-Derivative-Free Objective Function. *Journal of Chemical Theory and Computation* 2024, 20 (7), 2798–2811. DOI: 10.1021/acs.jctc.3c01246.
14. Schlegel, H. B. Exploring potential energy surfaces for chemical reactions: An overview of some practical methods. *Journal of Computational Chemistry* 2003, 24 (12), 1514–1527. DOI: <https://doi.org/10.1002/jcc.10231>.
15. Wang, X.; Mao, Y.; Wang, Z. Machine learning approaches for transition state prediction. *Chem Catalysis* 2025, 5 (8), 101458. DOI: <https://doi.org/10.1016/j.checat.2025.101458>.
16. Xu, Q.; Delin, A. A general optimization framework for mapping local transition-state networks. *npj Computational Materials* 2026, 12 (1), 112. DOI: 10.1038/s41524-026-01985-3.

17. Pattanaik, L.; Ingraham, J. B.; Grambow, C. A.; Green, W. H. Generating transition states of isomerization reactions with deep learning. *Physical Chemistry Chemical Physics* 2020, 22 (41), 23618–23626. DOI: 10.1039/D0CP04670A.
18. Jackson, R.; Zhang, W.; Pearson, J. TSNet: predicting transition state structures with tensor field networks and transfer learning. *Chemical Science* 2021, 12 (29), 10022–10040, 10.1039/D1SC01206A. DOI: 10.1039/D1SC01206A.
19. Makoś, M. Z.; Verma, N.; Larson, E. C.; Freindorf, M.; Kraka, E. Generative adversarial networks for transition state geometry prediction. *The Journal of Chemical Physics* 2021, 155 (2). DOI: 10.1063/5.0055094.
20. Choi, S. Prediction of transition state structures of gas-phase chemical reactions via machine learning. *Nature Communications* 2023, 14 (1), 1168. DOI: 10.1038/s41467-023-36823-3.
21. Schreiner, M.; Bhowmik, A.; Vegge, T.; Jørgensen, P. B.; Winther, O. NeuralNEB—neural networks can find reaction paths fast. *Machine Learning: Science and Technology* 2022, 3 (4), 045022. DOI: 10.1088/2632-2153/aca23e.
22. Duan, C.; Du, Y.; Jia, H.; Kulik, H. J. Accurate transition state generation with an object-aware equivariant elementary reaction diffusion model. *Nature Computational Science* 2023, 3 (12), 1045–1055. DOI: 10.1038/s43588-023-00563-7.
23. Kim, S.; Woo, J.; Kim, W. Y. Diffusion-based generative AI for exploring transition states from 2D molecular graphs. *Nature Communications* 2024, 15 (1), 341. DOI: 10.1038/s41467-023-44629-6.
24. Si, Y.; Ma, Y.; Yu, T.; Wu, Y.; Liu, Y.; Lai, W.; Zhang, Z.; Shi, J.; Guo, L.; Prezhdo, O. V.; et al. Transition state structure detection with machine learnings. *npj Computational Materials* 2025, 11 (1), 199. DOI: 10.1038/s41524-025-01693-4.
25. Schreiner, M.; Bhowmik, A.; Vegge, T.; Busk, J.; Winther, O. Transition1x - a dataset for building generalizable reactive machine learning potentials. *Scientific Data* 2022, 9 (1), 779. DOI: 10.1038/s41597-022-01870-w.
26. Grambow, C. A.; Pattanaik, L.; Green, W. H. Reactants, products, and transition states of elementary chemical reactions based on quantum chemistry. *Scientific Data* 2020, 7 (1), 137. DOI: 10.1038/s41597-020-0460-4.
27. Wander, B.; Shuaibi, M.; Kitchin, J. R.; Ulissi, Z. W.; Zitnick, C. L. CatTSunami: Accelerating Transition State Energy Calculations with Pretrained Graph Neural Networks. *ACS Catalysis* 2025, 15 (7), 5283–5294. DOI: 10.1021/acscatal.4c04272.
28. Zhang, D.; Smith, B.; Wu, H.; Nguyen, M.-T.; Rousseau, R.; Glezakou, V.-A. Data Analytics for Catalysis Predictions: Are We Ready Yet? *ACS Catalysis* 2024, 14 (10), 8073–8086. DOI: 10.1021/acscatal.3c05285.
29. Lavroff, R. H.; Morgan, Harry W. T.; Zhang, Z.; Poths, P.; Alexandrova, A. N. Ensemble representation of catalytic interfaces: soloists, orchestras, and everything in-between. *Chemical Science* 2022, 13 (27), 8003–8016. DOI: 10.1039/d2sc01367c.

30. Jangid, P.; Chaudhury, S.; Kolomeisky, A. Dynamic Fluctuations Can Counteract Catalytic Inhibition and Enhance Catalytic Performance. *ACS Catalysis* 2026. DOI: 10.1021/acscatal.6c05102.
31. Cheng, S.-W.; Tian, F.-Z.; Wang, G.-B.; Chen, H. M. Making Morphology Meaningful: Tracking Dynamic Catalyst Evolution under Working Conditions. *ACS Applied Materials & Interfaces* 2026, 18 (17), 23893–23904. DOI: 10.1021/acsami.6c04016.
32. Wu, H.; Zhang, J. Dynamic restructuring of electrocatalysts in the activation of small molecules: challenges and opportunities. *Chemical Communications* 2025, 61 (11), 2190–2202. DOI: 10.1039/d4cc05165c.
33. Zhang, D.; Makoś, M. Z.; Rousseau, R.; Glezakou, V.-A. RANGE: A robust adaptive nature-inspired global explorer of potential energy surfaces. *The Journal of Chemical Physics* 2025, 163 (15). DOI: 10.1063/5.0288910.
34. Curteanu, S.; Cartwright, H. Neural networks applied in chemistry. I. Determination of the optimal topology of multilayer perceptron neural networks. *Journal of Chemometrics* 2011, 25 (10), 527–549. DOI: <https://doi.org/10.1002/cem.1401>.
35. Kingma, D. P.; Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* 2014.
36. Kaufman, S.; Rosset, S.; Perlich, C.; Stitelman, O. Leakage in data mining: Formulation, detection, and avoidance. *ACM Trans. Knowl. Discov. Data* 2012, 6 (4), Article 15. DOI: 10.1145/2382577.2382579.
37. Karaboga, D.; Akay, B. A comparative study of Artificial Bee Colony algorithm. *Applied Mathematics and Computation* 2009, 214 (1), 108–132. DOI: <https://doi.org/10.1016/j.amc.2009.03.090>.
38. Jarngal, P.; Jackson, B. A.; Yuk, S. F.; Zhang, D.; Lee, M.-S.; Menziani, M. C.; Glezakou, V.-A.; Rousseau, R.; Piccini, G. Protonation dynamics of confined ethanol–water mixtures in H-ZSM-5 from machine learning-driven metadynamics. *Journal of Catalysis* 2026, 454, 116658. DOI: <https://doi.org/10.1016/j.jcat.2025.116658>.
39. Gao, J.; Zhang, D.; Zhu, Z.; Parker, G. D.; Rousseau, R.; Yu, X.-Y. In situ molecular imaging of ion clusters reveals the acid gas capture capacity and mechanism of water-lean ionic liquids. *Carbon Capture Science & Technology* 2026, 19, 100622. DOI: <https://doi.org/10.1016/j.ccst.2026.100622>.
40. Hermes, E. D.; Sargsyan, K.; Najm, H. N.; Zádor, J. Accelerated Saddle Point Refinement through Full Exploitation of Partial Hessian Diagonalization. *Journal of Chemical Theory and Computation* 2019, 15 (11), 6536–6549. DOI: 10.1021/acs.jctc.9b00869.
41. Hermes, E. D.; Sargsyan, K.; Najm, H. N.; Zádor, J. Sella, an Open-Source Automation-Friendly Molecular Saddle Point Optimizer. *Journal of Chemical Theory and Computation* 2022, 18 (11), 6974–6988. DOI: 10.1021/acs.jctc.2c00395.
42. Hjorth Larsen, A.; Jørgen Mortensen, J.; Blomqvist, J.; Castelli, I. E.; Christensen, R.; Dułak, M.; Friis, J.; Groves, M. N.; Hammer, B.; Hargus, C.; et al. The atomic simulation environment—a Python library for working with atoms. *Journal of Physics: Condensed Matter* 2017, 29 (27), 273002. DOI: 10.1088/1361-648X/aa680e.

43. Pan, J.; Strugovshchikov, E.; Salóm-Català, A.; Novell-Leruth, G.; Kaźmierczak, K.; Curulla-Ferré, D.; Carbó, J. J.; Godard, C.; Ricart, J. M. Propane Dehydrogenation on Pt Single-Atom and Pt₄ and Pt₃Sn Single-Cluster Supported on g-C₃N₄: A Theoretical Study. *The Journal of Physical Chemistry C* 2025, 129 (5), 2477–2487. DOI: 10.1021/acs.jpcc.4c07686.
44. Kovács, D. P.; Batatia, I.; Arany, E. S.; Csányi, G. Evaluation of the MACE force field architecture: From medicinal chemistry to materials science. *The Journal of Chemical Physics* 2023, 159 (4), 044118. DOI: 10.1063/5.0155322.
45. Kühne, T. D.; Iannuzzi, M.; Del Ben, M.; Rybkin, V. V.; Seewald, P.; Stein, F.; Laino, T.; Khaliullin, R. Z.; Schütt, O.; Schiffmann, F.; et al. CP2K: An electronic structure and molecular dynamics software package - Quickstep: Efficient and accurate electronic structure calculations. *The Journal of Chemical Physics* 2020, 152 (19). DOI: 10.1063/5.0007045.
46. Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Physical Review Letters* 1996, 77 (18), 3865–3868. DOI: 10.1103/PhysRevLett.77.3865.
47. Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple [Phys. Rev. Lett. 77, 3865 (1996)]. *Physical Review Letters* 1997, 78 (7), 1396–1396. DOI: 10.1103/PhysRevLett.78.1396.
48. Goedecker, S.; Teter, M.; Hutter, J. Separable dual-space Gaussian pseudopotentials. *Physical Review B* 1996, 54 (3), 1703–1710. DOI: 10.1103/PhysRevB.54.1703.
49. VandeVondele, J.; Hutter, J. Gaussian basis sets for accurate calculations on molecular systems in gas and condensed phases. *The Journal of Chemical Physics* 2007, 127 (11). DOI: 10.1063/1.2770708.
50. 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. *The Journal of Chemical Physics* 2010, 132 (15). DOI: 10.1063/1.3382344.
51. Darouich, S.; Toney, J. W.; Luo, W.; Kästner, J.; Niepert, M.; Kulik, H. J. Robust generative transition-state models for unseen chemistry. *Nature Computational Science* 2026. DOI: 10.1038/s43588-026-01034-5.
52. Keith, J. A.; Vassilev-Galindo, V.; Cheng, B.; Chmiela, S.; Gastegger, M.; Müller, K.-R.; Tkatchenko, A. Combining Machine Learning and Computational Chemistry for Predictive Insights Into Chemical Systems. *Chemical Reviews* 2021, 121 (16), 9816–9872. DOI: 10.1021/acs.chemrev.1c00107.
53. He, Z.; Yang, J.; Liu, L. Design of Supported Metal Catalysts and Systems for Propane Dehydrogenation. *JACS Au* 2024, 4 (11), 4084–4109. DOI: 10.1021/jacsau.4c00730.
54. Yang, F.; Zhang, J.; Shi, Z.; Chen, J.; Wang, G.; He, J.; Zhao, J.; Zhuo, R.; Wang, R. Advanced design and development of catalysts in propane dehydrogenation. *Nanoscale* 2022, 14 (28), 9963–9988. DOI: 10.1039/D2NR02208G.
55. Duraisamy, P. d.; Mallik, B. S. Decorating graphitic carbon nitride monolayer with cobalt-incorporated Pd₄ and Pt₄ nanoparticles for reversible hydrogen storage. *Surfaces and Interfaces* 2026, 92, 109302. DOI: <https://doi.org/10.1016/j.surfin.2026.109302>.

56. Hauser, A. W.; Gomes, J.; Bajdich, M.; Head-Gordon, M.; Bell, A. T. Subnanometer-sized Pt/Sn alloy cluster catalysts for the dehydrogenation of linear alkanes. *Physical Chemistry Chemical Physics* 2013, 15 (47), 20727–20734. DOI: 10.1039/c3cp53796j.
57. Zhu, J.; Yang, M.-L.; Yu, Y.; Zhu, Y.-A.; Sui, Z.-J.; Zhou, X.-G.; Holmen, A.; Chen, D. Size-Dependent Reaction Mechanism and Kinetics for Propane Dehydrogenation over Pt Catalysts. *ACS Catalysis* 2015, 5 (11), 6310–6319. DOI: 10.1021/acscatal.5b01423.
58. Sun, G.; Fuller, J. T., III; Alexandrova, A. N.; Sautet, P. Global Activity Search Uncovers Reaction Induced Concomitant Catalyst Restructuring for Alkane Dissociation on Model Pt Catalysts. *ACS Catalysis* 2021, 11 (3), 1877–1885. DOI: 10.1021/acscatal.0c05421.
59. Yang, M.-L.; Zhu, Y.-A.; Fan, C.; Sui, Z.-J.; Chen, D.; Zhou, X.-G. DFT study of propane dehydrogenation on Pt catalyst: effects of step sites. *Physical Chemistry Chemical Physics* 2011, 13 (8), 3257–3267. DOI: 10.1039/c0cp00341g.
60. Liu, X.; Yao, Z.; Lai, J.; Li, J.; Wang, J. Hydrogen coverage effects and structure sensitivity in propane dehydrogenation on Pt nanoparticles. *Chemical Engineering Science* 2026, 319, 122299. DOI: <https://doi.org/10.1016/j.ces.2025.122299>.
61. He, K.; Zheng, J.; Su, X.; Qu, Z.; Yang, Y.; Jiao, J.; Li, Q.; Song, L. First-principles insights into the negative kinetic effects of hydrogen in propane dehydrogenation over (Sub)nanometer platinum catalysts. *Applied Surface Science* 2026, 748, 167800. DOI: <https://doi.org/10.1016/j.apsusc.2026.167800>.
62. Rong, H.; Ji, S.; Zhang, J.; Wang, D.; Li, Y. Synthetic strategies of supported atomic clusters for heterogeneous catalysis. *Nature Communications* 2020, 11 (1), 5884. DOI: 10.1038/s41467-020-19571-6.
63. Wang, H.; Gao, Z.; Sun, B.; Mu, S.; Dang, F.; Guo, X.; Ma, D.; Shi, C. Engineering metal-support interaction to construct catalytic interfaces and redisperse metal nanoparticles. *Chem Catalysis* 2023, 3 (10). DOI: 10.1016/j.checat.2023.100768.
64. Sun, G.; Sautet, P. Active Site Fluxional Restructuring as a New Paradigm in Triggering Reaction Activity for Nanocluster Catalysis. *Accounts of Chemical Research* 2021, 54 (20), 3841–3849. DOI: 10.1021/acs.accounts.1c00413.

Figures




Figure 1

(a) An overview of the computational framework developed in this work for reaction pathway mapping, including three stages: the structural search for reactant (R) and product (P) configurations, the prediction of TS configuration for each appropriate R/P pair, and the reaction pathway mapping using connected R, TS, and P configurations. (b) The detailed two steps in the TS prediction stage: the generation of representative fragments for molecular or surface modeling of R and P configurations, and the prediction of TS configuration using converted graph from the representative fragment. (c) The

detailed steps within the graph neural networks from the input of representative fragments to the output of TS atomic coordinates. (d) The detailed message-passing steps in each attention layer of graph neural networks.




Figure 2

(a) Distribution of the number of moving atoms with their neighboring atoms under different atomic neighbor cutoffs (R_A) and graph size upper boundaries (D_{max}) in the two datasets. (b) Example of GNN predictive error on the testing subset using different graph parameters and GNN width and depth for the two datasets. (c) The GNN model capacity-accuracy relationship for the two datasets, with data points

including all the different combinations of graph setting and GNN hyperparameters. The dashed vertical lines mark the number of reactions included in the datasets.




Figure 3

(a) TS predictive error using different models for three datasets: OC20NEB (OCN), Transition1x (T1x), and Isomerization (Iso). Bars filled with patterns indicate RMSE values, and empty bars indicate D-MAE values. (b) The predictive performance of different models against their model complexity for two datasets (T1x and Iso). The number of parameters for Pattanaik's and Choi's models is an estimation based on their model architecture due to the lack of available specific model from their reports. The estimated degree of freedom (DOF) of data is highlighted in the shaded area.




Figure 4

(a) The modeling system size (nm) of the Pt₄/gC₃N₄ surface, and the snapshots of three Pt₄ cluster configurations. Two top-view snapshots highlight the planar and tetrahedral configurations. Pt, C and N atoms are in orange, gray, and blue color. (b) The energy landscape of the reactant state (R) for propane adsorption on tetrahedral Pt₄ cluster, mapped as a function of propane-Pt₄ distance (D_{Pt/C_3H_8}) and propane height to gC₃N₄ surface ($Z_{C_3H_8}$). Structure snapshots on the right side highlight representative structures labeled on the map. Propane molecule is indicated in green. (c) Energy landscape of first

intermediate state (I1) on tetrahedral cluster and representative configuration snapshots. (d) Energy landscape of second intermediate state (I2) on tetrahedral cluster and representative configuration snapshots. (e) Energy landscape of product state (P) on tetrahedral cluster and representative configuration snapshots.




Figure 5

Comparison of the PCA maps for reaction landscapes between (a) tetrahedral and (b) planar Pt_4 clusters. The graph on the right side in each case shows the connectivity of different reaction states. Nodes indicate unique structures, with green and purple edges indicating inter- and intra-state change.

(c) The number of unique binding configurations and the normalized number of edges on the two clusters. (d) The energy landscape of representative three-step and four-step reaction pathways, with the two lowest-barrier paths in each set highlighted by thicker lines.




Figure 6

(a) DFT-level reaction pathways for PDH on the planar and tetrahedral Pt₄ clusters, refined from the lowest barrier pathways in force field-based mapping. (b) Energy barriers of each elementary chemical reaction in the DFT-refined pathways. The “Overall” bars indicate the highest barriers of all the elementary steps in each pathway. Structural snapshots of representative binding configurations in each state is highlighted for (c) tetrahedral cluster and (d) planar cluster. C, H, and Pt atoms are shown in gray, white, and orange, respectively. The carbon nitride support is rendered as a wireframe for visual clarity.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [GOGNNSI.docx](#)