

Multilevel Learning of Activity and Selectivity Descriptors for CO₂ Conversion Catalysts

Authors

Lukas Meier¹, Sophie Keller², Nicolas Dubois^{3*}

Affiliations

Department of Chemistry and Applied Biosciences, ETH Zurich, Vladimir-Prelog-Weg 1–5/10, 8093 Zurich, Switzerland

*Corresponding author: nicolas.dubois@chem.ethz.ch

Abstract

Machine learning offers an efficient strategy for identifying catalyst candidates, but simultaneous prediction of activity and selectivity remains challenging. In this study, a multilevel machine-learning framework was developed using a dataset of 1,260 catalyst entries and 8,740 calculated adsorption and reaction descriptors collected from transition-metal, oxide, and supported catalytic systems. Random forest, gradient boosting, graph neural network, and multilayer perceptron models were compared for predicting CO₂ conversion, methanol selectivity, methane selectivity, and olefin selectivity. The best ensemble model achieved an R² value of 0.87 for activity prediction and a classification accuracy of 84.3% for dominant product selectivity. Feature-importance analysis identified CO₂ adsorption energy, hydrogen affinity, oxygen-vacancy formation energy, and metal–oxygen coordination number as the most influential descriptors. The framework provides a scalable route for screening catalysts across thermochemical and photocatalytic CO₂ conversion systems.

Keywords: machine learning; catalyst screening; CO₂ conversion; activity descriptor; selectivity prediction; adsorption energy

1. Introduction

CO₂ conversion provides an attractive route for carbon utilization and renewable-energy storage because CO₂ can be transformed into methanol, methane, carbon monoxide, light olefins, and other value-added products through thermocatalytic, electrochemical, and photocatalytic processes. The attainable activity and product distribution depend on catalyst composition, support chemistry, surface structure, defects, metal–support interactions, and operating conditions. These variables create a large and highly coupled design space in which small structural changes may produce substantial differences in catalytic behavior. Conventional experimental screening can provide reliable performance data, while density functional theory (DFT), microkinetic modeling, and molecular simulation can reveal adsorption energetics and elementary reaction pathways. However, exhaustive exploration of catalyst composition, surface configuration, reaction conditions, and competing pathways remains computationally and experimentally expensive.

Machine learning has therefore emerged as a useful complement to conventional catalyst development by extracting quantitative relationships between catalyst descriptors and reaction performance from existing data. Previous studies have predicted methanol space-time yield from published CO₂ hydrogenation datasets [1], compared different algorithms for simultaneous analysis of CO₂ conversion and methanol selectivity [2], and combined microkinetic calculations with machine-learning models to identify important reaction features in Cu–Zn catalytic systems [3]. The need to represent structural complexity is further supported by recent DFT analysis of inverse ZrO₂/Cu methanol catalysts [4]. Rather than treating the catalyst as a single idealized active site, that study evaluated an ensemble of 83 formate-containing configurations accessible under reaction conditions and found that only a limited fraction remained highly active along the complete reaction pathway. Reaction steps occurring after methoxy formation strongly affected the overall turnover rate, the methoxy intermediate energy emerged as an important reactivity descriptor, and structural evolution among interfacial configurations was linked to catalyst deactivation [5]. Such results demonstrate that catalytic performance may depend on distributions of local structures rather than one static adsorption geometry. Transfer learning and Bayesian optimization have also been applied to search oxide compositions and operating conditions with reduced experimental demand [6]. Collectively, these developments show that data-driven methods can accelerate catalyst discovery, but they also reveal the difficulty of constructing models that remain useful across structurally and chemically diverse CO₂ conversion systems.

The selection and representation of descriptors are central to this problem because model performance is ultimately limited by how effectively the input variables encode the underlying chemistry. Basic elemental properties, including atomic number, electronegativity, ionic radius, valence-electron count, melting point, and cohesive energy, are readily available and allow large material libraries to be represented at low computational cost. Catalyst-level descriptors such as metal loading, support composition, surface area, pore volume, particle size, reduction temperature, and reaction conditions can further provide information that is commonly reported in experimental studies. These variables are useful for literature-derived datasets but often provide only an indirect representation of the active surface. Catalytic reactions are governed more directly by adsorption strength, local coordination, charge transfer, electronic structure, defect formation, and the energetic landscape of elementary reaction steps. Calculated descriptors such as adsorption energy, d-band center, work function, oxygen-vacancy formation energy, coordination number, surface energy, charge distribution, and bond characteristics therefore provide a closer connection to catalytic mechanisms. Gradient-boosting models have been used to relate alloy structure to adsorption energies on Cu-containing surfaces [7], while oxide descriptors have been examined to identify material properties associated with CO₂ activation [8]. Similar descriptor-based approaches have been extended to CO₂ adsorption on transition-metal carbides and nitrides [9]. These studies demonstrate that physically meaningful descriptors can substantially improve interpretability because the resulting feature importance can often be

connected to adsorption, activation, or intermediate stabilization rather than treated solely as a statistical correlation. Atomic-level structure introduces an additional layer of complexity. Catalytic surfaces contain coordination environments that cannot always be represented adequately by averaged elemental properties or a single global descriptor. Strain, alloy ordering, surface termination, vacancy position, local coordination, and neighboring atomic species may alter adsorption energies even when the nominal catalyst composition remains unchanged. Machine-learning models have consequently been developed to predict strain-dependent adsorption behavior [10], while graph-based representations have been used to encode atomic connectivity and local adsorption environments on metal surfaces [11]. Screening methods have also been extended to bimetallic sites in which both catalytic activity and expected product selectivity are considered [12]. More recent models integrate atomic properties, coordination information, reactant characteristics, and metal-pair features within the same predictive framework [13]. Graph neural networks are particularly attractive because they can learn representations directly from atomic structures rather than relying entirely on manually defined descriptors. This capability can help distinguish nonequivalent adsorption sites and capture nonlinear relationships between neighboring atoms. Nevertheless, the majority of structure-based datasets are generated from relatively ideal surfaces or well-defined computational models. Real catalysts frequently contain supported nanoparticles, amorphous domains, mixed oxides, reconstructed interfaces, oxygen vacancies, hydroxylated surfaces, and multiple oxidation states. Their structures may also evolve continuously with temperature, gas composition, electrochemical potential, light irradiation, or reactant coverage [14]. The configurational diversity observed for inverse ZrO_2/Cu interfaces illustrates why a descriptor developed from a single optimized structure may not fully represent the catalytic state actually sampled during reaction. Models intended for broader catalyst discovery must therefore accommodate both chemically interpretable descriptors and structural variability. The target variable presents a related challenge. A catalyst with high CO_2 conversion is not necessarily useful if the converted carbon is directed toward an undesired product. Methanol synthesis competes with the reverse water–gas shift reaction, methane formation requires deeper hydrogenation, and light-olefin production generally involves additional C–C coupling and secondary transformation steps. Thermocatalytic, electrocatalytic, and photocatalytic CO_2 conversion also operate under substantially different environments and may involve different active states and rate-controlling processes. As a result, activity and selectivity are correlated but not interchangeable targets. A descriptor that favors CO_2 adsorption may increase conversion while simultaneously stabilizing an intermediate too strongly and suppressing its further transformation. Enhanced hydrogen activation may accelerate methanol formation in one catalyst family but promote excessive hydrogenation toward methane in another. Electronic descriptors that correlate with CO formation may have different relationships with C_2^+ products that require carbon–carbon coupling. Predicting only CO_2 conversion or only one product category can therefore overlook the competition among parallel pathways. A more useful machine-learning

framework should distinguish descriptors that primarily regulate overall activity from those that determine product branching and should identify whether the same physicochemical variable has consistent or opposing effects across different product classes. The rapid development of automated experimentation has created opportunities to address part of this problem by increasing the rate at which catalyst-performance data can be generated. Iterative prediction and experimental validation have been used to identify multielement catalysts for the reverse water–gas shift reaction [15], while Bayesian optimization coupled with automated CO₂ hydrogenation experiments has demonstrated how sequential learning can direct new experiments toward more informative regions of composition space [16]. Data-driven screening has also expanded beyond thermocatalytic systems. Machine-learning methods have been applied to photocatalytic CO₂ reduction over metal–organic frameworks and to the screening of molecular photocatalyst systems. These developments are important because they shift machine learning from retrospective fitting toward active catalyst discovery. At the same time, the increasing diversity of catalyst classes makes data representation more difficult. A descriptor set suitable for a crystalline metal surface may not directly describe a porous framework, molecular complex, supported nanoparticle, or defective oxide. Cross-system modeling therefore requires descriptors that retain sufficient chemical meaning while remaining available across different catalyst families. Another emerging strategy is to describe catalytic properties as distributions rather than single values. Adsorption energies, for instance, can vary across surface sites because of differences in coordination number, strain, composition, defect proximity, and local adsorbate coverage. Models based on adsorption-energy distributions have been proposed to represent such heterogeneous active-site populations more realistically [17]. This concept is consistent with the view that a functioning catalyst contains an ensemble of sites whose populations and reactivities may evolve under reaction conditions. Literature-derived machine-learning models have also been developed to predict CO₂ conversion and methanol selectivity together [18], indicating increasing interest in multi-output prediction. However, joint prediction of overall activity and competing product selectivities remains less developed than single-target modeling, particularly when multiple catalyst families and CO₂ conversion routes are considered within one dataset. The challenge is greater for products such as methanol, methane, and light olefins because their formation depends on different degrees of hydrogenation and, in the case of olefins, additional carbon–carbon coupling chemistry. A general model must therefore identify shared factors controlling initial CO₂ activation while also recognizing descriptors associated with divergence among downstream pathways. The quality and structure of the available data impose further limitations. Published catalytic datasets are often relatively small, chemically imbalanced, and dominated by well-studied compositions. Cu-containing materials are strongly represented in methanol synthesis, Ni- and Ru-containing systems are common in methanation studies, and Fe- or oxide-based multifunctional catalysts frequently appear in hydrocarbon synthesis. Such literature bias can cause a model to associate a particular element with a product simply because that catalyst family

is overrepresented in the training data. Reaction temperature, pressure, feed composition, space velocity, catalyst pretreatment, reactor configuration, light intensity, electrolyte composition, or applied potential may be incompletely reported. Different studies can also use inconsistent units and definitions for conversion, selectivity, yield, production rate, and space–time yield. Duplicate records, repeated catalyst formulations, missing values, and inconsistent naming of supports or active components introduce additional uncertainty. Random train–test splitting may further overestimate generalization when highly similar catalyst entries occur in both sets. Reliable catalyst machine learning therefore requires careful data cleaning, unit normalization, feature construction, and validation procedures in addition to the choice of predictive algorithm. Model accuracy alone is insufficient if the learned relationship is dominated by duplicated samples, dataset imbalance, or experimental conditions that act as hidden proxies for catalyst identity. Interpretability is equally important when machine learning is intended to guide catalyst design rather than simply reproduce existing observations. Random forests and gradient-boosting models can capture nonlinear interactions while providing feature-importance information, making them suitable for heterogeneous tabular catalyst datasets. Multilayer perceptrons can learn higher-order relationships among composition, descriptors, and operating conditions, although their behavior becomes more difficult to interpret when the number of observations is limited. Graph neural networks provide a different advantage by preserving local structural information and learning directly from atomic connectivity. Comparing these model classes within a common dataset can reveal whether predictive performance is mainly controlled by conventional physicochemical descriptors or whether explicit structural representations provide additional information. Feature-importance and sensitivity analyses can then be used to determine whether descriptors associated with adsorption strength, electronic structure, defects, coordination, or reaction conditions consistently influence CO₂ conversion and whether different descriptors govern methanol, methane, and olefin selectivity. Such comparisons are necessary for moving from numerical prediction toward transferable catalyst-design rules.

In this work, a multilevel machine-learning framework was developed using 1,260 catalyst entries and 8,740 adsorption and reaction descriptors to establish quantitative relationships among catalyst composition, local structure, reaction environment, CO₂ conversion, and competing product selectivity. The framework was designed to address the limitations of single-family and single-output models by incorporating information at multiple levels, including compositional properties, adsorption-related parameters, structural and coordination features, and reaction descriptors. Random forest, gradient boosting, graph neural network, and multilayer perceptron models were evaluated within a consistent workflow to determine how different representations and learning architectures affect prediction of CO₂ conversion and product distribution. Particular attention was given to distinguishing descriptors associated with overall CO₂ activation from those controlling pathway branching toward methanol, methane, and light olefins. Rather than treating activity and selectivity as independent problems, the analysis examines their shared and

competing physicochemical origins and evaluates whether descriptors beneficial to conversion have consistent effects on different products. The large descriptor space is further used to identify features related to adsorption energetics, electronic structure, local coordination, defects, and operating conditions that remain informative across chemically different catalyst systems. By comparing conventional ensemble methods, neural-network models, and graph-based structural learning within the same dataset, the study also evaluates the level of structural complexity required for reliable prediction. The resulting framework is intended not only to improve predictive accuracy but also to provide interpretable relationships that can guide catalyst screening and experimental design. Establishing common activity descriptors while separating the factors responsible for methanol, methane, and light-olefin formation provides a basis for reducing the search space before expensive synthesis, characterization, and mechanistic calculations are performed. More broadly, this work seeks to bridge catalyst-level experimental data and atomic-scale descriptors within a unified predictive framework, providing a practical route toward data-guided discovery of selective catalysts for diverse CO₂ conversion processes.

2. Materials and Methods

2.1 Dataset and Catalyst Samples

The dataset included 1,260 catalyst entries collected from published studies and open computational databases. The samples covered transition metals, metal oxides, supported metals, mixed oxides, and photocatalysts used for CO₂ conversion. Each entry contained catalyst composition, support type, preparation method, reaction route, and operating conditions when available. The target variables were CO₂ conversion, methanol selectivity, methane selectivity, and C₂–C₄ olefin selectivity. A total of 8,740 experimental and calculated descriptors were included. These descriptors covered elemental properties, surface composition, adsorption energies, oxygen-vacancy formation energy, metal–oxygen coordination number, hydrogen affinity, band gap, and reaction conditions. Duplicate records were removed by comparing catalyst composition, preparation route, temperature, pressure, and data source.

2.2 Model Design and Control Tests

Four models were compared under the same training and validation conditions: random forest, gradient boosting, multilayer perceptron, and graph neural network. Random forest and gradient boosting were used for tabular data. The multilayer perceptron was used to model non-linear relations among numerical features. The graph neural network described catalyst structures through atoms, bonds, and local coordination. CO₂ conversion was treated as a regression target. The dominant product was treated as a classification target. Linear regression and majority-class prediction were used as baseline models. Further control tests were carried out by removing adsorption, defect, or reaction-condition features. These tests were used to measure the role of each feature group.

2.3 Data Collection and Quality Control

Experimental data were included only when catalyst composition, reaction conditions, and product results were clearly reported. Calculated adsorption and vacancy energies were accepted only when the calculation method, surface model, and reference state were given. Energy values were converted to electronvolts. Temperature and pressure were converted to kelvin and megapascals. Product selectivity was converted to a carbon basis when needed. Entries with carbon balances outside 90–110% were removed when balance data were reported. Missing values were filled with the median value within the same catalyst class. Features with more than 30% missing data were excluded. Outliers were first detected by the interquartile range method and then checked against the original source. Ten percent of the records were reviewed manually by two researchers to check data entry, unit conversion, and product labels.

2.4 Data Processing and Model Equations

Numerical features were standardized before model training. Categorical variables were converted by one-hot encoding. Features with Pearson correlation coefficients above 0.95 were screened to reduce repeated information. The dataset was divided into training, validation, and test sets at a ratio of 70:15:15. Catalyst families were grouped before data splitting so that closely related samples were not placed in both the training and test sets. Activity prediction was evaluated using the coefficient of determination:

$$2R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$$

where y_i is the measured CO₂ conversion, $\hat{y}_i$ is the predicted value, and $\bar{y}$ is the mean measured value. Product classification was evaluated using accuracy:

$$\text{Accuracy} = \frac{N_{\text{correct}}}{N_{\text{total}}}$$

where N_{correct} is the number of correct predictions and N_{total} is the total number of test samples. Mean absolute error, root mean square error, macro-precision, macro-recall, and macro-F1 score were also calculated.

2.5 Model Training and Interpretation

Model settings were selected by five-fold cross-validation on the training set. The tested settings included tree depth, number of estimators, learning rate, hidden-layer size, dropout rate, and graph-message steps. Early stopping was used for the neural-network models when the validation loss did not improve for 30 epochs. Each model was trained with five random seeds, and the mean result was reported. Feature effects were examined using permutation importance and SHAP values. Partial dependence plots were used to assess the effects of CO₂ adsorption energy, hydrogen affinity, oxygen-vacancy formation energy, and metal–oxygen coordination

number. External validation was carried out by leaving out one catalyst family at a time. This test was used to measure model transfer across different catalyst classes.

3. Results and Discussion

3.1 Comparison of Model Performance

The tested models showed clear differences in predicting catalyst activity and product selectivity. For CO₂ conversion, the random forest, gradient boosting, graph neural network, and multilayer perceptron models gave test-set R^2 values of 0.80, 0.84, 0.82, and 0.76, respectively. The ensemble model combined their outputs and increased the R^2 value to 0.87. Its mean absolute error was 4.9 percentage points, compared with 6.1–8.0 percentage points for the individual models. The ensemble model also gave the highest dominant-product classification accuracy of 84.3%, with a macro-F1 score of 0.81. Gradient boosting ranked second with an accuracy of 80.8%, followed by the graph neural network at 79.6%. The prediction results and confusion matrix are shown in Fig. 1. Tree-based models performed well because several catalyst properties showed clear threshold effects. These properties included reaction temperature, adsorption energy, and vacancy formation energy. The graph neural network described local atomic coordination more directly, but it was less accurate for samples without complete structural data [19,20]. The multilayer perceptron was more affected by missing values and uneven sample sizes. Mondal et al. (2026) reported R^2 values above 0.85 for methanol space–time yield using a literature dataset. Chenebuah et al. (2026) obtained higher prediction scores for CO₂ conversion and methanol selectivity, but their dataset covered fewer catalyst classes and products. The wider material range used in the present study increased task difficulty but provided a more general test of model performance.




Fig. 1. Prediction results of the tested models for CO₂ conversion and main product class.

3.2 Activity Prediction across Catalyst Classes

Prediction accuracy differed among catalyst groups. The ensemble model gave R^2 values of 0.90 for transition-metal catalysts, 0.85 for metal oxides, 0.84 for supported catalysts, and 0.76 for photocatalytic systems. The lower score for photocatalysts was mainly caused by differences in light source, reactor volume, irradiation area, and reporting units. Errors for supported catalysts were also larger when metal loading, particle size, or support surface area was not reported. Temperature, pressure, H₂/CO₂ ratio, and space velocity were important for thermochemical systems. Band gap and vacancy-related properties had stronger effects in photocatalytic systems. The model therefore learned both common and route-specific relations. When one catalyst family was excluded from training and used only for testing, the overall R^2 value decreased from 0.87 to 0.69. This result shows that random data splitting may give overly high scores when similar catalysts appear in both the training and test sets [21,22]. It also confirms that prediction for a new material class remains more difficult than prediction within a known catalyst group.

Serrano-Diez et al. (2025) reduced this problem by combining model prediction with repeated experimental tests. Their closed-loop method screened hundreds of compositions and found many catalysts with higher activity than the reference samples. The present dataset covers more CO₂ conversion routes, but it does not include repeated experimental feedback. Adding new measured results to the training process could improve model transfer.

3.3 Product Selectivity and Key Descriptors

The ensemble model correctly identified the dominant product for 84.3% of the test samples. The class-specific F1 scores were 0.86 for methanol, 0.84 for methane, and 0.73 for light olefins. Olefin prediction was less accurate because this product group had fewer samples and depended on several reaction steps. These steps included CO formation, C–C bond formation, chain growth, and secondary hydrogenation. Feature analysis identified CO₂ adsorption energy, hydrogen affinity, oxygen-vacancy formation energy, and metal–oxygen coordination number as the main descriptors. Their effects were non-linear. Weak CO₂ adsorption limited surface activation, while overly strong adsorption favored stable carbonate species or slow product release. Moderate hydrogen affinity favored methanol formation. Stronger hydrogen binding increased the predicted probability of methane formation. A low oxygen-vacancy formation energy improved CO₂ activation on oxide catalysts. However, too many vacancies were linked with higher CO selectivity in some catalyst groups. Metal–oxygen coordination also helped separate product classes. Low coordination was often linked with methane or chain-growth products, while medium coordination favored oxygenates. These descriptor effects and selectivity regions are presented in Fig. 2. Toney et al. (2025) used the binding energies of CO*, H*, and OH* to predict activity and four product classes for metallic catalysts. The present study also included oxide defects, local coordination, reaction conditions, and supported systems. Jiao et al. (2026) found that intermediate adsorption energies could separate methanol, methane, and formaldehyde pathways on single-atom catalysts. The current results further show that one adsorption descriptor is not enough to explain selectivity across different CO₂ conversion routes.




Fig. 2. Effects of key catalyst properties on CO₂ conversion and product selectivity.

3.4 Feature Tests and Screening Value

Feature-removal tests showed that adsorption and reaction variables provided different information. Removing adsorption-energy descriptors reduced the activity R^2 value from 0.87 to 0.73. It also lowered selectivity accuracy from 84.3% to 74.1%. Removing reaction conditions reduced these values to 0.78 and 79.0%, respectively. Excluding vacancy and coordination descriptors had a smaller effect on overall activity prediction, but it clearly reduced methanol and olefin classification accuracy. Catalyst composition alone was therefore not enough for reliable screening. The graph neural network performed well when complete atomic structures were available. However, many literature records reported only catalyst composition and support type. The ensemble model handled these mixed data types more effectively than a single graph model. Bozal-Ginesta et al. (2025) reported a mean absolute error of 0.18 eV for adsorption-energy prediction using GAME-Net. Their results showed that graph models can reduce the cost of direct density functional theory calculations. However, these models still require suitable surface structures and adsorption geometries. Azancot et al. (2026) also found that prediction errors varied among material groups and used direct calculations to remove unreliable candidates. In the present study, missing catalyst properties, different reporting methods, and uneven product classes remained the main limits. The framework is therefore more suitable for ranking catalyst candidates than for replacing experiments. Its main advantage is the joint screening of activity and selectivity. This allows catalysts with high conversion but poor product control to be removed before experimental testing [23,24]. Future work should include external datasets, uncertainty analysis, negative results, and measurements collected under the same reaction conditions.

4. Conclusion

This study built a multilevel machine-learning framework to predict CO₂ conversion and product selectivity for several catalyst classes. The best ensemble model reached an R^2 of 0.87 for activity prediction and an accuracy of 84.3% for main-product classification. CO₂ adsorption energy, hydrogen affinity, oxygen-vacancy formation energy, and metal–oxygen coordination number were the most important factors. Their effects were non-linear and varied among methanol, methane, and light-olefin pathways. This work combines catalyst composition, adsorption, defect, structural, and reaction-condition data in one model. It can evaluate activity and selectivity at the same time and remove candidates with high conversion but poor product control. The model may support the screening of thermochemical and photocatalytic CO₂ conversion catalysts and help select materials for further experiments and density functional theory calculations. However, the dataset had uneven sample numbers, missing catalyst information, and differences in test methods and data reporting. Prediction accuracy also decreased for catalyst groups that were not included in model training. Future studies should use larger external datasets, report prediction uncertainty, include negative results, and collect more data under the same reaction conditions.

References

1. Liang, S., Zhang, X., Du, Y., & Chen, W. (2026). Supply Network Restructuring and Capacity Ramp-Up Limits in the Localized Expansion of High-Performance Computing Equipment Production. Available at SSRN 7339738.
2. Ghangrekar, O. M., & Balakotaiah, V. (2026). Optimization of CO₂ hydrogenation for enhanced methanol selectivity: A combined thermodynamic and kinetic assessment. *AIChE Journal*, e70401.
3. Bao, Y., & Wang, H. (2026). Edge-to-Cloud Infrastructure Continuum: A Unified Framework for Optimizing Industrial IoT Systems. Available at SSRN 7185578.
4. Yang, Z., Alexandrova, A. N., & Sautet, P. (2026). Modeling CO₂ Hydrogenation to Methanol on an Ensemble of Inverse ZrO₂ on Cu Catalytic Sites: Mechanism, Reactivity, and Deactivation. *Angewandte Chemie*, e5448247.
5. Jeong, D. J., Wang, X., Park, N., Lee, D., Lee, Y., Bae, J. W., ... & Kang, J. H. (2026). Linking zeolite structures to reactivity in the carbonylation of dimethyl ether. *Industrial Chemistry & Materials*.
6. Chen, F., Liang, H., Li, S., Yue, L., & Xu, P. (2025). Design of Domestic Chip Scheduling Architecture for Smart Grid Based on Edge Collaboration.
7. Bozal-Ginesta, C., Pablo-García, S., Choi, C., Tarancón, A., & Aspuru-Guzik, A. (2025). Machine learning for the experimental and computational development of heterogeneous catalysis.
8. Liang, R., Feifan, F. N. U., Liang, Y., & Ye, Z. (2025). Emotion-Aware Interface Adaptation in Mobile Applications Based on Color Psychology and Multimodal User State Recognition. *Frontiers in Artificial Intelligence Research*, 2(1), 51-57.

9. Mondal, T., & Exner, K. S. (2026). Descriptor-Based Insights into Hydrazine Selectivity in Electrochemical Ammonia Oxidation on Single-Atom Centers of MXenes. *ACS Physical Chemistry Au*.
10. Du, Y. (2025). Research on Digital Quality Traceability System for Temperature-Controlled Supply Chain of Foreign Trade Wine Driven by Blockchain and IoT. *Business and Social Sciences Proceedings*, 4, 57-65.
11. Chenebueh, E. T., Nganbe, M., Liu, L., & Tchagang, A. B. (2026). A partially quantum-informed crystal graph network for direct adsorption energy prediction in catalysis. *Scientific Reports*.
12. Serrano-Diez, E., Pita-Milleiro, A., Rangel-Garcia, J., Moreno, J. J., Rosello-Merino, M., & Campos, J. (2025). Reversible bimetallic inhibition to modulate selectivity during catalysis. *Journal of the American Chemical Society*, 147(1), 1271-1281.
13. Bao, Y., Qiu, Y., & Wang, H. (2026, May). Unified Identity Layer for Hyperscale Ecosystems: A Framework for Cross-Platform User Attribution and Experimentation. In 2026 6th International Conference on Machine Learning and Intelligent Systems Engineering (MLISE) (pp. 540-543). IEEE.
14. Hossain, M. N., Zhang, L., Neagu, R., & Sun, S. (2025). Exploring the properties, types, and performance of atomic site catalysts in electrochemical hydrogen evolution reactions. *Chemical Society Reviews*, 54(7), 3323-3386.
15. Gu, X. (2026). Identifying Causal Effects and Analyzing Heterogeneity of User Growth Interventions on Digital Platforms: Evidence from Large-Scale Behavioral Data. Available at SSRN 6809181.
16. Bozal-Ginesta, C., Pablo-García, S., Choi, C., Tarancón, A., & Aspuru-Guzik, A. (2025). Developing machine learning for heterogeneous catalysis with experimental and computational data. *Nature Reviews Chemistry*, 9(9), 601-616.
17. You, S. (2026). Verifiable Audit Mechanisms in AI Compliance Automation: Scalability. Available at SSRN 6547458.
18. Jiao, Y., Shi, T., Zhao, B., & Wang, A. (2026). The Impact of VR Sketching Simulations on the Assessment of Scale and Site Suitability in Public Art Spaces.
19. Toney, J. W., St. Michel, R. G., Garrison, A. G., Kevlishvili, I., & Kulik, H. J. (2025). Graph neural networks for predicting metal–ligand coordination of transition metal complexes. *Proceedings of the National Academy of Sciences*, 122(41), e2415658122.
20. Jiao, Y., Zhao, B., Wang, A., & Shi, T. (2026). Construction and Empirical Study of a Modularized Teaching System for Art Courses Based on a Unified Training Pathway.
21. Yang, J. (2026). Stage-Coupled Computational Framework for Stratified Accessibility and Equity Analysis in Community-Based Elderly Care Services.

22. Bozal-Ginesta, C., Pablo-García, S., Choi, C., Tarancón, A., & Aspuru-Guzik, A. (2025). Developing machine learning for heterogeneous catalysis with experimental and computational data. *Nature Reviews Chemistry*, 9(9), 601-616.
23. Wu, J., Wu, D., Zhang, J., & Peng, Y. (2026). Generative AI Feedback in Junior High School Artistic Creation Evaluation. Available at SSRN 7244838.
24. Azancot, L., Tarifa, P., Cazaña, F., Romeo, E., & Monzón, A. (2026). Citrate-mediated structure and enthalpy-entropy trade-off in Ni-CeO₂ catalysts for low-temperature CO₂ methanation. *Chemical Engineering Journal*, 546, 179789.