A Foundational Model for Reaction Networks on Metal Surfaces

Santiago Morandi[1,2,†], Oliver Loveday[1,2,†], Tim Renningholtz[1], Sergio Pablo-García[3,4,6], Rodrigo A. Vargas-Hernández[5], Ranga Rohit Seemakurthi[1], Pol Sanz Berman[1,2], Rodrigo García-Muelas[1], Alán Aspuru-Guzik[3,4,6,7,8,9], and Núria López[1*]

[1]Institute of Chemical Research of Catalonia, The Barcelona Institute of Science and Technology, Av. Països Catalans 16, 43007, Tarragona, Spain
[2]Department of Physical and Inorganic Chemistry, Universitat Rovira i Virgili, Campus Sescelades, N4 Block, C. Marcel·lí Domingo 1, 43007, Tarragona, Spain
[3]Department of Chemistry, University of Toronto, Lash Miller Chemical Laboratories 80 St. George Street, ON M5S 3H6, Toronto, Canada
[4]Department of Computer Science, University of Toronto, Sandford Fleming Building, 40 St. George Street, ON M5S 2E4, Toronto, Canada
[5]Department of Chemistry & Chemical Biology, MacMaster University 1280 Main Street West, L8S 4L8, Hamilton, Canada
[6]Vector Institute for Artificial Intelligence, 661 University Ave. Suite 710, ON M5G 1M1, Toronto, Canada
[7]Department of Materials Science and Engineering, University of Toronto, 184 College St., M5S 3E4, Toronto, Canada
[8]Lebovic Fellow, Canadian Institute for Advanced Research (CIFAR), 661 University Ave., M5G 1M1, Toronto, Canada
[9]Director, Acceleration Consortium
[†]Authors Equally Contributed to this Work
[*]Corresponding Author: nlopez@iciq.es

June 28, 2024

Abstract

Process optimization in heterogeneous catalysis relies on the control of competing reactions. The reaction mechanisms based on chemical knowledge can be evaluated *via* density functional theory unveiling experimental catalytic trends. However, this approach finds its limits when applied to complex reaction networks or large molecules, disregarding alternative paths and rare events. Here we present CARE, a foundational model for catalysis on metal surfaces with a rule-based reaction network generator for $C_xH_yO_z$ species built with GAME-Net-UQ, a graph neural network with uncertainty quantification targeting thermodynamic and kinetic parameters, coupled to microkinetic modeling. CARE reproduces experimental activity trends in methanol decomposition, selectivity to C_3 products in electrochemical reduction processes, and models the Fischer-Tropsch synthesis to C_6 products, including 370k reactions, breaking the current limits of network exploration. This comprehensive model opens the path towards the exploration of thermal and electrocatalytic surface processes previously not amenable to atomistic simulations.

1 Introduction

Catalytic processes rely on the control of complex competing multi-step chemical reactions networks (CRNs) [1, 2]. The core of a CRN consists of a set of species **S** linked by a set of reactions **R** defined by properties **P**, usually including, but not limited to their reaction and activation energy [3, 4]. For processes where the list of elementary steps can be hand-generated, density functional theory (DFT) can provide linear reaction profiles. But as network complexity grows [5], this approach finds its limits and it is necessary to leverage the CRNs abstract mathematical foundations [6–11].

CRN algorithms fall into two categories. CRN-PES (Potential Energy Surface) models require the on-the-fly *ab initio* energy evaluation of the species **S** involved along the reaction paths [12–17]. As such, they are only amenable to chemical spaces of the order of 10^2 intermediates. Alternatively, rule-based CRN models employ templates to define the attainable chemical space *via* a list of possible transformations [18]. These algorithms can decouple the network generation from the energy evaluation. Therefore, they lead to larger networks compared to CRN-PES, including multiple paths which could contribute to the catalytic activity or expand rare events responsible for catalyst degradation, thus getting closer than CRN-PES to achieve the completeness of the chemical space. However, the combinatorial explosion of the chemical space requires strategies to reduce the CRN complexity and/or substitute DFT with fast energy estimators.

The energy evaluation of the species **S** can be accelerated by employing group additivity (GA) methods [19–21], while linear scaling (LSR) [22, 23]

and Brønsted-Evans-Polanyi (BEP) relationships [24–26] can be used for characterizing the reaction properties **P**. An example of this approach is the Reaction Mechanism Generator (RMG) [27], that was applied to the CO_2 methanation, and hydrocarbon (CH_4, C_2H_x) oxidation on Pt(111) and Ni(111) surfaces [28]. Upon iterative species generation with graph-based approaches characterized with the aforementioned energy estimators (GA-LSR-BEP), a filtering is applied after every network expansion through microkinetics to quench the chemical space. Although the approach is self-consistent and obtains general reactivity trends, the accuracy of the species energy is low (about 0.5 eV) compared to DFT [29], and cannot tackle selectivity issues, limiting its applicability.

The competition and correlation between network completeness and the associated computational burden when addressing characterization of the chemical space is key to advance the field. Expanding the boundaries of the investigated chemical spaces requires more robust, accurate, and efficient energy regressors, which can be obtained *via* Machine Learning (ML) acceleration strategies [30–32]. Ulissi *et al.* [33] built a surrogate model with GA fingerprints fed to a Gaussian process and combined it with BEP relationships to obtain the rate-limiting steps for the conversion of syngas to C_2 products on Rh catalysts. This model was employed within an active learning strategy while keeping track of uncertainty to predict the most relevant steps to be refined with DFT. An alternative approach involved employing ML potentials from the Open Catalyst Project (OCP) [30] for refining initial guesses of the adsorption geometries subsequently relaxed with DFT [29], reducing the total computational time. Although speeding-

up DFT simulations, these models are still too computationally demanding for high-throughput *in-silico* exploration. As for the transition state (TS) and corresponding activation energy, going beyond the use of BEPs has been demonstrated via multiple ML regressors on dispersed datasets for $C_xH_yO_z$ species on metals, achieving MAEs in the order of 0.2 eV [34, 35]. A promising recent strategy [36] employs OCP models, explicitly trained for predicting adsorption energies, for targeting TS structures, achieving an MAE of 0.1 eV. Additional strategies involving ML potentials [37–39], graph-based approaches [40] and diffusion models [41] can reduce the prediction error, but have yet to be merged with CRN algorithms.

In this context, we present the Catalytic Automated Reaction Evaluator (CARE), a foundational model for heterogeneous catalysis consisting of (i) a rule-based CRN generator for processes involving $C_xH_yO_z$ species, (ii) GAME-Net-UQ, a robust and accurate graph neural network with uncertainty quantification for evaluating the thermodynamic and kinetic reaction parameters on 42 metal surfaces, and (iii) microkinetic modeling functionalities to analyze the reactivity. CARE is benchmarked against experimental data to rank catalytic activity in methanol decomposition, employed to study selectivity to C_3 products in electrochemical reduction on Cu, and in the Fischer-Tropsch process to C_6 products by evaluating 370k elementary reactions, demonstrating the wide impact of the methodology.

2 Results

CARE (*Catalytic Automated Reaction Evaluator*) is a foundational model able to construct reaction networks and to fast scan the activity and selectivity in heterogeneous and electrocatalysis for reactions involving $C_xH_yO_z$ species on transition metal surfaces. CARE consists of three independent modules: A template-based network generator, a fast evaluator of the thermodynamic and kinetic reaction parameters, and a robust microkinetic solver (**Fig. 1**).

Building the CRN in CARE proceeds in three steps, (i) the creation of a general network scheme, the CRN blueprint, composed of a list of species **S** and reactions **R**, (ii) the evaluation of the properties **P**, consisting of adsorbate placement and energy evaluation and (iii) the analysis of the CRN with microkinetic modeling (MKM).

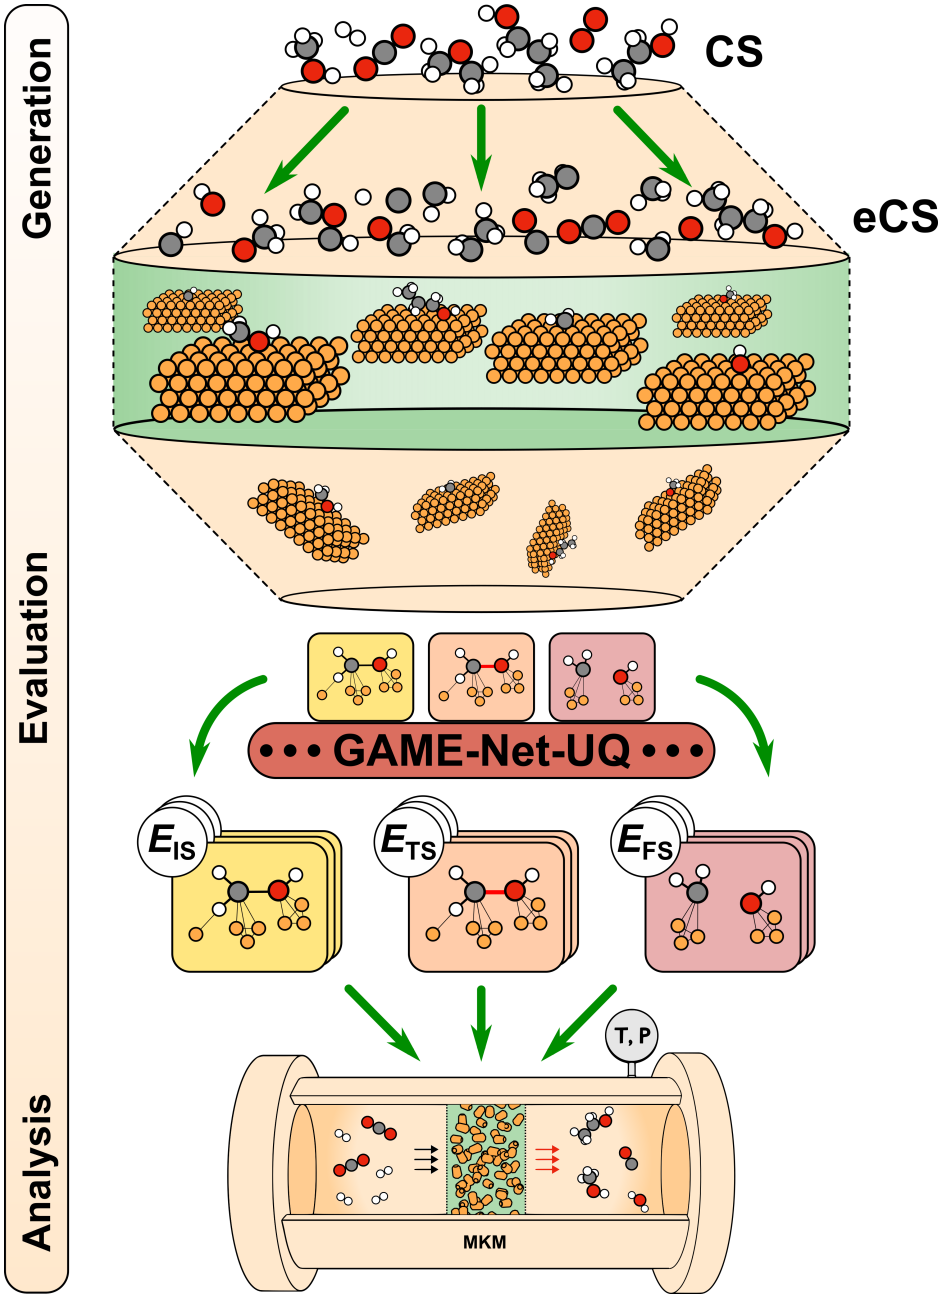

Fig. 1: CARE workflow from network construction to reactivity quantification CARE starts from a Chemical Space (CS) defined by the maximum number of C and O atoms and a molecular template defining how to build the carbon backbone. The bond-breaking template returns the extended Chemical Space (eCS) and all potential reactions. Species in the eCS are placed on the metal surface and their energy and those of the linking transition states are evaluated with GAME-Net-UQ, a graph neural network with uncertainty quantification. The reactivity of the network can be characterized with microkinetic modeling (MKM) under realistic reaction conditions.

2.1 Network blueprint generation

The first step for generating the CRN blueprint in CARE is the definition of the chemical space (CS). The size of the CS is determined by the backbone of the largest molecule obtained by establishing the maximum number of C and O atoms for the species, defined by the carbon and oxygen cutoffs, N_{C} and N_{O} respectively. The CS and the set of templates are implemented using RDKit [42] operations on SMILES [43]. To set up the CS, a pool of alkanes with up to N_{C} carbon atoms is generated. These are then oxygenated and new saturated molecules with up to N_{O} oxygen atoms are obtained (*e.g.*, ethers, alcohols, epoxides). This pool of species is complemented by a gas-phase reservoir of molecules usually encountered in $\mathrm{C}_x\mathrm{H}_y\mathrm{O}_z$ networks, consisting of $\mathrm{H_2}$, $\mathrm{H_2O}$, CO, $\mathrm{CO_2}$ and $\mathrm{O_2}$, composing the CS.

In a second step, the algorithm expands the CS by applying a series of bond-breaking templates generating simultaneously a set of reactions **R**, open-shell fragments, and unsaturated closed-shell molecules (*e.g.*, alkenes, ketones, acids). The newly obtained species plus the CS define the extended chemical space (eCS). In contrast to homogeneous catalysis, the number of possible reaction templates in heterogeneous catalysis is relatively small but the number of events can be very large, as in polymerization/depolymerization processes involving long-chain hydrocarbons. Therefore, the main reaction templates in CARE describe *X*-*Y* (*X, Y*=C, H, O) bond-breaking reactions, and adsorption/desorption of closed-shell and other key molecules. Additional reaction templates account for intramolec-

ular rearrangements such as 1,2-hydrogen shift, and a specific template for proton-coupled electron transfer (PCET) steps enabling modeling reaction networks under electrochemical conditions. The modularity of the network generator allows the inclusion of more reaction classes if required, for instance reactivity in solution at the interfaces of electrodes.

The ensemble of the eCS and the set of reactions **R** constitute the reaction network blueprint. It is catalyst agnostic and needs to be generated only once to model the same network on multiple catalysts and conditions. Extended details on the CRN blueprint generation algorithm are discussed in **Supplementary Note 1**.

2.2 Thermodynamic and kinetic parameters estimation

To estimate the thermodynamic and kinetic properties (**P**) of the obtained CRN blueprint, a fast energy regressor algorithm is needed. In CARE, this task is done by GAME-Net-UQ, a graph neural network (GNN) targeting the DFT energy of adsorbed species **S** and the transitions states linking them. Originally, GAME-Net [31] was designed to predict the energy of closed-shell molecules in gas phase and adsorbed on the closed-packed surface of 14 transition metal surfaces. The upgraded GAME-Net-UQ model includes three crucial advancements: (i) the extension to open-shell species, (ii) the estimation of transition state (TS) energy and (iii) the implementation of uncertainty quantification (UQ) for the predictions. Including these features keeps GAME-Net-UQ lightweight (558k parameters, roughly twice

the original GAME-Net), accurate, and robust, enabling extensive direct energy evaluations at minimal computational cost compared to state-of-the-art ML potentials containing up to 10^{7-8} parameters [44].

The graph dataset used to train and evaluate the model includes 12,608 entries, 80% of them are intermediates (2% in gas phase) and the remaining 20% are TSs representing bond-breaking events of species up to $C_2H_6O_2$ (**Supplementary Fig. 4**). To distinguish the TS graphs from intermediates, a boolean feature is added to the graph edges to label the bond involved in the reaction, following a similar approach to the virtual bond described by Wang *et al.* [40] (**Fig. 2a**). Consequently, the GNN architecture includes a topology-adaptive convolutional layer [45] utilizing the encoded edge attribute. The GNN is trained with graphs representing the DFT-optimized geometries (**Methods Section 4.1**) and a streamlined version of DockOnSurf [46] is employed as a pre-processing step to place adsorbates, replacing the need of DFT-optimized input geometries during inference. Adsorbate atoms with coordination lower than their valence are defined as anchoring points, while the surface adsorption sites are identified with ACAT [47].

The representation of the surface has been expanded to allow the study of structure sensitive processes. Surface effects are taken into account by including (i) species adsorbed on the 2^{nd} and 3^{rd} most stable surface facets of the metals in the training dataset (**Supplementary Fig. 5**), (ii) the generalized coordination number [23] of surface atoms as node feature and (iii) the 2-hop adsorbate nearest surface neighbors to distinguish hollow adsorption sites.

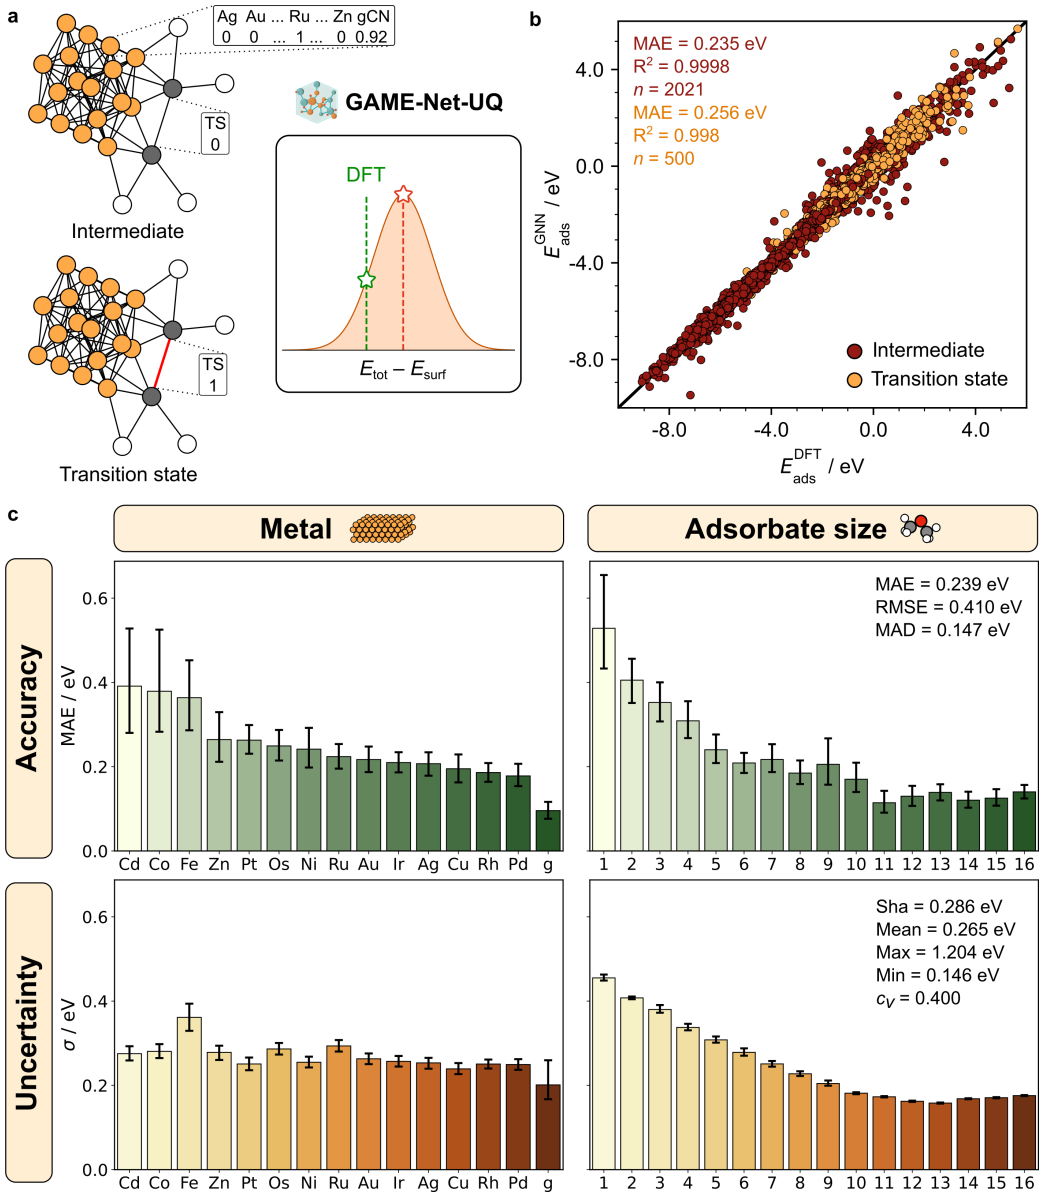

Fig. 2: GAME-Net-UQ performance. **a** Graph representation of intermediates and transition states (TS), and energy prediction as normal distribution. **b** Parity plot of the predicted *vs.* DFT adsorption energy for the test intermediates (brown) and TS (orange) data. **c** Mean absolute error (MAE) and predicted mean standard deviation grouped by metal and adsorbate size. Error bars represent the 95 % confidence interval, 'g' refers to gas-phase graphs. RMSE=Root Mean Squared Error. MAD=Median Absolute Deviation. Sha=Sharpness, c_v=Variation coefficient (dimensionless). Statistic metrics refer to the whole test set ($n = 2521$), which has been randomly sampled from the original data distribution.

As GAME-Net-UQ has been redefined as a mean-variance estimator [48], it returns DFT energy predictions as normal distributions $E \sim \mathcal{N}(\mu, \sigma^2)$ with the standard deviation σ representing the uncertainty (**Fig. 2a**). Details about training data, model architecture, UQ and training procedure can be found in **Supplementary Notes 3-6**.

In terms of accuracy, GAME-Net-UQ yields an MAE of 0.24 eV on a test set of 2521 data points randomly sampled from the initial dataset, 0.23 eV for intermediates and 0.26 eV for transition states (**Fig. 2b**). The performed ablation study (**Supplementary Note 7**) shows that labeling the edge involved in the TS improves the test MAE by 0.14 eV, while encoding the generalized coordination number in the graph nodes yields an overall MAE reduction of 0.10 eV. Data with small (high) error are reflected in small (high) uncertainty, resulting in a robust model with a 0.9% miscalibration area (**Supplementary Fig. 11c**). Small adsorbates, ferromagnetic metals, and Cd and Zn show higher error (**Fig. 2c**), but the model uncertainty for these systems is positively correlated. The sharpness of the model, representing the magnitude of its uncertainty on the test set, is 0.29 eV, while the variation coefficient, which quantifies the dispersion of the uncertainty estimates, is 0.4 (**Supplementary Fig. 11d**).

GAME-Net-UQ propagates the uncertainty to the reactions energy ΔE and activation barrier E_{act}, also defined as normal distributions (**Supplementary Note 8**). Since the GNN is trained on single adsorbates on the surface, bimolecular surface reactions $A^*+B^* \rightarrow C^*+^*$ are energetically evaluated in the bond-breaking direction ($C^*+^* \rightarrow A^*+B^*$) to construct the TS from the reactant graph by labelling the corresponding broken bond. The products

are then evaluated separately, assuming absence of lateral interactions. As fallback strategy, CARE automatically treats steps as barrierless if GAME-Net-UQ predicts TS energies between the initial and final states. For electrochemical processes, the computational hydrogen electrode (CHE) [49] approach is used to include the effect of applied potential and electrolyte pH to the energies. Reactions involving species not bonded to the surface, such as adsorption/desorption and PCET steps involving -OH removal and the formation of water, can be evaluated with GAME-Net-UQ as it has been trained with closed-shell gas phase molecules.

UQ can be exploited within an active learning strategy to refine with DFT the intermediates with the highest number of participated elementary reactions, the so-called 'hubs' [11], and to propagate the uncertainty to microkinetic simulations. The energy estimator interface in CARE can be connected to databases to avoid the estimation of the energy intermediates already evaluated at the DFT level. While this strategy leads to minimal computational savings, the actual gain is a reduced uncertainty and improved accuracy for the reactions involving the intermediates evaluated with DFT. As an example, employing the DFT energy of the adsorbed H^*, the main hub in the C_1O_2 CRN, positively affects the evaluation of 48% of the reactions in the network by reducing their error and associated uncertainty.

2.3 CRN analysis via microkinetic modeling

Having a fully characterized reaction network is insufficient to directly map the results to the experimental reactivity of materials. Inferring the reactivity from linear reaction profiles involving hundreds of intermediates and reactions is not straightforward since the CRNs in CARE (i) have undefined directionality at the global level as this depends on the definition of the initial reactants, (ii) are unconstrained to specific source and target species, and (iii) the catalytic activity depends on the applied operating conditions. Therefore, mean-field microkinetic modeling (MKM) functionalities [50] for both thermal and electrochemical processes are included in CARE.

Microkinetic modeling relies on the numerical steady-state solution of the ordinary differential equations (ODEs) systems defining the material balance for the species in the CRN. The default reactor model in CARE is a zero-conversion differential plug-flow reactor, which enables the extraction of macroscopic reaction rates and selectivity directly comparable to experiments, steady-state surface coverages, and rate-determining steps and species. The kinetic coefficients of the elementary reactions in the network are evaluated *via* transition state theory for all forward reactions except the adsorption steps, the coefficients of which are given by the Hertz-Knudsen equation [51]. The coefficients of the reverse reactions are obtained by imposing thermodynamic consistency (see **Supplementary Note 8** for implementation details).

Elementary reactions possess characteristic times ranging multiple or-

ders of magnitude, resulting into a stiff ODE system requiring specialized solvers. Popular Python libraries such as Scipy [52] lead to an intractable computational cost for kinetic simulations on huge CRNs and/or stiff ODEs. Therefore, CARE additionally implements the Julia solvers `DifferentialEquations.jl` and `DiffEqGPU.jl` [53, 54]. These outperform Scipy solvers and are prepared to be run on GPUs, providing a speed-up of $\times$20-100 compared to their CPU counterparts. CARE employs the backward differentiation formula (BDF) as default ODE solver (`BDF` in Scipy, `FBDF` in Julia). 64-bit floating-point arithmetic is implemented for microkinetic simulations on both CPU and GPU (although the latter is typically optimized for 32-bit floating-point).

Two fundamental aspects to consider in MKM simulations are the absolute and relative tolerances (`atol` and `rtol`) applied at each integration step, and the definition of the steady-state termination event. Initial `atol` and `rtol` are conservatively set to 10^{-16} and 10^{-8}. These values are iteratively increased by the solver in case it fails to converge. Steady-state is reached when the sum of the absolute derivatives of the species surface coverage is less than 10^{-10} (**Supplementary Fig. 13**). Fallback strategies are additionally available to facilitate the convergence of microkinetic simulations showing high stiffness and to reduce the total computational time. These include (i) removal of surface reactions with energy barriers higher than a specified threshold, (ii) considering all the reactions as barrierless (**Supplementary Fig. 14**), and (iii) steering the model towards specified target products by removing from the CRN species which are unlikely to be formed based on the user intuition. Details about the code implementation

can be found in **Methods Section 4.2**.

2.4 Application to industrially relevant problems

To highlight CARE capabilities, we have selected three applications of industrial relevance and increasing complexity. We start by benchmarking activity against an open experimental dataset. Obtaining systematic studies of experimentally measured reactivity is not straightforward but the CatTestHub [55] database presents a good example, changing the metal on the same carrier and under the same conditions for methanol decomposition. The second example illustrates a chemical reaction network in electrochemical environments. Selectivity is key to processes and typically depends on a set of two competing steps with small energy differences, in the range of few meV, thus being challenging in terms of accuracy, even more for data-driven estimators. To check the validity of CARE in predicting selectivity patterns in electrochemical applications, we analyzed the steps towards C_3 products on Cu(100) [56] in the C_3O_2 ($N_C = 3$, $N_O = 2$) CRN. Finally, we try to assess the capacity of the methods by analyzing a reaction network that cannot be addressed with DFT due to the number of species and steps involved, the Fischer-Tropsch process for producing fuels. For this case study, we generated and evaluated the C_6O_1 CRN on four typically employed metals (Co, Fe, Ni, and Ru) for this process, resulting in networks of 40k intermediates and 370k elementary reactions.

2.4.1 Methanol decomposition

First, a benchmark against experimental reactivity is established by studying the methanol decomposition to CO on different metal catalysts. This example focuses on a small CS (N_C=1, N_O=2) but covers a high-throughput investigation spanning multiple metal catalysts, predicting reactivity on different materials [55].

Methanol (MeOH) has gained importance for sustainable energy applications as hydrogen vector [57, 58]. Developing efficient catalysts for its controlled decomposition is therefore crucial. The public experimental database CatTestHub [55] contains 104 kinetic experiments for this reaction on 9 transition metal catalysts, covering multiple supports and operating conditions. Here we choose a smaller set with the same inert carrier and experimental conditions (**Supplementary Table 6**).

To benchmark our results, we first built the C_1O_2 CRN for all the metals on the (111), (110), and (100) surface facets for the face-centered cubic (fcc) metals, and (0001), (10$\bar{1}$0), (10$\bar{1}$1) for Ru. These CRNs include 28 surface species, 10 gas phase molecules, and 62 elementary reactions. We evaluate the network energetics with GAME-Net-UQ and run multiple microkinetic simulations at 473 K and 1 atm, as reported in the database. Details about the CRNs and MKM runs can be found in **Supplementary Note 9**.

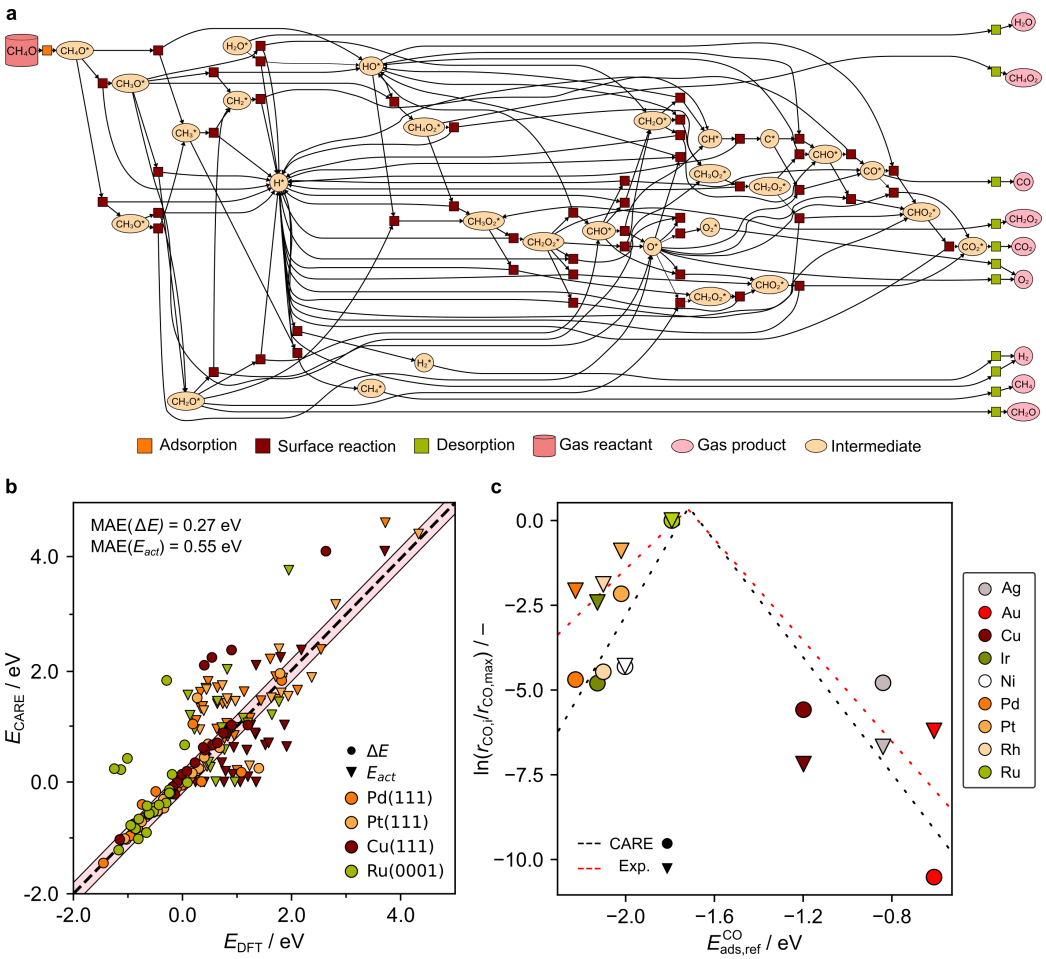

Fig. 3: Methanol decomposition on metal catalysts. **a** C_1O_2 CRN visualization. Species with same node label represent constitutional isomers. **b** Parity plot of the activation (triangles) and reaction (circles) energies obtained with GAME-Net-UQ and those from DFT simulations as reported in Ref. [59]. The highlighted region represents the parity line ±0.2 eV. n=88. **c** CO logarithmic normalized formation rate from microkinetic simulations (circles) and experiments (triangles) as function of the DFT CO adsorption energy, considering the contribution from all the available surface orientations weighted by their fractional area obtained from Ref. [60]. Experimental values are retrieved from CatTestHub [55]. T=473 K, P=1 bar, y_{MeOH}=10%. Nickel is depicted in white as experimental characterization reveals low dispersion compared to other metals (**Supplementary Table 6**).

Fig. 3a provides an overview of a C_1O_2 CRN. CARE represents CRNs and microkinetic simulations outputs as bipartite graphs with nodes repre-

senting species and elementary reactions, with edges $e(s, r)$ embedding the consumption/formation rate of species s due to reaction r.

Comparing GAME-Net-UQ predictions to available DFT data [59] on a subset of the studied metal surfaces, CARE shows an accuracy in terms of MAE of 0.27 eV for the reaction energy, and 0.55 eV for the activation energy (**Fig. 3b**). DFT studies typically focus on modeling reactive processes on the most stable surface facet. However, the observed catalytic activity in metal nanoparticles is attributed to the contributions from multiple exposed surfaces. The ability of GAME-Net-UQ to target multiple surfaces allows for predicting activity on different surfaces in a systematic way. Defining the overall reaction rate as a weighted sum of the rates on multiple surface facets leads to a general agreement with the experimental trend (**Fig. 3c**), while focusing on the most stable facet (*e.g.*, (111) facet for fcc metals) would underestimate these predictions (**Supplementary Fig. 16**).

Generating the 27 CRNs (3 surfaces for each of the 9 metals) and running the microkinetic simulations for all the metals and surface orientations took only 8 minutes on a 24-core CPU (3.2 CPU hours). Simulating all species and transition states in the network, adsorbed on diamagnetic (all except Ni) surfaces, with DFT would require at least 11000 CPU hours and 91 kWh (see **Supplementary Note 10**).

Our fast data-driven energy estimator yields a trend in agreement with experimental data. As for the absolute values, the predicted rates differ by 4-6 orders of magnitude from the experiments (**Supplementary Fig. S17**). This discrepancy, often encountered in mean-field kinetics

based on DFT data [61], stems from the infinite dilution in DFT simulations which minimizes the effect of lateral interactions, and the missing entropy contributions, factors that, if included, would likely increase the global reaction rates by destabilizing the reaction path.

2.4.2 Electrochemical reduction to C_3 products

The second example focuses on the electrochemical formation of C_3 products. Electrochemical reduction to C_3 products on Cu(100) has been studied with CARE by generating a C_3O_2 CRN. Compared to thermal processes, electrochemical reactions require the inclusion of the potential and pH contributions, which CARE implements *via* the computational hydrogen electrode (CHE) approach, allowing the definition of a set of specific electrochemical reactions between adsorbed species, charged particles as protons and electrons, and the solvent. The CRN has been modeled in neutral (pH=7) and alkaline (pH=13) conditions, at applied potentials of –0.4 and –1.0 V_{SHE}, and at ambient temperature and pressure (298 K, 1 bar). The network contains 893 surface species, 93 gas-phase molecules and 7709 elementary reactions, classified in five main types: Proton-coupled electron transfer (PCET), C-O, C-C and C-H bond-breaking/forming reactions, and [1,2]-H shifts (**Fig. 4a**). Further characterization of the C_3O_2 network can be found in **Supplementary Fig. 18**. The generation and energy evaluation of the CRNs took less than 5 minutes on a 24-cores CPU (2 CPU hours), making use of multiprocessing to evaluate the energy of the intermediates.

The global analysis of all reactions in the CRN (**Supplementary**

Fig. 19) shows that PCET steps cluster into two distinct groups according to their uncertainty. The cluster with higher uncertainty corresponds to -OH protonation leading to water formation and elimination, likely related to different orientations in the GNNs representing the three-dimensional structure.

As for pH effects, shifting towards a more alkaline pH (from pH 7 to pH 13) leads to an increase of the reaction energy for all PCET reactions. This behaviour is expected since increasing the pH implies a lower concentration of protons, thermodynamically disfavoring the protonation of the intermediates. As expected, lowering the applied potential (from –1.0 to –1.5 V_{SHE}) leads to a favorable thermodynamic shift of the energies for all PCET reactions.

A key problem in CO_2 reduction on Cu is fine-tuning its selectivity towards long-chain valuable products. CARE is able to reproduce selectivity trends, for instance, for C_2 products, ethylene and ethanol are typically observed but in the C_3 fraction only 1-propanol is observed and the route of propylene is blocked. CARE was used to reproduce the paths starting from two key intermediates, allyl alcohol and propionaldehyde. The selected paths leading to 1-propanol (**Fig. 4b**) and its competition with propylene (**Fig. 4c**), previously evaluated at the DFT level in Ref. [56] have been taken to analyze selectivity. The results obtained with CARE are comparable with those obtained with DFT, demonstrating the capacity of CARE to identify trends. When subjected to alkaline pH (**Fig. 4b**) and reductive conditions (**Fig. 4c**), thermodynamically favorable pathways towards the formation of 1-propanol are identified. These findings align with

previous DFT studies [56] and highlight the preferential formation of 1-propanol over propylene.

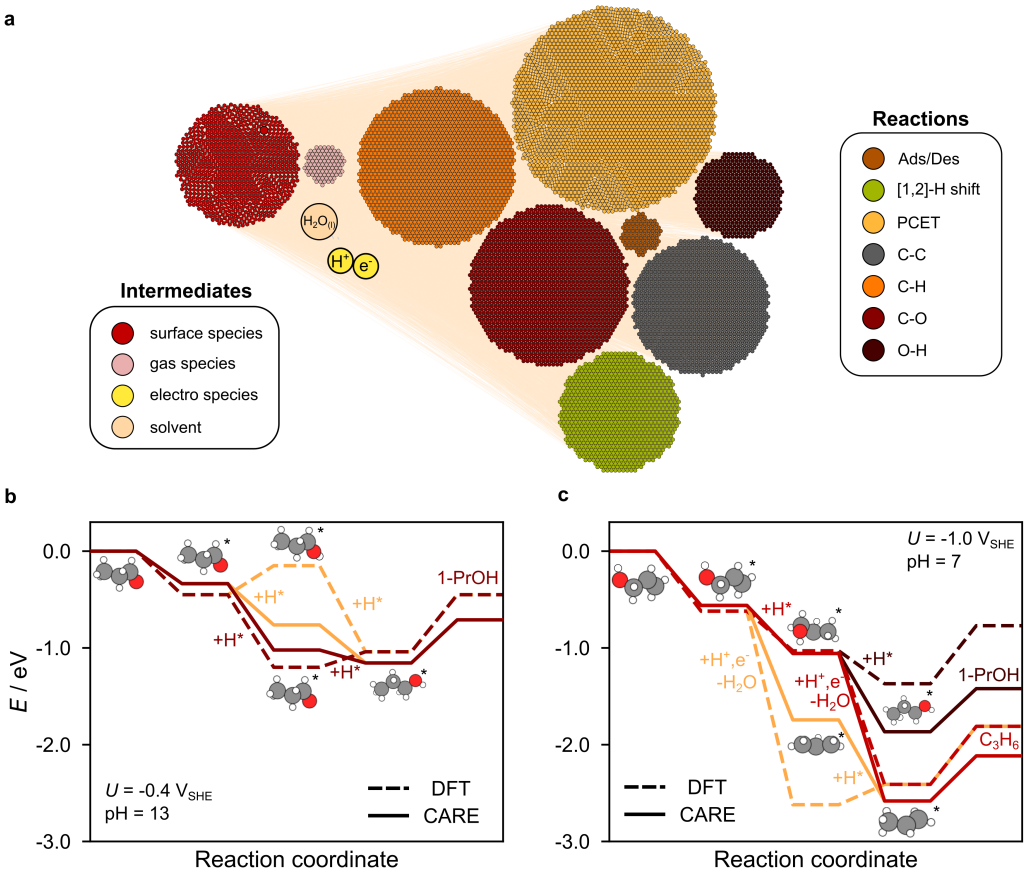

Fig. 4: Electrochemical C_3O_2 reaction network on Cu(100). **a** CRN visualization with intermediates (left) and elementary reactions (right) colored by phase and reaction type, respectively. Ads/Des=Adsorption/Desorption, PCET=Proton-Coupled Electron Transfer. Reaction profiles reproduced with CARE from Ref. [56] (dashed lines) for the electrocatalytic paths from **b** propionaldehyde to 1-propanol and **c** allyl alcohol to propylene and 1-propanol.

As for the selected paths at zero applied potential (U=0 V_{SHE}) (**Supplementary Fig. 20**) a shift in trend occurs, being in this case the path from allyl alcohol to propylene more thermodynamically favourable than that to 1-propanol, in line with DFT studies [56].

Therefore, CARE can deal with electrochemical systems by integrating a specific template and evaluating the energies taking into account pH and applied potential effects. In this way, selectivity trends comparable to DFT and experimental observations can be extracted.

2.4.3 Fischer-Tropsch process

This case study shows how CARE can construct huge CRNs, by focusing on the Fischer-Tropsch process [62], which converts syngas to a wide distribution of long-chain hydrocarbons (alkanes and olefins) used as synthetic clean fuels. To this end, a C_6O_1 ($N_C = 6$, $N_O = 1$) CRN has been generated and evaluated for Co(0001), Fe(110), Ru(0001) and Ni(111). The obtained CRNs include 38,913 surface species, 985 gas-phase molecules and 369,365 elementary reactions. Considering that simulating an intermediate adsorbed on ferromagnetic surfaces with DFT on average required 307 CPU hours and about 2.5 kWh on supercomputing facilities (**Supplementary Note 10**), simulating the extended chemical space of this CRN with DFT would have taken at least 10^6 kWh and 1360 CPU years, assuming sampling one configuration per adsorbate-surface pair. With CARE, each CRN blueprint generation plus the energy evaluation of its intermediates and transition states took 4 hours on a 24-core CPU (96 CPU hours), resulting in a speed up of $\times 10^6$ and a substantial reduction of the energy footprint.

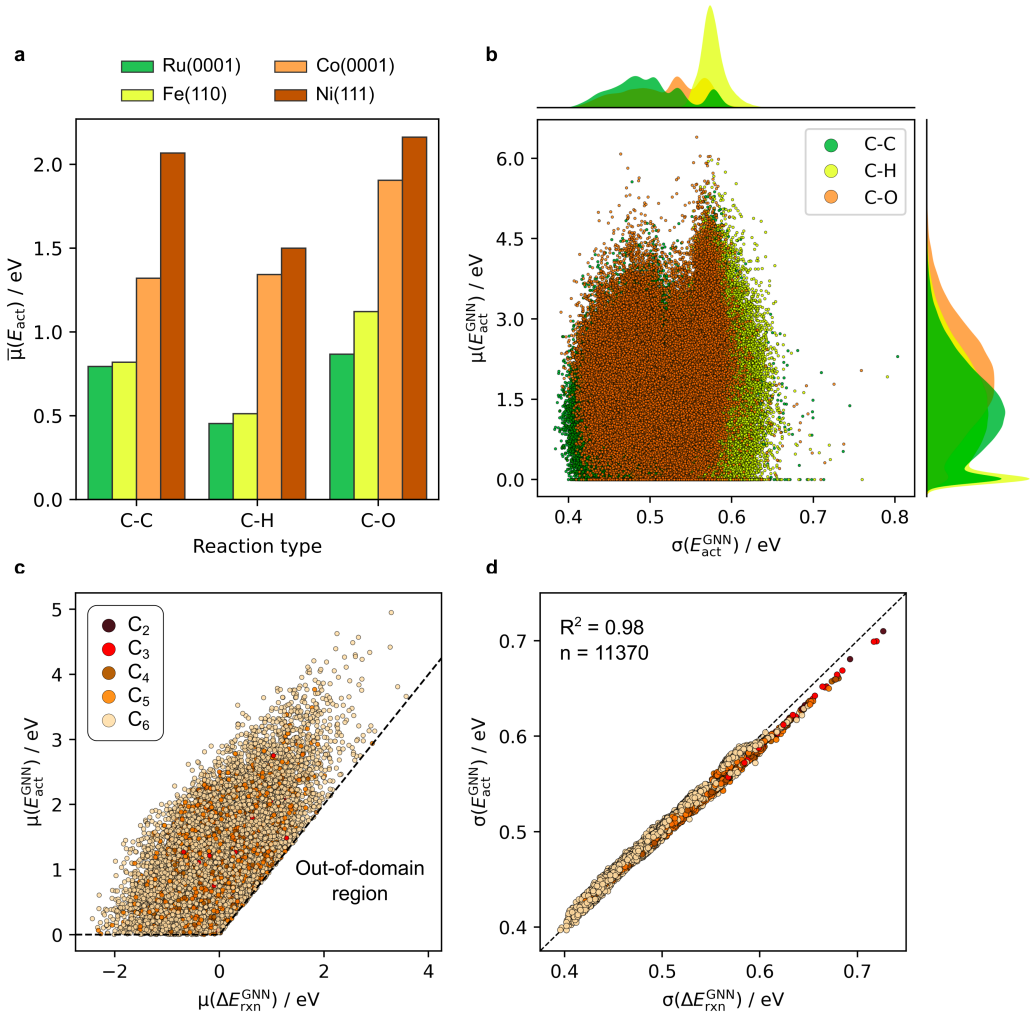

Fig. 5: Energy trends obtained from the C_6O_1 CRN for the Fischer-Tropsch process. **a** Bar-plot of the mean activation energy obtained with GAME-Net-UQ grouped by bond-forming type and metal surface. **b** Scatter-plot of the activation energy and related uncertainty (standard deviation) for all the surface reactions in the bond-forming direction of the C_6O_1 CRN on Co(0001) ($n = 357,222$). **c** Mean and **d** standard deviation of the predicted activation energy as function of the reaction energy for the C-C coupling reactions on Co(0001). $n_{\text{C-C}} = 11370$.

Fig. 5a provides energy trends in agreement with experimental evidence. Ru is the most active metal catalyst as it provides the lowest activation barriers among the four studied metals, and possesses a high

hydrogenation activity (C-H > C-C > C-O), however its cost limits its industrial application. Nickel is the least active catalyst and prefers C-H over C-C bond-forming reactions, while Cobalt, known for its selectivity towards long-chain products, is the only metal whose mean energy barrier for C-C coupling events is smaller than the competitive C-H and C-O bond-forming reactions, explaining its higher chain growth probability compared to the other metals.

In terms of uncertainty, a standard deviation between 0.4-0.6 eV (**Fig. 5b**) for the energy barriers is obtained with GAME-Net-UQ for Co. The greater uncertainty in the activation energy, compared to values for intermediate or transition state structures, arises from the propagation of uncertainty (*i.e.*, standard deviation σ) when deriving reaction properties from species, which involves performing addition and subtraction operations on normal distributions. The smallest uncertainty is obtained for C-C coupling events, with a minimum standard deviation of 0.4 eV, while hydrogenation reactions show an uniform uncertainty distribution around 0.6 eV. Similar trends are observed for the other metals (**Supplementary Fig. 21**).

A relevant consequence of GAME-Net-UQ targeting both intermediates and TSs is the perfect correlation between the uncertainty found for the reaction energy ΔE and activation barrier E_{act}, as observed in the C-C coupling events on Co(0001) (**Fig. 5d**). This feature originates from the similar performance of GAME-Net-UQ for intermediates and TS (0.24 *vs.* 0.26 eV).

Although the presented CRNs require prohibitive resources to be simulated at the microkinetic level, the case study highlights the capability of

CARE to generate and evaluate huge networks cheaply including all the possible bond-breaking reactions that could eventually take place on the catalyst surface.

3 Discussion

CARE presents a fast and robust software tool for evaluating the catalytic activity of metal catalysts *via* a reaction network generator powered by a data-driven energy estimator plugged into a simple reactor model. Being the first of this kind, we have identified several potential challenges in the field.

We have approached the reaction network completeness with a set of templates including all potential bond-breaking reactions leading to all the available intermediate fragments, even those that might be discarded initially based on chemical intuition. CARE is modular in nature, allowing the user to implement new templates for increasing the granularity in the definition of the explored chemical space. This systematic approach reduces bias in CRNs, but leads to larger reaction lists that can be a bottleneck when analyzing and manipulating the constructed networks. However, there are three definitive advantages in the methodology. The first is the modularity of the algorithm, allowing the implementation of additional molecular and reaction templates, thus getting as close to network completeness as desired, for instance, by including competing reactions on surfaces and solution. CARE can consider rare events easily, providing the possibility for studying their implication for the origin of catalyst deactivation [63], an

aspect typically overlooked by traditional DFT simulations but crucial for improving the long-term stability of the catalyst. The third element is the substantial reduction of the computational cost by accelerating the energy evaluation of the reaction networks with GAME-Net-UQ.

Our data-driven energy estimator is fast and relatively cheap due to its small size and direct approach for predicting the DFT energy of relaxed structures, differently than machine learning potentials, which have two orders of magnitude more parameters and require multiple iterations to converge. The uncertainty quantification of the predictions provides an additional feature which can be propagated to reaction properties and microkinetic modeling, or exploited for active learning purposes and Bayesian optimization. The relatively low computational cost of our approach and the balance in the uncertainty quantification between intermediates and transition states removes the need for empirical rules like linear-scaling relationships, with GAME-Net-UQ covering a wider diversity in terms of metals, surfaces, and species, and being more accurate. The implemented interface of CARE for this model allows the energy estimation of thousands of intermediates/transition states in less than 1 CPU hour, representing a viable ML strategy for fast property prediction. Again, challenges in this area are linked to the introduction of lateral interactions (like solvents or high coverage), more detailed configurational analysis for adsorbates and expanding the structural diversity of metal surfaces. These aspects could be addressed by implementing lateral interactions via GNN [64], active learning and Δ-ML schemes [65]. As for the energy estimator, this could be first-principles or data-driven, allowing the study of processes on catalysts

different from the ones presented here such as alloys or metal oxides.

As for the challenges in the microkinetic part, our implementation with Julia allows to go beyond limitations in capacity as number of intermediates and reactions. Future improvements would require a strategy to include entropic contributions and a more detailed representation of the electrochemical events could be introduced. Clever CRN pruning strategies are needed particularly for very stiff ODEs to improve the predictive power and detailed rate and selectivity control tools can also be beneficial in unveiling important catalytic pathways. Finally, more complex reactor models able to capture relevant phenomena at the nanoparticle scale could be implemented.

In summary, CARE opens the path to the exploration and characterization of catalytic processes on metals with improved network completeness and reduced bias in an automated and robust and accurate manner. Our work paves the way towards the study of processes involving complex compounds such as plastics and biomass at a reasonable computational cost.

4 Methods

4.1 Density Functional Theory

The DFT simulations for developing GAME-Net-UQ have been performed with the Vienna Ab-initio Simulation Package, VASP 5.4.4 [66]. The Perdew–Burke–Ernzerhof (PBE) [67] functional with reparameterized D2 [68] dispersion correction for metals was used [69]. Core electrons were rep-

resented by projector-augmented wave pseudopotentials [70] and valence electrons were represented by plane waves with a kinetic energy cutoff of 450 eV. Electronic convergence was set to 10^{-5} eV and atomic positions were converged until residual forces fell below $0.03\,\mathrm{eV}\cdot\mathrm{\AA}^{-1}$.

The metals provided with CARE include 8 face-centered cubic (fcc), 1 body-centered cubic (bcc), and 5 hexagonal close-packed (hcp) structures on the three most stable *hkl* surface facets for each one, (**Supplementary Fig. 2**). Metal surfaces were modeled by four to ten layer slabs (**Supplementary Table 1**), where the half uppermost layers were fully relaxed and the bottom ones were fixed to the bulk distances. A surface coverage concentration of $0.02\,\mathrm{molecules}\cdot\mathrm{\AA}^{-2}$ was defined for all the adsorption structures, a reasonable value to neglect lateral interactions. The vacuum between the slabs was set between 13 and 16 Å and the dipole correction was applied in z direction [71]. The Brillouin zone was sampled by a Γ-centered $3\times3\times1$ k-points mesh generated through the Monkhorst–Pack method [72]. Transition states were obtained with the improved dimer method [73].

4.2 Microkinetic modeling

Once the CRNs are constructed, three ingredients are required for running microkinetic simulations:

1. Stoichiometric matrix $\nu \in \mathbb{Z}^{n_\mathrm{S},n_\mathrm{R}}$ representing the reaction network with n_S species (including the surface active site "*") and n_R elementary reactions. $\nu_{i,j}$ is the stoichiometric coefficient of species i in reaction j. As elementary reactions are mono- or bi-molecular,

stoichiometric coefficients lie in the interval [-2;+2]. This matrix is intrinsically sparse, and assuming that the reactions in the CRN are of the kind $A + B \rightarrow C + D$, its sparsity s_ν (*i.e.*, percentage of zero values) follows the relationship (**Supplementary Fig. 12d**):

$$s_\nu = \left(1 - \frac{4}{n_\mathrm{S}}\right) \times 100\% \tag{1}$$

Compressed sparse row/column (CSR/CSC) formats and integer data type (`Int8`) have to be preferred to optimize matrix multiplications for CRNs with more than 400 intermediates ($s_\nu = 99\%$).

2. Operating conditions, such as temperature, pressure, and applied potential and electrolyte pH for electrochemical systems. These values define the thermodynamic and kinetic constants of the elementary reactions (**Supplementary Note 4**).

3. Reactor model. The default in CARE is a zero-conversion differential plug-flow reactor (`DifferentialPFR`), defined by the ODE:

$$\begin{cases} \dfrac{d\theta_i}{dt} = \displaystyle\sum_{j=1}^{n_R} \nu_{i,j} r_j & \text{adsorbed species} \\ \dfrac{dP_i}{dt} = 0 & \text{gas species} \end{cases} \tag{2}$$

where θ_i is the fractional surface coverage of species i, $\nu_{i,j}$ the stoichiometric coefficient of species i in the reaction j, P_i the partial pressure of the gas-phase reactant i. r_j is the net rate of reaction j,

defined as:

$$r_j = \vec{r}_{f,j} - \vec{r}_{r,j} = \vec{k}_{f,j} \prod_{i=1}^{n_S} y_i^{|min(\nu_{i,j},0)|} - \overleftarrow{k}_{r,j} \prod_{i=1}^{n_S} y_i^{|max(\nu_{i,j},0)|} \tag{3}$$

where $r_{f,j}$ and $r_{r,j}$ are the rates of reaction j in the forward and reverse direction, $\vec{k}_{f,j}$ and $\overleftarrow{k}_{r,j}$ the respective kinetic constants, and y_i refers to θ_i or P_i, depending on the phase of species i. The initial conditions assume an empty surface exposed to the gas mixture defined by the user:

$$\begin{cases} \theta_*(t=0) = 1 & \text{active site} \\ \theta_k(t=0) = 0 & \text{surface species} \\ P_i(t=0) = P_i^0 & i \in \text{gas reactants} \\ P_j(t=0) = 0 & j \in \text{gas products} \end{cases} \tag{4}$$

4.3 Computational Tools

CARE is written in Python 3.11, the CRN algorithm mainly depends on RDKit 2023.09.5, the Atomic Simulation Environment (ASE) 3.22.1 [74], and NetworkX 3.2.1 [75]. GAME-Net-UQ has been developed with Pytorch Geometric 2.4.0 [76] and Pytorch 2.0.1 [77], and its training has been performed on a NVIDIA RTX A2000 12GB GPU with CUDA 12.3. Microkinetic modeling functionalities are implemented with Scipy 1.12.0 [52] and DifferentialEquations.jl [53] (Julia 1.10.2). Visualizations have been created with Matplotlib 3.8.0 [78], Seaborn 0.12.2 [79], and Inkscape 1.3.2.

5 Code availability

CARE implementation has been publicly released under the MIT license and is available on github.com/LopezGroup-ICIQ/care. The code for training and evaluating GAME-Net-UQ is also available on GitHub at github.com/LopezGroup-ICIQ/gamenet_uq. To improve the reproducibility of this work, the version of the codes will be frozen and uploaded in Zenodo as static repository.

6 Data availability

The DFT data used to develop GAME-Net-UQ will be available as ASE database in Zenodo when published and in ioChem-BD at 10.19061/iochem-bd-1-257 (FG-dataset, already accessible) and 10.19061/iochem-bd-1-257 (new data, under embargo until publication).

7 Acknowledgements

This publication was created as part of NCCR Catalysis (grant number 180544), a National Centre of Competence in Research funded by the Swiss National Science Foundation. O. L. acknowledges the Joan Oró Predoctoral Programme of the Generalitat de Catalunya, and the European Social Fund Plus (2023 FI-1 00769). T.R acknowledges the support of the Erasmus+ program of the European Union. S. P.-G. acknowledges support from the U.S. Department of Energy, Office of Science, Subaward by University of Minnesota, Project title: Development of Machine Learning and Molecular

Simulation Approaches to Accelerate the Discovery of Porous Materials for Energy-Relevant Applications (DE-SC0023454).

R.R.S. acknowledges funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no. 754510. S.M., O.L., R.R.S., P. S. B. and N.L. thank the Spanish Ministry of Science and Innovation (PID2021-122516OBI00 and PRE2022-101291) and Severo Ochoa (CEX2019-000925-S) for the financial support.

A.A.-G. acknowledges support from the Canada 150 Research Chairs Program as well as Anders G. Fröseth. The authors thank BSC-RES for generously providing computational resources.

8 Competing Interests Statement

The authors declare no competing interests.

9 Author contributions

S.M., Conceptualization, Data Curation, Formal Analysis, Investigation, Methodology, Software, Visualization, Writing – Original Draft Preparation. O.L., Conceptualization, Formal Analysis, Investigation, Methodology, Software, Visualization, Writing – Original Draft Preparation. T.R., Data Curation, Software. S. P.-G., Conceptualization, Writing - Review & Editing. R.V.H., Conceptualization, Writing - Review & Editing. R.R.S., Software, Writing – Review & Editing. P.S, Software, Visualization, Writ-

ing – Review & Editing. R.G.M., Data Curation. A. A.-G., Funding Acquisition. N.L., Conceptualization, Funding Acquisition, Project Administration, Writing – Original Draft Preparation, Supervision.

References

1. Wen, M., Spotte-Smith, E. W. C., Blau, S. M., McDermott, M. J., Krishnapriyan, A. S. & Persson, K. A. Chemical reaction networks and opportunities for machine learning. *Nature Computational Science* **3,** 12–24 (2023).

2. Unsleber, J. P. & Reiher, M. The Exploration of Chemical Reaction Networks. *Annual Review of Physical Chemistry* **71,** 121–142 (2020).

3. Broadbelt, L. J., Stark, S. M. & Klein, M. T. Computer Generated Pyrolysis Modeling: On-the-Fly Generation of Species, Reactions, and Rates. *Industrial & Engineering Chemistry Research* **33,** 790–799 (1994).

4. Feinberg, M. *Foundations of Chemical Reaction Network Theory* (Springer Cham, 2019).

5. Garay-Ruiz, D. & Bo, C. Revisiting Catalytic Cycles: A Broader View through the Energy Span Model. *ACS Catalysis* **10,** 12627–12635 (2020).

6. Newman, M. E. J. The Structure and Function of Complex Networks. *SIAM Review* **45,** 167–256 (2003).

7. Tyson, J. J. & Novák, B. Functional Motifs in Biochemical Reaction Networks. *Annual Review of Physical Chemistry* **61,** 219–240 (2010).

8. Wong, A. S. Y. & Huck, W. T. S. Grip on complexity in chemical reaction networks. *Beilstein Journal of Organic Chemistry* **13,** 1486–1497 (2017).

9. Hashemi, A., Bougueroua, S., Gaigeot, M.-P. & Pidko, E. A. ReNeGate: A Reaction Network Graph-Theoretical Tool for Automated Mechanistic Studies in Computational Homogeneous Catalysis. *Journal of Chemical Theory and Computation* **18,** 7470–7482 (2022).

10. Steiner, M. & Reiher, M. Autonomous Reaction Network Exploration in Homogeneous and Heterogeneous Catalysis. *Topics in Catalysis* **65,** 6–39 (2022).

11. Stocker, S., Csányi, G., Reuter, K. & Margraf, J. T. Machine learning in chemical reaction space. *Nature Communications* **11,** 5505 (2020).

12. Unsleber, J. P., Grimmel, S. A. & Reiher, M. Chemoton 2.0: Autonomous exploration of chemical reaction networks. *Journal of Chemical Theory and Computation* **18,** 5393–5409 (2022).

13. Zhao, Q. & Savoie, B. M. Simultaneously improving reaction coverage and computational cost in automated reaction prediction tasks. *Nature Computational Science* **1,** 479–490 (2021).

14. Zhao, Q., Xu, Y., Greeley, J. & Savoie, B. M. Deep reaction network exploration at a heterogeneous catalytic interface. *Nature Communications* **13,** 4860 (2022).

15. Jafari, M. & Zimmerman, P. M. Uncovering reaction sequences on surfaces through graphical methods. *Physical Chemistry Chemical Physics* **20,** 7721–7729 (2018).

16. Maeda, S., Harabuchi, Y., Takagi, M., Saita, K., Suzuki, K., Ichino, T., Sumiya, Y., Sugiyama, K. & Ono, Y. Implementation and performance of the artificial force induced reaction method in the GRRM17 program. *Journal of Computational Chemistry* **39,** 233–251 (2018).

17. Chen, D., Shang, C. & Liu, Z.-P. Machine-learning atomic simulation for heterogeneous catalysis. *npj Computational Materials* **9,** 1–9 (2023).

18. Rangarajan, S., Bhan, A. & Daoutidis, P. Language-oriented rule-based reaction network generation and analysis: Description of RING. *Computers & Chemical Engineering* **45,** 114–123 (2012).

19. Benson, S. W. & Buss, J. H. Additivity Rules for the Estimation of Molecular Properties. Thermodynamic Properties. *The Journal of Chemical Physics* **29,** 546–572 (1958).

20. Sabbe, M. K., Saeys, M., Reyniers, M.-F., Marin, G. B., Speybroeck, V. V. & Waroquier, M. Group Additive Values for the Gas Phase Standard Enthalpy of Formation of Hydrocarbons and Hydrocarbon Radicals. *The Journal of Physical Chemistry A* **109,** 7466–7480 (2005).

21. Wittreich, G. R. & Vlachos, D. G. Python Group Additivity (pGrAdd) Software for Estimating Species Thermochemical Properties. *Computer Physics Communications* **273,** 108277 (2022).

22. Abild-Pedersen, F., Greeley, J., Studt, F., Rossmeisl, J., Munter, T. R., Moses, P. G., Skúlason, E., Bligaard, T. & Nørskov, J. K. Scaling Properties of Adsorption Energies for Hydrogen-Containing Molecules on Transition-Metal Surfaces. *Physical Review Letters* **99,** 016105 (2007).

23. Calle-Vallejo, F., Loffreda, D., Koper, M. T. M. & Sautet, P. Introducing structural sensitivity into adsorption–energy scaling relations by means of coordination numbers. *Nature Chemistry* **7,** 403–410 (2015).

24. Bronsted, J. N. Acid and Basic Catalysis. *Chemical Reviews* **5,** 231–338 (1928).

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

26. Nørskov, J. K., Bligaard, T., Logadottir, A., Bahn, S., Hansen, L. B., Bollinger, M., Bengaard, H., Hammer, B., Sljivancanin, Z., Mavrikakis, M., Xu, Y., Dahl, S. & Jacobsen, C. J. H. Universality in Heterogeneous Catalysis. *Journal of Catalysis* **209,** 275–278 (2002).

27. Liu, M., Grinberg Dana, A., Johnson, M. S., Goldman, M. J., Jocher, A., Payne, A. M., Grambow, C. A., Han, K., Yee, N. W., Mazeau, E. J., Blondal, K., West, R. H., Goldsmith, C. F. & Green, W. H. Reaction Mechanism Generator v3.0: Advances in Automatic Mechanism Generation. *Journal of Chemical Information and Modeling* **61,** 2686–2696 (2021).

28. Kreitz, B., Lott, P., Studt, F., Medford, A. J., Deutschmann, O. & Goldsmith, C. F. Automated Generation of Microkinetics for Hetero-

geneously Catalyzed Reactions Considering Correlated Uncertainties. *Angewandte Chemie International Edition* **62,** e202306514 (2023).

29. Kreitz, B., Lott, P., Bae, J., Blöndal, K., Angeli, S., Ulissi, Z. W., Studt, F., Goldsmith, C. F. & Deutschmann, O. Detailed Microkinetics for the Oxidation of Exhaust Gas Emissions through Automated Mechanism Generation. *ACS Catal.* **12.** Publisher: American Chemical Society, 11137–11151 (2022).

30. Chanussot, L., Das, A., Goyal, S., Lavril, T., Shuaibi, M., Riviere, M., Tran, K., Heras-Domingo, J., Ho, C., Hu, W., Palizhati, A., Sriram, A., Wood, B., Yoon, J., Parikh, D., Zitnick, C. L. & Ulissi, Z. Open Catalyst 2020 (OC20) Dataset and Community Challenges. *ACS Catalysis* **11,** 6059–6072 (2021).

31. Pablo-García, S., Morandi, S., Vargas-Hernández, R. A., Jorner, K., Ivković, Ž., López, N. & Aspuru-Guzik, A. Fast evaluation of the adsorption energy of organic molecules on metals via graph neural networks. *Nature Computational Science* **3,** 433–442 (2023).

32. Suvarna, M. & Pérez-Ramírez, J. Embracing data science in catalysis research. *Nature Catalysis,* 1–12 (2024).

33. Ulissi, Z. W., Medford, A. J., Bligaard, T. & Nørskov, J. K. To address surface reaction network complexity using scaling relations machine learning and DFT calculations. *Nature communications* **8,** 14621 (2017).

34. Göltl, F. & Mavrikakis, M. Generalized Brønsted-Evans-Polanyi Relationships for Reactions on Metal Surfaces from Machine Learning. *ChemCatChem* **14,** e202201108 (2022).

35. Hutton, D. J., Cordes, K. E., Michel, C. & Göltl, F. Machine learning-based prediction of activation energies for chemical reactions on metal surfaces. *Journal of Chemical Information and Modeling* **63,** 6006–6013 (2023).

36. Wander, B., Shuaibi, M., Kitchin, J. R., Ulissi, Z. W. & Zitnick, C. L. *CatTSunami: Accelerating Transition State Energy Calculations with Pre-trained Graph Neural Networks* 2024. arXiv: 2405.02078 [cond-mat.mtrl-sci].

37. Garrido Torres, J. A., Jennings, P. C., Hansen, M. H., Boes, J. R. & Bligaard, T. Low-Scaling Algorithm for Nudged Elastic Band Calculations Using a Surrogate Machine Learning Model. *Physical Review Letters* **122,** 156001 (2019).

38. Stocker, S., Jung, H., Csányi, G., Goldsmith, C. F., Reuter, K. & Margraf, J. T. Estimating Free Energy Barriers for Heterogeneous Catalytic Reactions with Machine Learning Potentials and Umbrella Integration. *Journal of Chemical Theory and Computation* (2023).

39. Schaaf, L. L., Fako, E., De, S., Schäfer, A. & Csányi, G. Accurate energy barriers for catalytic reaction pathways: an automatic training protocol for machine learning force fields. *npj Computational Materials* **9,** 180 (2023).

40. Wang, B., Gu, T., Lu, Y. & Yang, B. Prediction of energies for reaction intermediates and transition states on catalyst surfaces using graph-based machine learning models. *Molecular Catalysis* **498,** 111266 (2020).

41. Kim, S., Woo, J. & Kim, W. Y. Diffusion-based generative AI for exploring transition states from 2D molecular graphs. *Nature Communications* **15,** 341 (2024).

42. Landrum, G. *et al. rdkit/rdkit: 2023_09_3 (Q3 2023) Release* version Release_2023_09_3. 2023.

43. Weininger, D. Smiles, a Chemical Language and Information System. 1. Introduction To Methodology and Encoding Rules. *Journal of Chemical Information and Modeling* **28,** 31–36 (1988).

44. Liao, Y.-L., Wood, B., Das, A. & Smidt, T. *EquiformerV2: Improved Equivariant Transformer for Scaling to Higher-Degree Representations* 2024.

45. Du, J., Zhang, S., Wu, G., Moura, J. M. F. & Kar, S. *Topology Adaptive Graph Convolutional Networks* 2018.

46. Martí, C., Blanck, S., Staub, R., Loehlé, S., Michel, C. & Steinmann, S. N. DockOnSurf: A Python Code for the High-Throughput Screening of Flexible Molecules Adsorbed on Surfaces. *Journal of Chemical Information and Modeling* **61,** 3386–3396 (2021).

47. Han, S., Lysgaard, S., Vegge, T. & Hansen, H. A. Rapid mapping of alloy surface phase diagrams via Bayesian evolutionary multitasking. *npj Computational Materials* **9,** 139 (2023).

48. Hirschfeld, L., Swanson, K., Yang, K., Barzilay, R. & Coley, C. W. Uncertainty Quantification Using Neural Networks for Molecular Property Prediction. *Journal of Chemical Information and Modeling* **60,** 3770–3780 (2020).

49. Nørskov, J. K., Rossmeisl, J., Logadottir, A., Lindqvist, L., Kitchin, J. R., Bligaard, T. & Jónsson, H. Origin of the Overpotential for Oxygen Reduction at a Fuel-Cell Cathode. *The Journal of Physical Chemistry B* **108,** 17886–17892 (2004).

50. Motagamwala, A. H. & Dumesic, J. A. Microkinetic Modeling: A Tool for Rational Catalyst Design. *Chemical Reviews* **121,** 1049–1076 (2021).

51. Chorkendorff, I. & Niemantsverdriet, J. *Concepts of modern catalysis and kinetics* (Wiley-VCH Verlag, 2003).

52. Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17,** 261–272 (2020).

53. Rackauckas, C. & Nie, Q. DifferentialEquations.jl – A Performant and Feature-Rich Ecosystem for Solving Differential Equations in Julia. *The Journal of Open Research Software* **5** (2017).

54. Utkarsh, U., Churavy, V., Ma, Y., Besard, T., Srisuma, P., Gymnich, T., Gerlach, A. R., Edelman, A., Barbastathis, G., Braatz, R. D. & Rackauckas, C. Automated translation and accelerated solving of differential equations on multiple GPU platforms. *Computer Methods in Applied Mechanics and Engineering* **419,** 116591 (2024).

55. Burte, A., Page, B., Nair, A., Grabow, L., Dauenhauer, P. & Scott, S. CatTestHub: A Benchmarking Database of Experimental Heterogeneous Catalysis for Evaluating Advanced Materials. *ChemRxiv* (2024).

56. Pablo-García, S., Veenstra, F. L. P., Ting, L. R. L., García-Muelas, R., Dattila, F., Martín, A. J., Yeo, B. S., Pérez-Ramírez, J. & López, N. Mechanistic Routes Toward C_3 Products in Copper-Catalysed CO_2 Electroreduction. *Catalysis Science & Technology* **12,** 409–417 (2022).

57. Olah, G. A. Beyond Oil and Gas: The Methanol Economy. *Angewandte Chemie International Edition* **44,** 2636–2639 (2005).

58. Deka, T. J., Osman, A. I., Baruah, D. C. & Rooney, D. W. Methanol fuel production, utilization, and techno-economy: a review. *Environmental Chemistry Letters* **20,** 3525–3554 (2022).

59. García-Muelas, R., Li, Q. & López, N. Density Functional Theory Comparison of Methanol Decomposition and Reverse Reactions on Metal Surfaces. *ACS Catalysis* **5,** 1027–1036 (2015).

60. Tran, R., Xu, Z., Radhakrishnan, B., Winston, D., Sun, W., Persson, K. A. & Ong, S. P. Surface energies of elemental crystals. *Scientific Data* **3,** 160080 (2016).

61. Honkala, K., Hellman, A., Remediakis, I. N., Logadottir, A., Carlsson, A., Dahl, S., Christensen, C. H. & Nørskov, J. K. Ammonia Synthesis from First-Principles Calculations. *Science* **307,** 555–558 (2005).

62. Mahmoudi, H., Mahmoudi, M., Doustdar, O., Jahangiri, H., Tsolakis, A., Gu, S. & LechWyszynski, M. A review of Fischer Tropsch synthe-

sis process, mechanism, surface chemistry and catalyst formulation. *Biofuels Engineering* **2,** 11–31 (2017).

63. Martín, A. J., Mitchell, S., Mondelli, C., Jaydev, S. & Pérez-Ramírez, J. Unifying views on catalyst deactivation. *Nature Catalysis* **5,** 854–866 (2022).

64. Deshpande, S., Maxson, T. & Greeley, J. Graph theory approach to determine configurations of multidentate and high coverage adsorbates for heterogeneous catalysis. *npj Computational Materials* **6** (2020).

65. Ramakrishnan, R., Dral, P. O., Rupp, M. & von Lilienfeld, O. A. Big Data Meets Quantum Chemistry Approximations: The δ-Machine Learning Approach. *Journal of Chemical Theory and Computation* **11,** 2087–2096 (2015).

66. Kresse, G. & Furthmüller, J. Efficiency of Ab-Initio Total Energy Calculations for Metals and Semiconductors Using a Plane-Wave Basis Set. *Computational Materials Science* **6,** 15–50 (1996).

67. Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Physical Review Letters* **77,** 3865–3868 (1996).

68. Grimme, S., Ehrlich, S. & Goerigk, L. Effect of the Damping Function in Dispersion Corrected Density Functional Theory. *Journal of Computational Chemistry* **32,** 1456–1465 (2011).

69. Almora-Barrios, N., Carchini, G., Błoński, P. & López, N. Costless Derivation of Dispersion Coefficients for Metal Surfaces. *Journal of Chemical Theory and Computation* **10,** 5002–5009 (2014).

70. Blöchl, P. E. Projector Augmented-Wave Method. *Physical Review B* **50,** 17953–17979 (1994).

71. Neugebauer, J. & Scheffler, M. Adsorbate-Substrate and Adsorbate-Adsorbate Interactions of Na and K Adlayers on Al(111). *Physical Review B* **46,** 16067–16080 (1992).

72. Monkhorst, H. J. & Pack, J. D. Special Points for Brillouin-Zone Integrations. *Physical Review B* **13,** 5188–5192 (1976).

73. Heyden, A., Bell, A. T. & Keil, F. J. Efficient methods for finding transition states in chemical reactions: Comparison of improved dimer method and partitioned rational function optimization method. *The Journal of Chemical Physics* **123,** 224101 (2005).

74. Larsen, A. H. *et al.* The atomic simulation environment—a Python library for working with atoms. *Journal of Physics: Condensed Matter* **29,** 273002 (2017).

75. NetworkX developer team. *NetworkX* 2014.

76. Fey, M. & Lenssen, J. E. *Fast Graph Representation Learning with PyTorch Geometric* 2019.

77. Paszke, A. *et al.* in *Advances in Neural Information Processing Systems 32* 8024–8035 (Curran Associates, Inc., 2019).

78. Hunter, J. D. Matplotlib: a 2d Graphics Environment. *Computing in Science & Engineering* **9,** 90–95 (2007).

79. Waskom, M. L. seaborn: statistical data visualization. *Journal of Open Source Software* **6,** 3021 (2021).