Science for Energy and Environment
https://www.sciltp.com/journals/see

Review

Check for updates

Artificial Intelligence in Energy Catalysis: From Catalyst Screening to Decision-Making

Qixin Zhou [1], Yujie Mao [2,3], Wenjie Guo [4], Haifeng Shen [5], Haoying Wang [4], Yan Guo [2,*], and Yongfa Zhu [4,*]

[1] Department of Materials Science and Engineering, City University of Hong Kong, Hong Kong SAR, China
[2] Department of Environmental Science and Engineering, University of Science and Technology of China, Hefei 230052, China
[3] Key Laboratory of Interface Science and Engineering in Advanced Materials, Ministry of Education, Taiyuan University of Technology, Taiyuan 030024, China
[4] Department of Chemistry, Tsinghua University, Beijing 00190, China
[5] School of Chemical Engineering, Adelaide University, Adelaide, SA 5005, Australia
* Correspondence: yanguo@ustc.edu.cn (Y.G.); zhuyf@tsinghua.edu.cn (Y.Z.)

How To Cite: Zhou, Q.; Mao, Y.; Guo, W.; et al. Artificial Intelligence in Energy Catalysis: From Catalyst Screening to Decision-Making. *Science for Energy and Environment* **2026**, *3*(1), 3. https://doi.org/10.53941/see.2026.100003

Received: 10 April 2026
Revised: 8 May 2026
Accepted: 15 May 2026
Published: 18 May 2026

Abstract: Machine learning is increasingly used to rank catalysts and accelerate discovery, but in energy catalysis the practical value of a model is rarely determined by ranking alone. Activity, selectivity and stability emerge from working catalysts and interfaces that depend on potential, solvent, temperature, pressure, transport and degradation, and these can differ substantially from the nominal materials used to build models. In this Review, we examine when screening-based workflows remain sufficient and when machine learning adds value beyond catalyst screening. We organize the field around three linked shifts, from static descriptors to operando-relevant representations, from offline prediction to closed-loop discovery, and from single-objective activity optimization to sustainability-aware multi-objective optimization. Using CO_2 electroreduction, acidic oxygen evolution and higher alcohol synthesis as case studies, we show how these shifts can change catalyst ranking, experimental priorities and validation at the device or reactor level. We finally outline practical criteria for data reporting, extrapolative testing, uncertainty quantification and validation on the platform that defines success.

Keywords: artificial intelligence; energy catalysis; catalyst screening; operando catalysis; multi-objective optimization

1. Introduction

Energy catalysis underpins technologies that convert electricity and renewable feedstocks into fuels and chemicals, and artificial intelligence (AI) and machine learning are increasingly used to accelerate discovery in this space. Early studies largely treated AI as a tool for accelerating catalyst ranking or building surrogates around density functional theory calculations [1–7]. More recent reviews and roadmaps have broadened this agenda to include data infrastructure, richer representations and tighter coupling between computation and experiment [8–11]. Recent studies also emphasize integrated AI methods, shared data infrastructures, large AI models, universal descriptor frameworks, and autonomous catalyst-discovery workflows [12–15]. Within this broader landscape, the present Review is organized around a decision-centered perspective that examines when AI changes an experimental recommendation relative to a screening-first baseline and uses that perspective to compare representation, experiment-selection strategy,

and validation endpoint across distinct catalytic systems. That broader framing is especially important for energy catalysis, because catalytic performance is defined not only by composition and nominal structure but also by working states that emerge under applied potential, solvent organization, adsorbate coverage, temperature, pressure, mass transport and reactor hydrodynamics [16–19]. Figure 1 therefore frames this Review along a decision pathway from nominal materials to working interfaces, devices and processes.

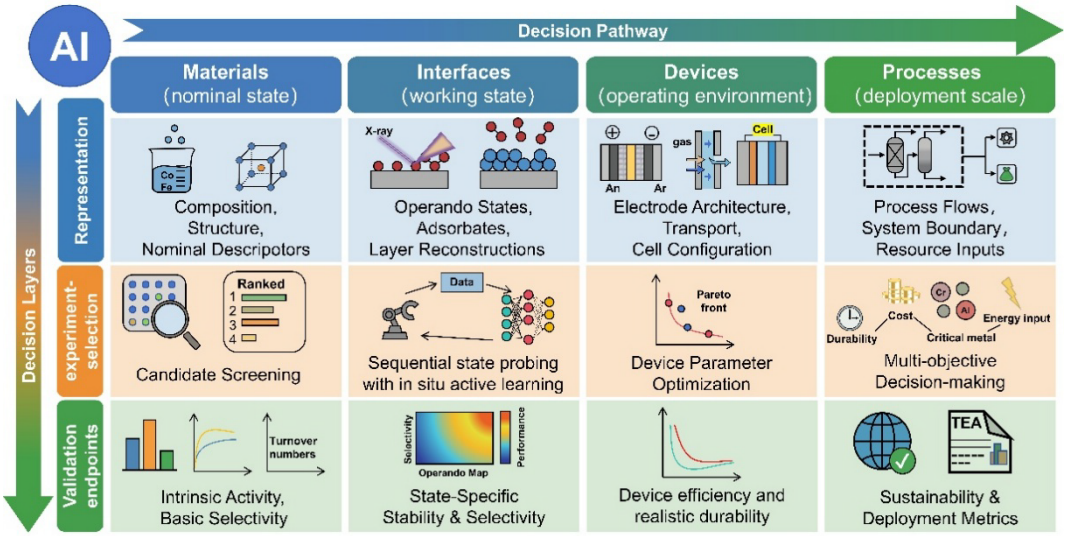

Figure 1. Decision pathway for artificial intelligence in energy catalysis. The review follows catalytic decisions from nominal materials to working interfaces, devices and processes. Three questions recur across this pathway, including representation, experiment-selection strategy and validation endpoint. Artificial intelligence moves beyond screening when it changes at least one of these decisions for a stated catalytic endpoint.

Seen in this way, energy catalysis produces evidence at several coupled levels. Theory and computation identify candidate materials and plausible reaction energetics. Synthesis and pretreatment determine particle size, defect density, facet exposure, support interactions, residual species, and metastable phases before testing even begins [18,20–22]. Recent chemistry-led catalyst studies show that fields, confinement, proton channels, and atomic precision dual sites pre-organize charge/proton/multielectron pathways before AI-based test interpretation [23–27]. Testing conditions, including electrolyte identity and concentration, applied potential or current protocol, temperature, pressure, flow field, residence time, and reactor geometry, can then reshape apparent activity, selectivity, and stability through local-environment and transport effects [16,17,19,28]. Operando measurements interrogate the evolving catalyst-electrolyte or catalyst-gas interface under bias, temperature or reactive flow. Device and reactor tests expose transport losses, integration constraints and degradation pathways that are hidden in idealized half-cell or surface models. Techno-economic and life-cycle analyses then define the practical boundary conditions for scale, cost, element availability and environmental burden. AI becomes most useful when it links these levels in a way that changes an experimental recommendation. The central issue is therefore not prediction alone, but decision-making. Which catalyst should be synthesized next, which experiment should be prioritized, which validation platform should be decisive, and whether a candidate remains attractive once durability, criticality or process constraints are taken into account.

In this Review, we take a decision-centered view of AI in energy catalysis. We use artificial intelligence as an umbrella term that includes machine learning models, data-driven optimization and automation-assisted experimental workflows. These inputs may include microscopic state variables such as charge distribution, adsorption-site identity, or adsorbate geometry when these evolve under operation [17,22,28–30]. And we use machine learning when the discussion concerns a specific predictive, generative or uncertainty-aware model. We further distinguish nominal catalyst descriptions from operando-relevant representations, by which we mean model inputs that are conditioned on both the catalyst and its working environment. We also use sustainability-aware optimization to describe workflows that consider catalytic performance together with durability, critical-element demand, process variables or life-cycle burden [8,9,16,31–33]. Here, the sustainability-aware objective includes at least one deployment-relevant non-activity criterion, such as durability, critical-element demand, mass-specific activity, cost, circularity, or life-cycle burden, alongside catalytic performance. Throughout the Review, screening-first workflows denote schemes that

rank candidates mainly within nominal descriptor space, whereas beyond-screening evidence means evidence that an added representation, automation layer, or broader objective changes the experimental recommendation under a stated catalytic endpoint.

From this perspective, the practical value of AI can be assessed across three recurring decision layers. The first is representation, whether the model captures the catalyst state that actually controls reactivity under the intended conditions. The second is experiment-selection strategy, whether the workflow improves how costly, sequential experiments are chosen. The third is the validation endpoint, whether the decisive criterion is aligned with the real use case and not limited to a single activity metric measured on an idealized platform. The sections that follow are organized around these three layers. We first define the boundary conditions under which screening-first workflows remain effective and then examine representative case studies in carbon dioxide electroreduction, acidic oxygen evolution and higher alcohol synthesis, where richer representations, closed-loop discovery or broader objectives lead to a different and more useful experimental recommendation.

2. Screening-First Workflows and Their Limits

Screening-first artificial intelligence still leaves a gap between what is optimized computationally and what governs performance experimentally. Many workflows retain the logic of high-throughput screening. A nominal bulk or surface structure is reduced to a compact descriptor set, and model quality is judged mainly by interpolation within that representation space [2–5,9]. Under reaction conditions, catalysts can change oxidation state, adsorbate coverage, vacancy population, surface termination, local strain, or support interaction; in more extreme cases they can reconstruct, dissolve and redeposit, or form new active phases [16,17,34–36]. This strategy provides efficient prioritization, and it remains a productive entry point for large candidate spaces. Its limits become clear when the catalytically relevant state is generated by the operating environment itself. In electrocatalysis, applied potential, double-layer structure, solvent orientation, ion pairing, and gas-liquid transport can all alter intermediate stabilization and reaction barriers. In thermocatalytic process and flow-reactor systems, temperature, feed composition, residence time, and reactor hydrodynamics can reshape coverages, selectivity, and apparent stability [16,17,19]. Models trained on static or ex situ representations can then learn correlations that hold within the training set but do not hold under the intended endpoint. Constant-potential machine learning molecular dynamics at the Ag-H_2O interface showed, for example, that water reorganization controls carbon dioxide reduction barriers, while multi-modal synchrotron workflows showed that experimentally measured spectral descriptors can recover structure-performance relations that composition alone does not capture [17,37].

Data structure is equally important. Catalysis is constrained by fragmented, weakly annotated, and non-interoperable datasets that cannot be aggregated straightforwardly across laboratories or across theory and experiment [31–33,38–41]. Synthesis provenance, pretreatment history, reactor configuration, calibration procedures, and failed runs are often not preserved in reusable form. The field is responding with infrastructure as well as algorithms. Findable, accessible, interoperable and reusable data infrastructures, standardized upload schemas, local data infrastructures, and benchmarking databases are converting catalytic measurements into reusable digital objects with traceable context. Text-mined corpora for carbon dioxide electroreduction show that even the published literature can become a reusable record of materials, synthesis, and testing context [42,43]. This development strengthens learning because it improves the definition of the experimental object that a model sees.

Objective choice is the third source of mismatch. Activity, or a proxy such as adsorption energy or overpotential, remains the dominant benchmark for isolated compositions [1,5,10]. Practical decisions in energy catalysis depend on multiple outputs at the same time. Carbon dioxide electroreduction couples activity to product distribution and interfacial transport. Acidic oxygen evolution couples current density to durability and noble metal utilization. In reactor-integrated thermocatalysis, catalyst formulation is intrinsically tied to operating conditions and reactor architecture [19,44]. Artificial intelligence is therefore most useful when it supports a decision that already includes the relevant constraints. This broader decision frame makes screening-first workflows easier to interpret, because it clarifies when a coarse ranking is sufficient and when richer evidence is worth the cost.

Screening-first workflows continue to deliver value in three common situations. They are effective in early composition triage when the aim is to remove clearly poor candidates from a very large space. They are effective within chemically coherent families when the active state remains close to the nominal structure and the ranking landscape is smooth enough that small descriptor errors do not change which material is

synthesized next. They are also effective when the practical decision occurs before device integration and before the additional complexity of operando data, automation, or sustainability analysis can influence the recommendation [34,45–47]. In these settings, screening-first artificial intelligence provides a rational and economical way to compress search space, identify unpromising regions, and focus experimental effort.

Operando-relevant workflows serve a different decision class. They become important when the working state departs from the nominal structure, when interface-sensitive observables shift candidate ranking, or when the decisive endpoint involves degradation, transport, or device integration [16,17,37,48]. These studies also face strong constraints. Paired datasets that link structure, operando observables, and performance are still sparse. A measured signal may track only one projection of the active state. Operando cells can alter transport, bubble management, or local composition. The practical contribution of an operando-relevant model therefore becomes strongest when the added modality changes a candidate ranking, an experimental queue, or a mechanistic conclusion that matters to the intended endpoint.

Closed-loop discovery has its own decision class and its own limits. Active learning and automated experimentation are most persuasive when the next experiment is expensive, when the search space is high dimensional, and when the order of subsequent experiments is reported clearly enough for acceleration to be interpreted [49–53]. These workflows remain capital intensive, laboratory specific, and sensitive to calibration, hardware reliability, and human supervision. Their strongest claims arise when a closed loop identifies candidates outside the initial candidate set, reduces the number of experiments needed to reach a target, or shifts the explored region of composition and condition space in a way that can be reproduced.

Sustainability-aware optimization also requires an explicit boundary. Raw material cost, criticality indices, short stability tests, or simplified life cycle estimates can all be useful proxies, yet each one speaks to a different time scale and decision horizon [54–60]. Short-duration tests are prone to misranking candidates whose device-level lifetime diverges significantly. Material price can mislead when loading or regeneration dominates total burden. A life cycle estimate can be informative for early triage but still miss separations, balance of plant, or replacement frequency. The most useful sustainability frame therefore matches the proxy to the real decision horizon and reports clearly which burdens are included.

These boundary conditions support a simple standard for evidence. A beyond-screening claim is strong when the added representation, automation layer, or sustainability frame changes the experimental recommendation under a stated endpoint. A beyond-screening claim is weaker when the extra information enriches interpretation but leaves the ranked candidate set unchanged. Figure 2 summarizes these decision boundaries, Table 1 defines when beyond-screening evidence is needed, and Table 2 lists common sustainability proxies together with their predictable biases.

Table 1. Conditions under which screening-first workflows are sufficient and conditions that justify beyond-screening evidence.

Decision Context	Screening First is Often Sufficient	Beyond-Screening Evidence is Needed When	Stronger Evidence Threshold
Early composition triage	The task is to remove clearly poor candidates from a very large design space	Small changes in rank determine which materials are synthesized next	External validation shows that ranking is stable after adding the new modality
Descriptor refinement for stable catalysts	The working state stays close to the nominal structure	Potential, coverage, or restructuring changes the active state	Operando data alter candidate ranking and the resulting synthesis choice
Electrocatalysis under interfacial sensitivity	Adsorbate energetics dominate and transport is secondary	Double-layer structure, solvent, ion pairing, or local pH changes the pathway	The model improves prediction under the intended current density and electrolyte
Degradation-limited discovery	Short activity tests are enough for the decision	Dissolution, reconstruction, or support failure changes lifetime ranking	Durability or mass-normalized performance is validated at the device level
Reactor-integrated thermocatalysis	Material ranking is insensitive to contact pattern and space velocity	Residence time, heat transfer, or geometry changes the apparent optimum	Reactor or process validation changes which candidate is preferred
Resource-constrained discovery	Cost and criticality do not change the top ranked set	Material scarcity or loading changes the practical optimum	The final recommendation is stable after adding resource-aware objectives

Notes: The last column states the minimum additional evidence that supports a beyond-screening claim for the stated decision context. The table is organized by decision type and is intended for study design as well as retrospective evaluation.

Table 2. Common sustainability proxies in energy catalysis and their likely bias.

Proxy Used in Optimization	What the Proxy Captures	Typical Bias or Blind Spot	Most Defensible Use
Raw material cost	Immediate material price signal	Ignores loading, fabrication, regeneration, and balance of plant	Early down-selection among otherwise similar formulations
Criticality index	Supply risk and geopolitical exposure	Does not reflect lifetime extension or recycling route	Comparing candidates that rely on scarce elements
Mass-specific activity	Noble metal utilization	Can obscure short lifetime or poor device integration	Electrocatalyst triage under explicit loading constraints
Short duration stability test	Initial resistance to deactivation	Frequently mischaracterizes long-term lifetime and restart behavior	Early elimination of clearly unstable candidates
Cell voltage at fixed current	Device-relevant efficiency snapshot	Mixes catalyst, transport, and assembly effects	Comparing full devices built under matched architecture
Simplified life cycle estimate	First-order environmental burden	Often omits separations, replacement frequency, and upstream infrastructure	Screening stage when the decision horizon is stated explicitly

Notes: Each proxy corresponds to a different decision horizon. A defensible sustainability claim states both the proxy and the horizon explicitly, and it reports the main burdens that remain outside the proxy.

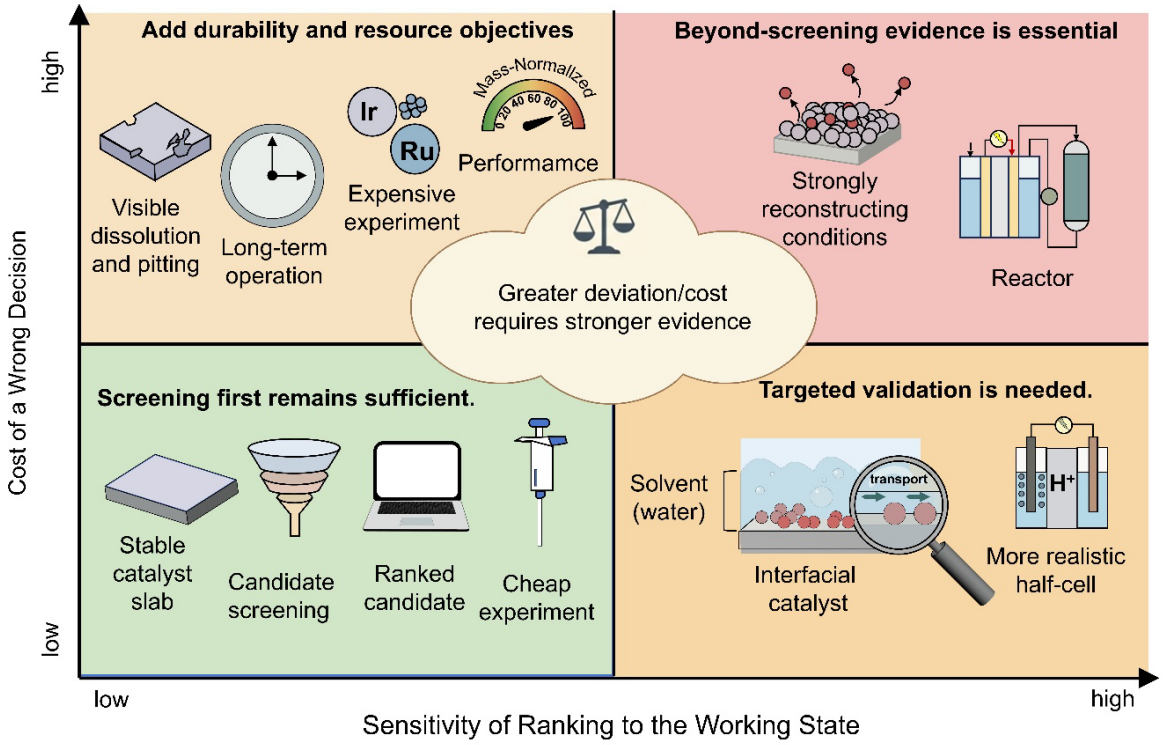

Figure 2. Boundary conditions for screening-first workflows. Screening-first approaches remain highly useful for early triage and stable ranking tasks. Beyond-screening evidence becomes decisive when the working state diverges from the nominal state, when ranking errors carry a high cost, or when validation depends on interface, degradation, device, or reactor context.

3. Three Shifts beyond Screening

Figure 3 organizes the three shifts discussed below and makes explicit that each shift acts on a different layer of the decision pathway, namely representation, discovery workflow, or optimization target.

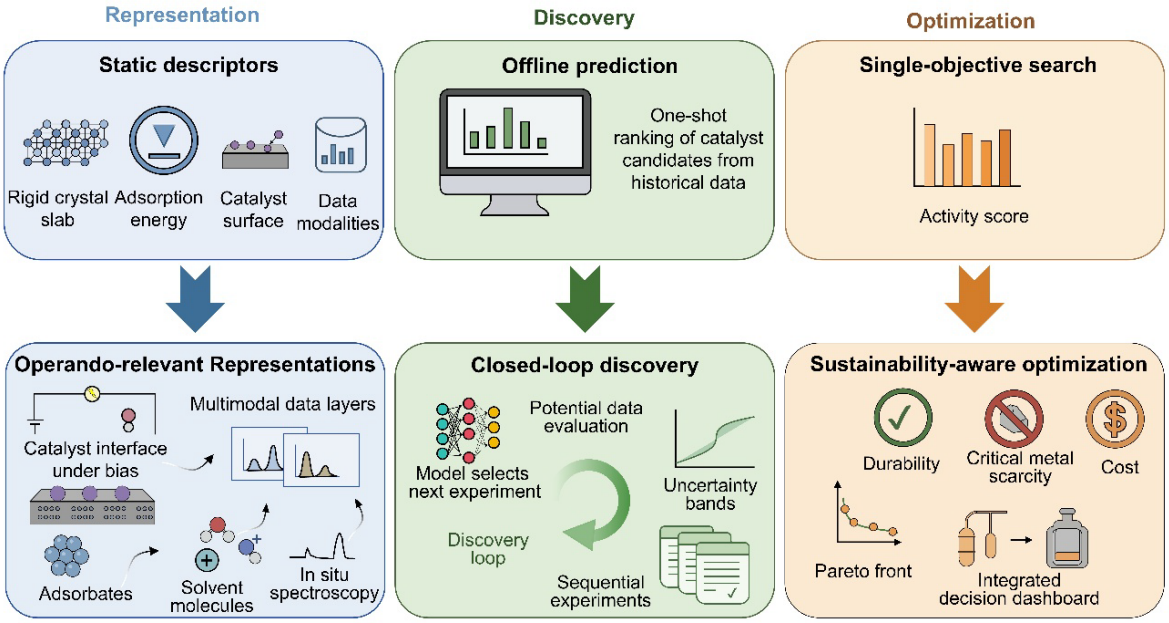

Figure 3. Three linked shifts beyond screening. The first shift expands representation from nominal structures to state-conditioned inputs. The second shift changes discovery workflow from offline ranking to sequential experiment design. The third shift broadens the objective from a single activity proxy to a sustainability-aware, multi-objective decision.

3.1. From Static Descriptors to Operando-Relevant Representations

Descriptors in catalysis were historically constructed as compressed surrogates for a nominal catalyst structure, such as composition, coordination environment, adsorption energy, or other local fingerprints [4–6]. Energy catalysis requires a broader object because the catalytically relevant state is generated jointly by the material and its operating ensemble, including electrode potential, solvent organization, ionic composition, adsorbate coverage, water activity, temperature, and pressure [16,17]. This broader object also includes microscopic characteristics that evolve with the operating system, such as charge redistribution or band alignment, adsorption-site identity, adsorbate molecular geometry, and defect or vacancy population [20,22,28–30]. Operando-relevant representation extends model input from a nominal structure to a state-conditioned description that can respond to the intended environment.

This expansion is visible in atomistic learning frameworks that move beyond fixed-composition and constant-charge surface models. Atomistic learning in the electronically grand-canonical ensemble treats system energy as a function of both atomic positions and electrode potential [61]. Related constant-potential approaches show that electrocatalytic barriers can shift through nuclear quantum effects, and grand-canonical studies of nitrate adsorption and dissociation show that the electrochemical state can change adsorption and dissociation trends across metals and dilute alloys [29,30]. At a different scale, local-atomic-environment descriptor frameworks likewise show how richer local representations can generalize across electrochemical catalyst families and guide more transferable catalyst design [14]. Reusable interface datasets, topology-guided active phase sampling, reactive neural networks for hydrated zeolites, and reviews of machine learning potentials all point in the same direction [34,62–64]. The modelled object now includes a distribution of accessible states under a specified potential, solvent loading, or interfacial environment.

Experiment also enters the representation layer more directly than before. Multimodal synchrotron workflows can treat X-ray absorption near-edge structure, extended X-ray absorption fine structure, X-ray diffraction, and related observables as direct inputs to structure-performance models [37]. This step matters because it turns spectral signatures into model features that can be compared directly with simulation-derived descriptors. It also changes the way agreement is judged. Agreement is no longer limited to an average activity value. Agreement can include whether the model is consistent with the observed working state.

Operando-relevant representation therefore creates two concrete benefits for energy catalysis. It improves the chance that a model is aligned with the state that actually controls reactivity, and it provides a clearer bridge between simulation, operando measurement, and endpoint-specific validation. These benefits matter most in systems that show potential-driven restructuring, explicit solvent effects, or strong

coupling between local interface and macroscopic output. The representation shift is therefore a decision shift because it changes which physical object the model is allowed to rank.

3.2. From Offline Prediction to Closed-Loop Discovery

Offline prediction treats model construction and experimentation as temporally separated activities. Historical data are assembled, a surrogate is trained, and candidate catalysts are then validated in a downstream campaign. This arrangement is efficient for retrospective ranking. Energy catalysis, however, is a sequential research problem in which each new experiment updates the informative region of composition and condition space [8,50,51]. Closed-loop discovery addresses that structure directly by coupling surrogate modelling, experiment-selection strategy, experiment execution, and retraining within a single workflow.

This shift changes what counts as model success. The important quantity is the quality of the next experiment selected under a limited budget. Active learning and self-driving laboratories matter in energy catalysis because experiments are expensive, candidate spaces are high dimensional, and the decisive endpoint often depends on a small number of informative measurements [49–53]. A useful closed loop therefore reports the acquisition rule, the baseline policy, the budget, the intervention log, and the stopping criterion. These details make it possible to compare workflows across laboratories and to judge whether automation actually changed the discovery trajectory.

Several energy catalysis studies show this effect clearly. In acidic oxygen evolution, a multistage workflow integrating data mining, active learning, and domain adaptation systematically narrowed composition and process space before culminating in proton-exchange-membrane-relevant validation of a Ru-Mn-Ca-Pr oxide catalyst [45]. In a separate study, density functional theory combined with Bayesian learning identified surface Ir-doped TiO_2 as a low-iridium oxygen evolution catalyst with strongly improved iridium mass-specific activity [46]. In thermocatalysis, active learning reduced an estimated Fe-Co-Cu-Zr composition and condition space of about five billion combinations to eighty-six experiments while enabling Pareto analysis across higher alcohol productivity and undesired by-products [44]. These examples show a concrete decision change because the model influences which experiment comes next and how the campaign spends its budget.

Closed-loop discovery also broadens the optimization target. Reactor optimization and the Reac-Discovery platform show that catalyst composition, operating conditions, and reactor geometry can all enter the same iterative loop [19,65]. Autonomous Pareto-front mapping in catalysis extends this logic to simultaneous optimization of yield and selectivity under controlled budget [66]. The workflow becomes stronger as it aligns experiment-selection strategy with the true endpoint, because the experiment queue then serves device or reactor performance directly instead of serving a surrogate metric only.

3.3. From Single-Objective Activity Maximization to Sustainability-Aware Optimization

Sustainability-aware optimization expands the objective function from activity alone to a set of constraints that are already present in energy technology. In this Review, that means optimizing activity or selectivity together with at least one deployment-relevant non-activity criterion, such as durability, critical-element demand, mass-specific performance, cost, circularity, or life-cycle burden. Recent perspectives on sustainable materials and circularity argue that model-guided discovery becomes more useful when durability, criticality, cost, and environmental burden are considered within the same optimization frame [54]. Catalysis research has carried this logic for many years through life cycle assessment and related process analyses [55–58]. Energy catalysis is now beginning to encode these constraints early enough that they influence which candidates are explored.

The key gain is not conceptual breadth by itself. The gain is a more realistic ranking problem. Sustainability assessment of electrochemical carbon dioxide conversion catalysts showed that catalyst choice governs selectivity, durability, supply risk, and life cycle burden at the same time, and that improved stability can lower several burdens together [59]. A related analysis of water electrolysers and fuel cells formalized material sustainability across supply risk, environmental impact, social impact, resource depletion, circularity, and substitutability [60]. These studies make durability and critical element demand part of the optimization target.

Multi-objective methods already support this broader target. Machine learning and genetic algorithms have been used to map Pareto fronts for high-entropy-alloy oxygen evolution catalysts under simultaneous activity and stability objectives [67]. Machine-learning-guided design of tungsten single atoms in

oxyhydroxides added a cost-aware and noble-metal-free frame to water electrolysis catalyst discovery [68]. Reviews of catalysis and machine learning now present multi-objective optimization as a routine design principle with broad applicability across catalyst design and process selection [69].

The strongest sustainability-aware studies also match the proxy to the decision horizon. A short stability test is useful for early elimination. Mass-specific activity is useful when loading constraints shape device design. A first-order life cycle estimate is useful when the scope of the estimate is stated explicitly. A sustainability-aware workflow becomes most convincing when it states which horizon it addresses, which proxy it uses, and which burdens remain outside the model. This is a constructive shift because it gives the field a clearer way to compare catalyst recommendations that differ in activity, durability, criticality, and process context.

4. Representative Case Studies in Energy Catalysis

The three case studies below make the artificial intelligence role explicit. In each case, the key question is whether artificial intelligence changes the representation, the experiment-selection strategy, the validation endpoint, or a combination of the three. The thermocatalysis subsection uses higher alcohol synthesis as the primary case study, while reverse water-gas shift and related process optimization serve as closely connected examples of catalyst-condition-reactor co-design. Table 3 maps representative studies onto those functions, and Figure 4 aligns the same cases visually across chemistry, main artificial intelligence role, and decisive endpoint.

Table 3. What artificial intelligence contributes in representative energy catalysis studies.

System and Representative Study	Main Artificial Intelligence Contribution	Data Modality or Experiment-Selection Strategy	Decision Changed	Validation Endpoint
Carbon dioxide electroreduction, active learning across intermetallics [70]	Early-stage ranking	Density functional theory-derived descriptors with active learning	Which alloy surfaces enter the synthesis queue	Computational triage
Carbon dioxide electroreduction, active motif map [71]	Representation of activity and selectivity together	Structural motifs linked to experimental validation	Which alloy motifs are worth testing for product control	Alloy synthesis and product selectivity
Carbon dioxide electroreduction, constant-potential molecular dynamics [17]	Operando representation	Explicit solvent and potential-dependent interface sampling	Whether interfacial water changes the preferred pathway	Mechanistic barrier change under bias
Acidic oxygen evolution, multistage active learning [45]	Experiment-selection strategy	Batchwise model update with domain adaptation	Which composition and process region is explored next	Proton exchange membrane single cell
Acidic oxygen evolution, Bayesian low-iridium discovery [46]	Resource-aware objective	Density functional theory plus Bayesian search under iridium utilization constraints	Which catalyst is preferred after mass-specific activity is counted	Half cell and mass-normalized activity
Higher alcohol synthesis, active learning [44]	Multi-objective decision strategy	Sequential experiments with Pareto analysis	Which catalyst and conditions are preferred after by-products are counted	Long-duration reactor performance
Reverse water-gas shift, extrapolative learning [72]	Out-of-domain search	Iterative model-to-experiment cycles	Whether elements absent from the initial candidate set should be explored	Experimental catalyst discovery
Reactor design and Reac-Discovery [19,65]	Process and geometry optimization	Multi-fidelity or semi-autonomous reactor search	Which reactor geometry is preferred for the chemistry	Reactor output and space-time yield

Notes: The table is organized by the part of the decision pathway changed by artificial intelligence. Some studies modify more than one layer, but the listed contribution identifies the main reason the study changes an experimental recommendation.

4.1. Carbon Dioxide Electroreduction

Carbon dioxide electroreduction provides a clear sequence of artificial intelligence roles. Early studies used active learning and surrogate models to screen large catalyst spaces. Later studies used richer representations to couple activity and selectivity. The newest contributions add explicit interfacial sampling and potential dependence. This sequence makes the development of the field easy to see because the chemistry stayed broadly similar while the decision layer changed.

Early active learning across intermetallics showed that automated density functional theory and machine learning schemes could prioritize large alloy spaces without exhaustive enumeration [70]. Active

machine learning later guided discovery of Cu-Al electrocatalysts with high ethylene selectivity at technologically relevant current density [73]. Active motif representations then predicted both activity and selectivity across metallic catalysts and supported experimental validation of Cu-Ga and Cu-Pd alloys [71]. More recent inverse-design workflows expanded the search beyond known prototypes and enabled validation of predicted alloys with high Faradaic efficiency [74]. Very recent work on Cu-based single-atom alloys further suggests that large-model-assisted and autonomous workflows can extract transferable design principles for CO_2 electroreduction alongside the identification of individual candidates [15]. In all of these studies, artificial intelligence compressed candidate space and changed which candidates were synthesized and tested next. The main contribution lay in representation and ranking.

The chemistry of the system explains why richer representation became necessary. Carbon dioxide electroreduction spans several product classes and several control variables at once

$$CO_2 + 2H^+ + 2e^- \rightarrow CO + H_2O$$

$$2CO_2 + 12H^+ + 12e^- \rightarrow C_2H_4 + 4H_2O$$

Once product distribution becomes central, simple activity ranking is no longer sufficient. Kinetic analysis combined with rotating ring-disk voltammetry showed that the effective carbon-monoxide adsorption free energy on carbon dioxide electroreduction active sites varies with catalyst identity, cation identity and concentration, applied potential, and surface structure [75]. Operando and theory-led studies on copper catalysts showed that ethylene and ethanol formation proceed through distinct intermediates and distinct site motifs [76]. This evidence clarifies the task for artificial intelligence. The model must capture a selectivity landscape that is conditioned on interface and coverage, not only on nominal composition.

Constant-potential machine learning molecular dynamics at the Ag-H_2O interface carried this step further by showing that applied potential reshapes interfacial water orientation and the hydrogen-bond network, which then changes carbon dioxide reduction barriers through solvent dynamics [17]. Explicit-solvent simulations on copper showed that alkali cations alter carbon dioxide activation and proton accessibility [77]. Device-level studies added a transport dimension by showing that local control of the H_2O to CO_2 ratio supports high-rate multi-carbon formation [78]. Dynamic restructuring studies on Ag-Cu interfaces and on copper single crystals showed that the working catalyst evolves away from the nominal surface model under operating conditions [79,80]. In this case, artificial intelligence changes the catalytic decision most strongly when it follows the interface that actually operates under bias. Representation is therefore the decisive layer, and product distribution under realistic bias is the decisive endpoint. Thus, the dominant decision layer in CO_2 electroreduction is representation.

4.2. Acidic Oxygen Evolution and Proton Exchange Membrane Water Electrolysis

Acidic oxygen evolution offers a different sequence of artificial intelligence roles. Here the central issue is not only activity. The field must allocate scarce iridium, maintain durability under strongly oxidizing conditions, and validate at the level of the proton exchange membrane device. Artificial intelligence contributes most when it couples these constraints to the discovery workflow.

The fundamental half reaction is well known

$$2H_2O \rightarrow O_2 + 4H^+ + 4e^-$$

The working catalyst, however, is strongly dynamic. Under proton exchange membrane operating conditions, iridium oxides remain the most reliable reference materials, yet activity and durability are governed by operando structural evolution [35]. Comparative analyses showed that aqueous model cells can underestimate practical lifetime, while device-level perspectives on proton exchange membrane water electrolysis emphasize current density, catalyst-layer utilization, porous-transport layers, and full-cell assembly as part of the validation context [81–83]. These results place validation endpoint at the center of the decision pathway.

Artificial intelligence has already changed how the search is run. Machine-learning-enabled screening has been used to predict acidic stability directly and to prioritize candidates under explicit durability constraints [84,85]. A multistage workflow combining data mining, active learning, and domain adaptation converged on a Ru-Mn-Ca-Pr oxide from a large composition and process space and established its relevance through proton exchange membrane single-cell validation [45]. Density functional theory combined with Bayesian learning identified surface Ir-doped TiO_2 as a low-iridium oxygen evolution catalyst with strongly improved iridium mass-specific activity relative to commercial IrO_2 [46]. In this

system, the main artificial intelligence contribution lies in experiment-selection strategy and resource-aware objective design. The model directs which composition and loading region is explored next, and it does so under a constraint that already includes scarcity.

The broader mechanistic literature explains why this endpoint-centered strategy is necessary. Ni-stabilized RuO_2 established that non-iridium anodes can function durably in proton exchange membrane water electrolysis [86]. Operando synchrotron studies showed that Mn-modified RuO_2 can switch the dominant pathway to an oxide-mediated mechanism based on asymmetric dual-active sites [36]. Other work showed that strain heterogeneity in RuO_2, high-entropy RuO_2, and support-tuned iridium reconstruction all reshape the balance between activity and structural persistence [87–89]. At the current frontier of validation, single-faceted IrO_2 monolayers and low-iridium Ru-based anodes are being assessed directly in durable device operation [90,91]. This body of work shows that the relevant object is an operando trajectory with an explicit device endpoint. Artificial intelligence becomes most useful when the workflow directs experiments toward that trajectory and treats iridium utilization, durability, and single-cell output as part of the optimization target from the beginning. Accordingly, the dominant decision layers in acidic oxygen evolution are acquisition policy experiment-selection strategy and endpoint-matched validation under durability and iridium-utilization constraints.

4.3. Thermocatalysis: Reverse Water-Gas Shift, Higher Alcohol Synthesis, and Reactor/Process Co-Optimization

Thermocatalysis changes the what must be optimized. The observed rate and selectivity are often attributes of a catalyst, condition, and reactor tuple defined by feed ratio, pressure, temperature, space velocity, pretreatment history, and transport regime. Data-driven thermocatalysis therefore moved early toward representations that already mix materials and operating variables. Interpretable machine learning for Rh-Mn-P-SiO_2 catalysts required both promoter descriptors and reaction conditions to recover higher alcohol space-time yield, and descriptor design for carbon dioxide to methanol catalysis used adsorption-energy distributions across multiple facets and sites because single-site descriptors do not represent heterogeneous thermal catalysts adequately [92,93].

The reaction network itself explains this coupling

$$CO_2 + H_2 \rightarrow CO + H_2O$$

$$nCO + 2nH_2 \rightarrow C_nH_{2n+1}OH + (n - 1)H_2O$$

Once chain growth, water formation, methane production, and carbon dioxide recycling are all considered, the relevant optimum becomes inherently multi-objective. Artificial intelligence therefore contributes most when it changes the search strategy across catalyst composition, operating condition, and reactor context at the same time.

Closed-loop experimentation already shows that this coupled design space can be navigated directly. In the reverse water-gas shift reaction, extrapolative machine learning executed repeated model-to-experiment cycles and identified catalysts that outperformed known references, including candidates containing elements absent from the initial set [72]. In direct carbon dioxide hydrogenation to methanol, Bayesian optimization was coupled to automated synthesis and high-throughput fixed-bed testing under a multi-objective function that included methanol-oriented performance, by-products, and cost [94]. These studies show a clear contribution in experiment selection. The workflow directs the experimental sequence toward under-sampled regions and updates that queue after each informative result.

Higher alcohol synthesis provides the clearest system-level example. Active learning navigated an Fe-Co-Cu-Zr composition and condition space of about five billion combinations in eighty-six experiments and replaced single-metric optimization with a Pareto front over productivity and undesired carbon dioxide and methane formation [44]. Later operando work linked the surface Fe_5C_2 to Fe_3O_4 + Cu ratio to higher alcohol and olefin formation [95]. Reactor optimization and the Reac-Discovery platform extend the same logic to geometry and process configuration [19,65]. In this case, artificial intelligence changes ranking, experiment-selection strategy, and endpoint together. The preferred recommendation is not simply a catalyst. It is a catalyst-condition-reactor tuple evaluated by long-duration reactor performance and process-aware trade-offs.

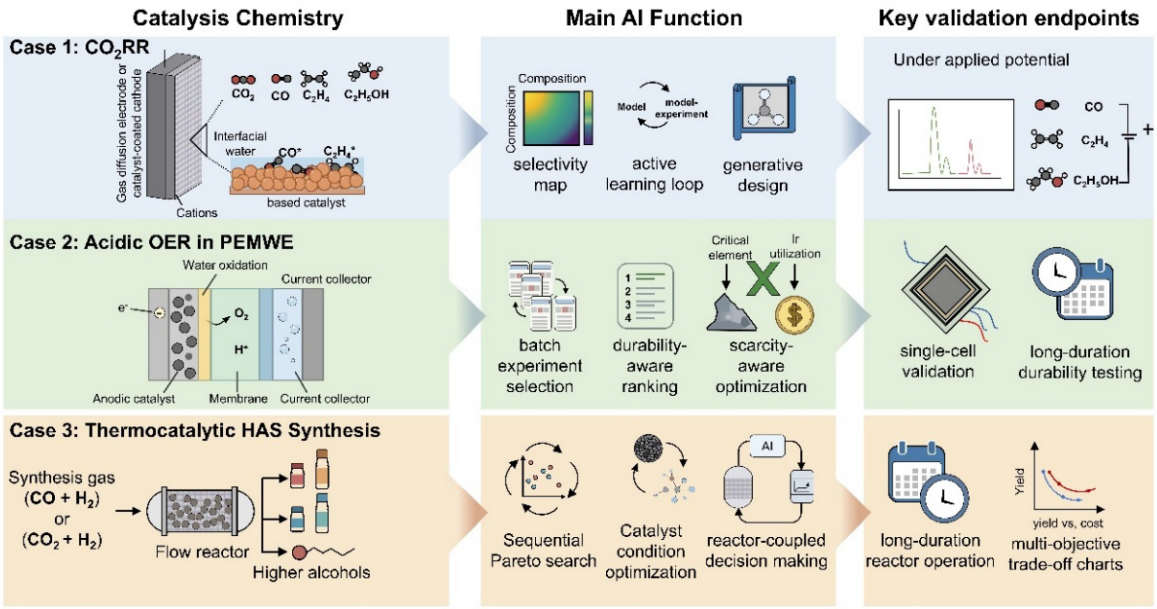

Figure 4. How artificial intelligence changes decisions in three energy catalysis case studies. Carbon dioxide electroreduction is dominated by representation and selectivity-aware ranking. Acidic oxygen evolution is dominated by experiment-selection strategy under durability and scarcity constraints. Higher alcohol synthesis is dominated by multi-objective search across catalyst, condition, and reactor choices. The validation endpoint column shows why each workflow requires a different level of validation.

5. Outlook and a Practical Benchmark Framework

The most useful next step for the field is to convert broad principles into reporting and validation rules that can be reused across studies. Table 4 summarizes a practical benchmark framework for artificial intelligence in energy catalysis, and Figure 5 condenses the same framework into a layered checklist that can be used during study design, peer review, or post-publication comparison.

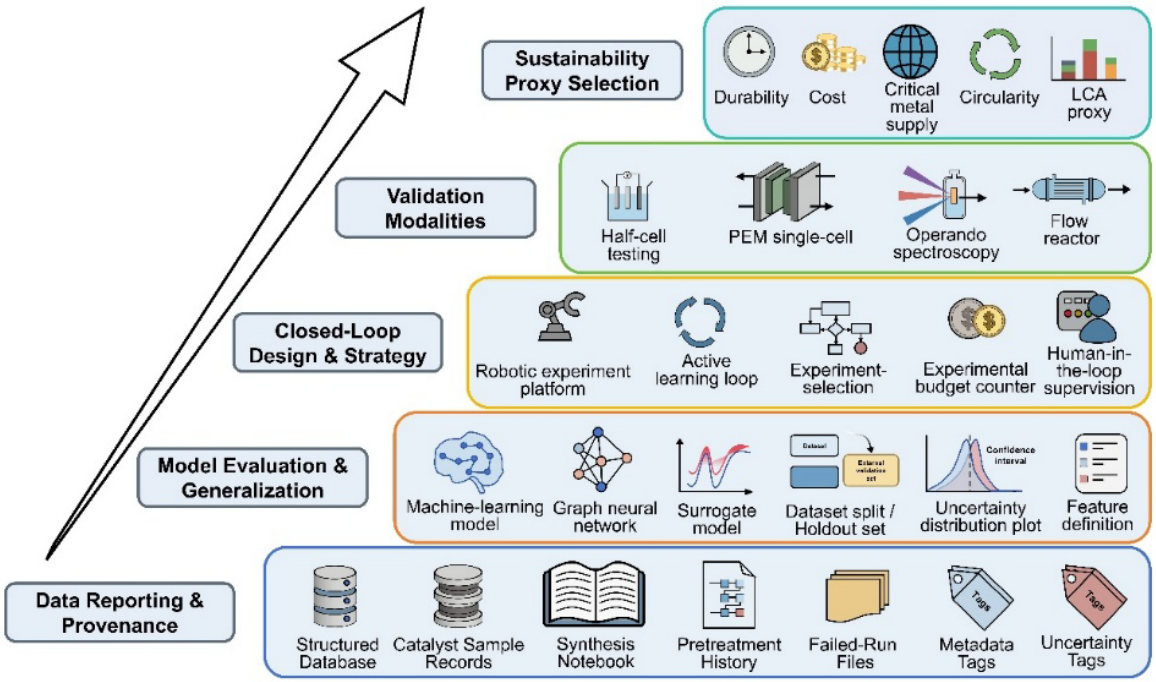

Figure 5. Framework for artificial intelligence in energy catalysis. A strong study aligns reusable data, reliable model evaluation, benchmarked automation, realistic validation, and a sustainability frame that matches the decision horizon. The figure is intended as a checklist that can be applied during study design, review, and comparison across publications.

Table 4. A practical benchmark framework for artificial intelligence in energy catalysis.

Layer	Minimum Reporting	Stronger Evidence	Typical Failure Mode
Data	Composition, synthesis, pretreatment, cell or reactor configuration, feed or electrolyte, failed runs, output uncertainty	Machine-readable schema, versioned provenance, cross-laboratory comparability, preserved negative data	Performance values detached from context
Model	External split, feature definition, decision objective, uncertainty estimate	Leave-family-out or leave-laboratory-out tests, calibrated uncertainty, ablation across data modalities	High random-split accuracy with poor extrapolation
Closed loop	Batch size, experiment-selection rule, baseline selection strategy, total experimental budget, human-intervention record	Acceleration factor and enhancement factor against explicit baselines, reproducible workflow software, hardware-failure accounting	Claimed acceleration without a comparable baseline
Validation	Stated endpoint such as half cell, single cell, or reactor output	Operando corroboration of the modelled state plus device- or reactor-level validation	Mechanistic interpretation without endpoint-matched validation
Sustainability	At least one explicit proxy among durability, criticality, cost, or life cycle burden	Joint treatment of durability, criticality, cost, and a proxy matched to the decision horizon	Proxy mismatch between study and real decision

Notes: The framework is intended to be used prospectively and retrospectively. Minimum reporting defines the information needed to interpret the claim, while stronger evidence defines the information needed to compare the claim across studies.

Progress should first be judged by whether the field becomes cumulatively learnable. The immediate bottleneck is no longer only measurement scarcity. The decisive issue is the lack of machine-readable experimental objects that preserve provenance, configuration, and performance outputs in interoperable form [32,33,40,41]. By machine-readable experimental objects, we mean structured digital records (such as standardized JSON or XML entries, electronic-lab-notebook exports, or database rows) that encode composition, synthesis route, pretreatment history, cell or reactor geometry, feed or electrolyte composition, calibration metadata, raw and processed outputs, units, uncertainty, and failed runs as explicit fields instead of narrative prose alone. Some of the most important advances over the next few years will therefore come from shared data structures, versioned metadata, and reusable benchmark sets [13,31–33]. These developments make it possible for a model trained in one study to remain interpretable in the next.

A second criterion is extrapolative utility. In self-driving experimentation, acceleration needs an explicit baseline, and acceleration factor and enhancement factor provide a clearer view of performance than simple counts of optimization cycles [50,51]. At the model level, conventional random hold-out tests or k-fold cross-validation can estimate interpolation error, often reported as an average RMSE, but they can overstate practical utility when training and test data are drawn from the same catalyst family, laboratory, or endpoint distribution. In energy catalysis, the central question is not only whether a model fits known data, but whether it generalizes beyond the distribution on which it was trained. Useful uncertainty quantification should therefore be treated as a core result and reported as part of the main claim. Useful uncertainty is accurate, robust, and informative for out-of-domain deployment [96,97]. Symbolic-regression-guided discovery of acid-stable oxides illustrates the value of uncertainty-aware search by recovering promising candidates outside the initial design space [98]. The most informative benchmark suites will therefore emphasize structured external validation, including leave-family-out, leave-laboratory-out, and leave-endpoint-out tests, instead of relying solely on random splits.

A third criterion is endpoint-matched validation. A model does not become experimentally decisive simply because it correlates with a measurement from a cell or reactor. The endpoint must match the intended use case. Best-practice analyses of electrocatalytic operando methods show that reactor geometry, fluid flow, species diffusion, and bubble management all influence whether an operando measurement remains commensurate with practical operation [16]. Strong evidence therefore takes the form of a chain. Uncertainty-aware prediction identifies the candidate region. Operando or state-sensitive measurements confirm that the relevant catalyst state is represented faithfully. Device- or reactor-level validation then determines whether the recommendation survives realistic throughput.

A fourth criterion is transferability across laboratories and automation platforms. Closed-loop platforms often rely on laboratory-specific automation, calibration routines, and software infrastructure, which makes apparent acceleration difficult to reproduce elsewhere [49,51–53]. A reusable claim therefore reports the reference selection strategy, the number of failed or repeated runs, the degree of human intervention, and the sensitivity of the outcome to the initial candidate set. These practical details determine whether a workflow can be transferred into another laboratory without rewriting the entire decision pathway.

A fifth criterion is alignment between sustainability proxy and decision horizon. Early discovery can benefit from simple proxies, but the proxy must be selected deliberately. Durability supports long-horizon device choices. Criticality and mass-specific performance matter when loading or material scarcity is decisive. A first-order life cycle proxy is most useful when the goal is comparative early-stage triage and when the omitted burdens are listed. Studies become easier to interpret when they state this alignment explicitly, because readers can see which decision is being supported and which future analysis is still required.

Figure 5 summarizes these five layers as a framework. The framework is useful because it gives the field a common language for progress. Data quality supports model evaluation. Model evaluation supports closed-loop design. Closed-loop design supports endpoint-matched validation. Sustainability framing then determines whether the validated catalyst remains attractive under real resource and deployment constraints.

6. Conclusions and Outlook

Five priorities emerge from the analysis above which were summarized in Figure 6. Experimentalists can strengthen the field by reporting provenance, failed runs, and validation context in machine-readable form. Model developers can increase usefulness by optimizing for extrapolation and calibrated uncertainty. Database builders can preserve operating conditions, negative data, and versioned metadata so that catalytic knowledge accumulates across papers. Device and reactor engineers can define the decisive endpoint early enough to shape the search space. Sustainability analysts can match proxy choice to durability, cost, criticality, and the intended decision horizon.

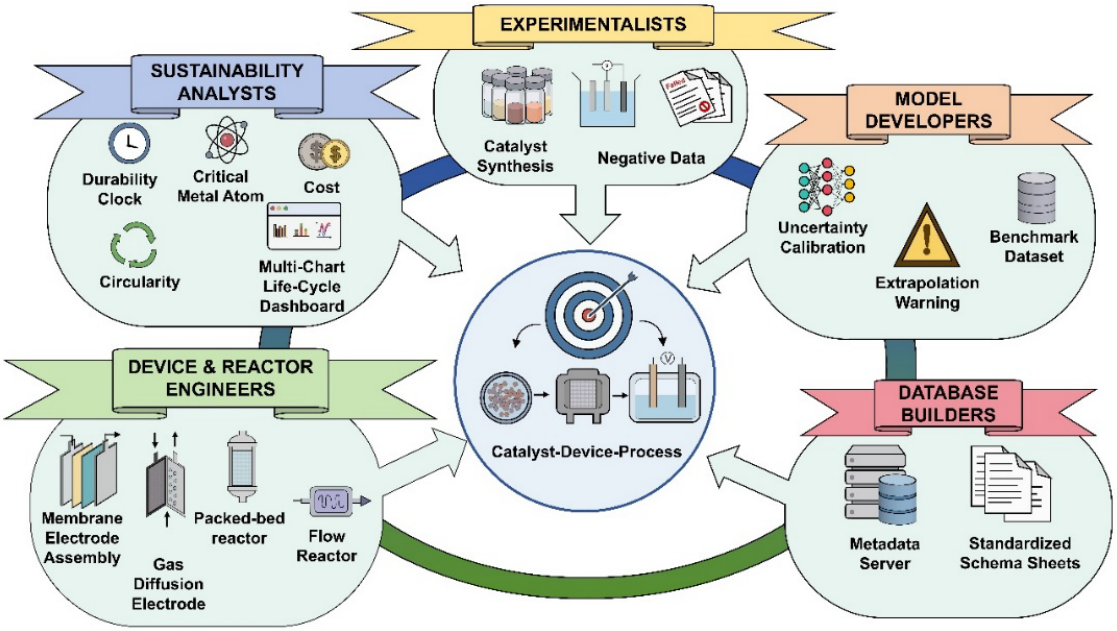

Figure 6. Priorities for the next phase of artificial intelligence in the field of energy catalysis. Progress in beyond-screening studies depends on coordination across experimental reporting, model development, database building, device or reactor validation, and sustainability assessment. The central objective is to change the experimental decision under a realistic endpoint.

These priorities describe a positive program for the next phase of the field. Energy catalysis already has the ingredients required for stronger artificial intelligence workflows, including rich theory, operando measurement, automated experimentation, and device-level testing. Emerging large-model and autonomous AI workflows may further accelerate literature synthesis, hypothesis generation, and campaign planning, but their catalytic value will still depend on reusable data and endpoint-matched validation [12,15]. The current opportunity lies in aligning these ingredients around the same experimental decision. When representation follows the working state, experiment-selection strategy follows the true bottleneck, validation follows the real endpoint, and sustainability framing follows the intended horizon, artificial intelligence becomes a practical tool for catalyst discovery and process design. Under those conditions, the most meaningful measure of progress is straightforward. Artificial intelligence matters in energy catalysis when it helps researchers choose a better experiment, validate a better catalyst, and support a better device or process recommendation.

Author Contributions

Q. Z.: writing—original draft preparation, visualization; Y. M.: writing—reviewing and editing, visualization; W. G.: writing—reviewing and editing, visualization; H. S.: writing—reviewing and editing, visualization; H. W.: writing—reviewing and editing, visualization; Y. G.: supervision, writing—reviewing and editing; Y. Z.: supervision, writing—reviewing and editing. All authors have read and agreed to the published version of the manuscript.

Funding

This research was funded by National Key Research and Development Project of China 2020YFA0710304.

Data Availability Statement

The data are available from the corresponding author upon reasonable request.

Conflicts of Interest

The authors declare no conflict of interest.

Use of AI and AI-Assisted Technologies

No AI tools were utilized for this paper.

References

1. Kitchin, J.R. Machine Learning in Catalysis. *Nat. Catal.* **2018**, *1*, 230–232. https://doi.org/10.1038/s41929-018-0056-y.
2. Li, Z.; Wang, S.; Xin, H. Toward Artificial Intelligence in Catalysis. *Nat. Catal.* **2018**, *1*, 641–642. https://doi.org/10.1038/s41929-018-0150-1.
3. Freeze, J.G.; Kelly, H.R.; Batista, V.S. Search for Catalysts by Inverse Design. Artificial Intelligence, Mountain Climbers, and Alchemists. *Chem. Rev.* **2019**, *119*, 6595–6612. https://doi.org/10.1021/acs.chemrev.8b00759.
4. Toyao, T.; Maeno, Z.; Takakusagi, S.; et al. Machine Learning for Catalysis Informatics. Recent Applications and Prospects. *ACS Catal.* **2020**, *10*, 2260–2297. https://doi.org/10.1021/acscatal.9b04186.
5. Esterhuizen, J.A.; Goldsmith, B.R.; Linic, S. Interpretable Machine Learning for Knowledge Generation in Heterogeneous Catalysis. *Nat. Catal.* **2022**, *5*, 175–184. https://doi.org/10.1038/s41929-022-00744-z.
6. Mou, L.-H.; Han, T.; Smith, P.E.S.; et al. Machine Learning Descriptors for Data-Driven Catalysis Study. *Adv. Sci.* **2023**, *10*, 2301020. https://doi.org/10.1002/advs.202301020.
7. Abraham, B.M.; Jyothirmai, M.V.; Sinha, P.; et al. Catalysis in the Digital Age. Unlocking the Power of Data with Machine Learning. *WIREs Comput. Mol. Sci.* **2024**, *14*, e1730. https://doi.org/10.1002/wcms.1730.
8. Xin, H.; Kitchin, J.R.; López, N.; et al. Roadmap for Transforming Heterogeneous Catalysis with Artificial Intelligence. *Nat. Catal.* **2026**, *9*, 102–111. https://doi.org/10.1038/s41929-026-01479-x.
9. Bozal-Ginesta, C.; Pablo-García, S.; Choi, C.; et al. Developing Machine Learning for Heterogeneous Catalysis with Experimental and Computational Data. *Nat. Rev. Chem.* **2025**, *9*, 601–616. https://doi.org/10.1038/s41570-025-00740-4.
10. Zeng, Y.; Wang, J.; Li, F.; et al. AI-Accelerated Discovery of Electrocatalyst Materials. *ACS Mater. Au* **2026**, *6*, 72–89. https://doi.org/10.1021/acsmaterialsau.5c00135.
11. Chen, H.; Kätelhön, E.; Lu, Y.; et al. 30 Years of AI for Electrocatalysis. Where We Are and What's Next. *eScience* **2026**, *6*, 100515. https://doi.org/10.1016/j.esci.2025.100515.
12. Zhang, D.; Chen, Y.; Liu, C.; et al. Accelerating Catalyst Materials Discovery with Large Artificial Intelligence Models. *Angew. Chem. Int. Ed.* **2026**, *65*, e26150. https://doi.org/10.1002/anie.202526150.
13. Zhang, D.; Jia, X.; Wang, Y.; et al. Digital Materials Ecosystem: From Databases to AI Agents for Autonomous Discovery. *Chem. Sci.* **2026**, *17*, 5782–5804. https://doi.org/10.1039/D5SC09229A.
14. Liu, Z.; Liu, Y.; Zhang, Y.; et al. Universal Catalyst Design Framework for Electrochemical Hydrogen Peroxide Synthesis Facilitated by Local Atomic Environment Descriptors. *Angew. Chem. Int. Ed.* **2026**, *65*, e18027. https://doi.org/10.1002/anie.202518027.
15. Wang, X.; Li, Z.; Zhang, D.; et al. Catalysis AI Agent Guides Discovering the Universal Design Principle of Cu-Based Single-Atom Alloy Catalysts for CO_2 Electroreduction. *Angew. Chem. Int. Ed.* **2026**, *65*, e24612. https://doi.org/10.1002/anie.202524612.

16. Prajapati, A.; Hahn, C.; Weidinger, I.M.; et al. Best Practices for *in-Situ* and Operando Techniques within Electrocatalytic Systems. *Nat. Commun.* **2025**, *16*, 2593. https://doi.org/10.1038/s41467-025-57563-6.
17. Tian, X.; Tosello Gardini, A.; Raucci, U.; et al. Electrochemical Potential-Driven Water Dynamics Control CO_2 Electroreduction at the Ag-H_2O Interface. *Nat. Commun.* **2025**, *16*, 10636. https://doi.org/10.1038/s41467-025-65630-1.
18. Jia, X.; Wang, T.; Zhang, D.; et al. Advancing Electrocatalyst Discovery through the Lens of Data Science. State of the Art and Perspectives. *J. Catal.* **2025**, *447*, 116162. https://doi.org/10.1016/j.jcat.2025.116162.
19. Savage, T.; Basha, N.; McDonough, J.; et al. Machine Learning-Assisted Discovery of Flow Reactor Designs. *Nat. Chem. Eng.* **2024**, *1*, 522–531. https://doi.org/10.1038/s44286-024-00099-1.
20. Peng, X.; Chen, L.; Liu, Y.; et al. Strain Engineering of Two-Dimensional Materials for Energy Storage and Conversion Applications. *Chem. Synth.* **2023**, *3*, 47. https://doi.org/10.20517/cs.2023.34.
21. Zhang, Y.; Liang, H. Crystal Facet Engineering of Electrocatalysts for Nitrate Reduction to Ammonia: Recent Advances and Future Perspectives. *Chem. Synth.* **2024**, *4*, 39. https://doi.org/10.20517/cs.2023.74.
22. Xia, T.; Wang, X.; Wan, J.; et al. Asymmetrically Coordinated Single-Atom Catalysts: From Synthetic Strategy to Structure-Activity Relationship. *Chem. Synth.* **2025**, *5*, 74. https://doi.org/10.20517/cs.2025.08.
23. Guo, Y.; Zhou, Q.; Nan, J.; et al. Perylenetetracarboxylic Acid Nanosheets with Internal Electric Fields and Anisotropic Charge Migration for Photocatalytic Hydrogen Evolution. *Nat. Commun.* **2022**, *13*, 2067. https://doi.org/10.1038/s41467-022-29826-z.
24. Zhou, Q.; Guo, Y.; Zhu, Y. Photocatalytic Sacrificial H_2 Evolution Dominated by Micropore-Confined Exciton Transfer in Hydrogen-Bonded Organic Frameworks. *Nat. Catal.* **2023**, *6*, 574–584. https://doi.org/10.1038/s41929-023-00972-x.
25. Guo, Y.; Zhou, Q.; Wang, L.; et al. Enhanced Hydrogen Peroxide Photosynthesis via Charge-Complementary π-Electron Sites. *Nat. Commun.* **2025**, *16*, 6297. https://doi.org/10.1038/s41467-025-61452-3.
26. Chai, H.; Guo, Y.; Nan, J.; et al. Proton-Conducting Hydrogen-Bonded Nanochannels Encapsulating Perylene-Based Photocatalysts for Solar Hydrogen Evolution. *J. Mater. Chem. A* **2026**, *14*, 10794–10799. https://doi.org/10.1039/D5TA10040B.
27. Zhou, Q.; Guo, Y.; Zhu, Y. Reticular Copper Dual Sites Embedded with Semiconductor Particles for Selective CO_2-to-C_2H_4 Photoreduction. *Nat. Catal.* **2025**, *8*, 728–739. https://doi.org/10.1038/s41929-025-01369-8.
28. Liu, W. Shining Light on Electrochemistry: A Synchrotron-Based X-ray Spectroscopic Interrogation. *Chem. Synth.* **2024**, *4*, 13. https://doi.org/10.20517/cs.2023.49.
29. Sun, M.; Jin, B.; Yang, X.; et al. Probing Nuclear Quantum Effects in Electrocatalysis via a Machine-Learning-Enhanced Grand-Canonical Constant-Potential Approach. *Nat. Commun.* **2025**, *16*, 3600. https://doi.org/10.1038/s41467-025-58871-7.
30. Sweeney, D.M.; Tran, B.; Goldsmith, B.R. A Grand-Canonical Study of the Potential Dependence of Nitrate Adsorption and Dissociation across Metals and Dilute Alloys. *Commun. Chem.* **2025**, *8*, 182. https://doi.org/10.1038/s42004-025-01579-y.
31. Gauthier, A.; Vancauwenberghe, L.; Cousty, J.-C.; et al. A FAIR Research Data Infrastructure for High-Throughput Digital Chemistry. *Digit. Discov.* **2025**, *4*, 3502–3514. https://doi.org/10.1039/D5DD00297D.
32. Mendes, P.S.F.; López, N.; De, S.; et al. Data as a Key Resource in Catalysis. A Community Account. *ChemCatChem* **2025**, *17*, e01226. https://doi.org/10.1002/cctc.202501226.
33. Schumann, J.; Näsström, H.; Götte, M.; et al. Enabling Open and FAIR Catalysis Data with Standardized Data Structures. *Nat. Catal.* **2026**, *9*, 225–229. https://doi.org/10.1038/s41929-026-01508-9.
34. Zheng, S.; Zhang, X.-M.; Liu, H.-S.; et al. Active Phase Discovery in Heterogeneous Catalysis via Topology-Guided Sampling and Machine Learning. *Nat. Commun.* **2025**, *16*, 2542. https://doi.org/10.1038/s41467-025-57824-4.
35. Mom, R.V.; Falling, L.J.; Kasian, O.; et al. Operando Structure-Activity-Stability Relationship of Iridium Oxides during the Oxygen Evolution Reaction. *ACS Catal.* **2022**, *12*, 5174–5184. https://doi.org/10.1021/acscatal.1c05951.
36. Ji, Q.; Tang, B.; Zhang, X.; et al. Operando Identification of the Oxide Path Mechanism with Different Dual-Active Sites for Acidic Water Oxidation. *Nat. Commun.* **2024**, *15*, 8089. https://doi.org/10.1038/s41467-024-52471-7.
37. Baliyan, A.; Verma, S.; Sasakawa, K.; et al. Machine Learning-Guided Multimodal Synchrotron Analysis Workflow for Fuel Cell Electrocatalyst Discovery. *Commun. Chem.* **2025**, *8*, 376. https://doi.org/10.1038/s42004-025-01800-y.
38. Lam, E.; Maury, T.; Preiss, S.; et al. General Data Management Workflow to Process Tabular Data in Automated and High-Throughput Heterogeneous Catalysis Research. *Digit. Discov.* **2025**, *4*, 539–547. https://doi.org/10.1039/D4DD00350K.
39. Moshantaf, A.; Wesemann, M.; Beinlich, S.; et al. Advancing Catalysis Research through FAIR Data Principles Implemented in a Local Data Infrastructure. A Case Study of an Automated Test Reactor. *Catal. Sci. Technol.* **2024**, *14*, 6186–6197. https://doi.org/10.1039/D4CY00693C.

40. Salim, I.; Berrouk, A.S. Toward Efficient Reporting for Catalytic Methane Reforming Data. A Framework and Review for Standardized Scientific Data Reporting. *NPJ Comput. Mater.* **2026**, *12*, 85. https://doi.org/10.1038/s41524-025-01953-3.
41. Burte, A.S.; Nair, A.; Grabow, L.C.; et al. CatTestHub. A Benchmarking Database of Experimental Heterogeneous Catalysis for Evaluating Advanced Materials. *J. Catal.* **2025**, *442*, 115902. https://doi.org/10.1016/j.jcat.2024.115902.
42. Wang, L.; Gao, Y.; Chen, X.; et al. A Corpus of CO_2 Electrocatalytic Reduction Process Extracted from the Scientific Literature. *Sci. Data* **2023**, *10*, 175. https://doi.org/10.1038/s41597-023-02089-z.
43. Chen, X.; Gao, Y.; Wang, L.; et al. Large Language Model Enhanced Corpus of CO_2 Reduction Electrocatalysts and Synthesis Procedures. *Sci. Data* **2024**, *11*, 347. https://doi.org/10.1038/s41597-024-03180-9.
44. Suvarna, M.; Zou, T.; Chong, S.H.; et al. Active Learning Streamlines Development of High Performance Catalysts for Higher Alcohol Synthesis. *Nat. Commun.* **2024**, *15*, 5844. https://doi.org/10.1038/s41467-024-50215-1.
45. Ding, R.; Liu, J.; Hua, K.; et al. Leveraging Data Mining, Active Learning, and Domain Adaptation for Efficient Discovery of Advanced Oxygen Evolution Electrocatalysts. *Sci. Adv.* **2025**, *11*, eadr9038. https://doi.org/10.1126/sciadv.adr9038.
46. Niu, X.; Chen, Y.; Sun, M.; et al. Bayesian Learning-Assisted Catalyst Discovery for Efficient Iridium Utilization in Electrochemical Water Splitting. *Sci. Adv.* **2025**, *11*, eadw0894. https://doi.org/10.1126/sciadv.adw0894.
47. Hung, L.; Yager, J.A.; Monteverde, D.; et al. Autonomous Laboratories for Accelerated Materials Discovery. A Community Survey and Practical Insights. *Digit. Discov.* **2024**, *3*, 1273–1279. https://doi.org/10.1039/D4DD00059E.
48. Tom, G.; Schmid, S.P.; Baird, S.G.; et al. Self-Driving Laboratories for Chemistry and Materials Science. *Chem. Rev.* **2024**, *124*, 9633–9732. https://doi.org/10.1021/acs.chemrev.4c00055.
49. Scheurer, C.; Reuter, K. Role of the Human-in-the-Loop in Emerging Self-Driving Laboratories for Heterogeneous Catalysis. *Nat. Catal.* **2025**, *8*, 13–19. https://doi.org/10.1038/s41929-024-01275-5.
50. Volk, A.A.; Abolhasani, M. Performance Metrics to Unleash the Power of Self-Driving Labs in Chemistry and Materials Science. *Nat. Commun.* **2024**, *15*, 1378. https://doi.org/10.1038/s41467-024-45569-5.
51. Adesiji, A.D.; Wang, J.; Kuo, C.-S.; et al. Benchmarking Self-Driving Labs. *Digit. Discov.* **2026**, *5*, 14–27. https://doi.org/10.1039/D5DD00337G.
52. Zhu, M.; Mroz, A.; Gui, L.; et al. Discrete and Mixed-Variable Experimental Design with Surrogate-Based Approach. *Digit. Discov.* **2024**, *3*, 2589–2606. https://doi.org/10.1039/D4DD00113C.
53. Salazar-Villacis, P.; Benyahia, B. The ADePT Framework for Assessing Autonomous Laboratory Robotics. *Commun. Chem.* **2026**, *9*, 99. https://doi.org/10.1038/s42004-026-01932-9.
54. Han, N.; Guo, H.; Zhang, B.; et al. Artificial Intelligence as a Driver of Sustainable Materials and Circularity. *Nat. Rev. Mater.* **2026**. https://doi.org/10.1038/s41578-026-00912-8.
55. Holman, P.A.; Shonnard, D.R.; Holles, J.H. Using Life Cycle Assessment to Guide Catalysis Research. *Ind. Eng. Chem. Res.* **2009**, *48*, 6668–6674. https://doi.org/10.1021/ie801934s.
56. Artz, J.; Müller, T.E.; Thenert, K.; et al. Sustainable Conversion of Carbon Dioxide. An Integrated Review of Catalysis and Life Cycle Assessment. *Chem. Rev.* **2018**, *118*, 434–504. https://doi.org/10.1021/acs.chemrev.7b00435.
57. Charalambous, C.; Moubarak, E.; Schilling, J.; et al. A Holistic Platform for Accelerating Sorbent-Based Carbon Capture. *Nature* **2024**, *632*, 89–94. https://doi.org/10.1038/s41586-024-07683-8.
58. Parveen, F.; Slater, A.G. Digitalisation of Catalytic Processes for Sustainable Production of Biobased Chemicals and Exploration of Wider Chemical Space. *Catal. Sci. Technol.* **2025**, *15*, 1689–1701. https://doi.org/10.1039/D4CY01525H.
59. Wang, C.; Lai, H.; Warkentin, H.; et al. Catalysts for Electrochemical CO_2 Conversion. Material Sustainability Perspective. *NPJ Mater. Sustain.* **2025**, *3*, 22. https://doi.org/10.1038/s44296-025-00065-9.
60. Xun, D.; Liu, M.; Hao, H.; et al. Sustainable Supply of Critical Materials for Water Electrolysers and Fuel Cells. *Commun. Earth Environ.* **2025**, *6*, 627. https://doi.org/10.1038/s43247-025-02621-6.
61. Chen, X.; El Khatib, M.; Lindgren, P.; et al. Atomistic Learning in the Electronically Grand-Canonical Ensemble. *NPJ Comput. Mater.* **2023**, *9*, 73. https://doi.org/10.1038/s41524-023-01007-6.
62. Zhuang, Y.-B.; Liu, C.; Zhu, J.-X.; et al. An Artificial Intelligence Accelerated Ab Initio Molecular Dynamics Dataset for Electrochemical Interfaces. *Sci. Data* **2025**, *12*, 997. https://doi.org/10.1038/s41597-025-05338-5.
63. Erlebach, A.; Šípka, M.; Saha, I.; et al. A Reactive Neural Network Framework for Water-Loaded Acidic Zeolites. *Nat. Commun.* **2024**, *15*, 4215. https://doi.org/10.1038/s41467-024-48609-2.
64. Omranpour, A.; Elsner, J.; Lausch, K.N.; et al. Machine Learning Potentials for Heterogeneous Catalysis. *ACS Catal.* **2025**, *15*, 1616–1634. https://doi.org/10.1021/acscatal.4c06717.
65. Tinajero, C.; Zanatta, M.; Sánchez-Velandia, J.E.; et al. Reac-Discovery. An Artificial Intelligence-Driven Platform for Continuous-Flow Catalytic Reactor Discovery and Optimization. *Nat. Commun.* **2025**, *16*, 9062. https://doi.org/10.1038/s41467-025-64127-1.

66. Bennett, J.A.; Orouji, N.; Khan, M.B.; et al. Autonomous Reaction Pareto-Front Mapping with a Self-Driving Catalysis Laboratory. *Nat. Chem. Eng.* **2024**, *1*, 240–250. https://doi.org/10.1038/s44286-024-00033-5.
67. Zhang, C.; Lu, R.; Sun, Q.; et al. Finding the Pareto Front for High-Entropy-Alloy Catalysts. *Chem. Sci.* **2026**, *17*, 4744–4752. https://doi.org/10.1039/D5SC06100H.
68. Kim, J.; Kwon, I.S.; Lim, J.; et al. Machine-Learning-Guided Tungsten Single Atoms Promote Oxyhydroxides for Noble-Metal-Free Water Electrolysis. *Nat. Commun.* **2026**, *17*, 2344. https://doi.org/10.1038/s41467-026-68735-3.
69. Casillo, E.; Scattolin, T.; Nolan, S.P. Catalysis Meets Machine Learning. A Guide to Data-Driven Discovery and Design. *Chem. Commun.* **2025**, *61*, 18247–18272. https://doi.org/10.1039/D5CC05274B.
70. Tran, K.; Ulissi, Z.W. Active Learning across Intermetallics to Guide Discovery of Electrocatalysts for CO_2 Reduction and H_2 Evolution. *Nat. Catal.* **2018**, *1*, 696–703. https://doi.org/10.1038/s41929-018-0142-1.
71. Mok, D.H.; Li, H.; Zhang, G.; et al. Data-Driven Discovery of Electrocatalysts for CO_2 Reduction Using Active Motifs-Based Machine Learning. *Nat. Commun.* **2023**, *14*, 7303. https://doi.org/10.1038/s41467-023-43118-0.
72. Wang, G.; Mine, S.; Chen, D.; et al. Accelerated Discovery of Multi-Elemental Reverse Water-Gas Shift Catalysts Using Extrapolative Machine Learning Approach. *Nat. Commun.* **2023**, *14*, 5861. https://doi.org/10.1038/s41467-023-41341-3.
73. Zhong, M.; Tran, K.; Min, Y.; et al. Accelerated Discovery of CO_2 Electrocatalysts Using Active Machine Learning. *Nature* **2020**, *581*, 178–183. https://doi.org/10.1038/s41586-020-2242-8.
74. Song, Z.; Fan, L.; Lu, S.; et al. Inverse Design of Promising Electrocatalysts for CO_2 Reduction via Generative Models and Bird Swarm Algorithm. *Nat. Commun.* **2025**, *16*, 1053. https://doi.org/10.1038/s41467-024-55613-z.
75. Cui, Z.; Aztergo, K.D.; Hwang, J.; et al. Determining CO Adsorption Free Energies on CO_2 Electroreduction Active Sites through Kinetic Analysis. *Nat. Catal.* **2025**, *8*, 1117–1127. https://doi.org/10.1038/s41929-025-01427-1.
76. Zhan, C.; Dattila, F.; Rettenmaier, C.; et al. Key Intermediates and Cu Active Sites for CO_2 Electroreduction to Ethylene and Ethanol. *Nat. Energy* **2024**, *9*, 1485–1496. https://doi.org/10.1038/s41560-024-01633-4.
77. Zhang, Z.; Li, H.; Shao, Y.; et al. Molecular Understanding of the Critical Role of Alkali Metal Cations in Initiating CO_2 Electroreduction on Cu(100) Surface. *Nat. Commun.* **2024**, *15*, 612. https://doi.org/10.1038/s41467-024-44896-x.
78. Chen, J.; Qiu, H.; Zhao, Y.; et al. Selective and Stable CO_2 Electroreduction at High Rates via Control of Local H_2O to CO_2 Ratio. *Nat. Commun.* **2024**, *15*, 5893. https://doi.org/10.1038/s41467-024-50269-1.
79. Gao, X.; Jiang, Y.; Liu, J.; et al. Intermediate-Regulated Dynamic Restructuring at Ag-Cu Biphasic Interface Enables Selective CO_2 Electroreduction to C_{2+} Fuels. *Nat. Commun.* **2024**, *15*, 10331. https://doi.org/10.1038/s41467-024-54630-2.
80. Cheng, D.; Nguyen, K.L.C.; Sumaria, V.; et al. Structure Sensitivity and Catalyst Restructuring for CO_2 Electro-Reduction on Copper. *Nat. Commun.* **2025**, *16*, 4064. https://doi.org/10.1038/s41467-025-59267-3.
81. Knöppel, J.; Möckl, M.; Escalera-López, D.; et al. On the Limitations in Assessing Stability of Oxygen Evolution Catalysts Using Aqueous Model Electrochemical Cells. *Nat. Commun.* **2021**, *12*, 2231. https://doi.org/10.1038/s41467-021-22296-9.
82. Tao, H.B.; Liu, H.; Lao, K.; et al. The Gap between Academic Research on Proton Exchange Membrane Water Electrolysers and Industrial Demands. *Nat. Nanotechnol.* **2024**, *19*, 1074–1076. https://doi.org/10.1038/s41565-024-01699-x.
83. Wang, C.R.; Stansberry, J.M.; Mukundan, R.; et al. Proton Exchange Membrane Water Electrolysis. Cell-Level Considerations for Gigawatt-Scale Deployment. *Chem. Rev.* **2025**, *125*, 1257–1302. https://doi.org/10.1021/acs.chemrev.3c00904.
84. Park, W.; Noh, J.; Gu, G.H.; et al. Machine Learning-Enabled Fast Exploration of Stable and Active Single-Atom Catalysts for Oxygen Evolution Reaction. *Innov. Mater.* **2024**, *2*, 100072. https://doi.org/10.59717/j.xinn-mater.2024.100072.
85. Abed, J.; Heras-Domingo, J.; Sanspeur, R.Y.; et al. Pourbaix Machine Learning Framework Identifies Acidic Water Oxidation Catalysts Exhibiting Suppressed Ruthenium Dissolution. *J. Am. Chem. Soc.* **2024**, *146*, 15740–15750. https://doi.org/10.1021/jacs.4c01353.
86. Wu, Z.-Y.; Chen, F.-Y.; Li, B.; et al. Non-Iridium-Based Electrocatalyst for Durable Acidic Oxygen Evolution Reaction in Proton Exchange Membrane Water Electrolysis. *Nat. Mater.* **2023**, *22*, 100–108. https://doi.org/10.1038/s41563-022-01380-5.
87. Cao, X.; Miao, L.; Jia, W.; et al. Strain Heterogeneity in RuO_2 for Efficient Acidic Oxygen Evolution Reaction in Proton Exchange Membrane Water Electrolysis. *Nat. Commun.* **2025**, *16*, 6217. https://doi.org/10.1038/s41467-025-58570-3.
88. Qian, F.; Cao, D.; Chen, S.; et al. High-Entropy RuO_2 Catalyst with Dual-Site Oxide Path for Durable Acidic Oxygen Evolution Reaction. *Nat. Commun.* **2025**, *16*, 6894. https://doi.org/10.1038/s41467-025-61763-5.
89. Zhang, K.; Liang, X.; Wang, Y.; et al. Support-Tuned Iridium Reconstruction with Crystalline Phase Dominating Acidic Oxygen Evolution. *Nat. Commun.* **2025**, *16*, 8164. https://doi.org/10.1038/s41467-025-63541-9.

90. Yang, D.; Zhang, C.; Qin, Y.; et al. Single-Faceted IrO_2 Monolayer Enabling High-Performing Proton Exchange Membrane Water Electrolysis beyond 10,000 h Stability at 1.5 A cm^{-2}. *Nat. Commun.* **2025**, *16*, 7236. https://doi.org/10.1038/s41467-025-62665-2.
91. Qiu, C.; Sellers, C.; Wu, Z.-Y.; et al. Low-Iridium Stabilized Ruthenium Oxide Anode Catalyst for Durable Proton-Exchange Membrane Water Electrolysis. *Nat. Nanotechnol.* **2025**, *20*, 1787–1795. https://doi.org/10.1038/s41565-025-02030-y.
92. Suvarna, M.; Preikschas, P.; Pérez-Ramírez, J. Identifying Descriptors for Promoted Rhodium-Based Catalysts for Higher Alcohol Synthesis via Machine Learning. *ACS Catal.* **2022**, *12*, 15373–15385. https://doi.org/10.1021/acscatal.2c04349.
93. Pisal, P.; Krejčí, O.; Rinke, P. Machine Learning Accelerated Descriptor Design for Catalyst Discovery in CO_2 to Methanol Conversion. *NPJ Comput. Mater.* **2025**, *11*, 213. https://doi.org/10.1038/s41524-025-01664-9.
94. Ramirez, A.; Lam, E.; Gutierrez, D.P.; et al. Accelerated Exploration of Heterogeneous CO_2 Hydrogenation Catalysts by Bayesian-Optimized High-Throughput and Automated Experimentation. *Chem Catal.* **2024**, *4*, 100888. https://doi.org/10.1016/j.checat.2023.100888.
95. Ge, Y.; Martín, A.J.; Suvarna, M.; et al. Descriptors toward Higher Alcohols and Olefins Formation in CO_2 Hydrogenation on Promoted Iron Catalysts. *Appl. Catal. B Environ. Energy* **2025**, *378*, 125578. https://doi.org/10.1016/j.apcatb.2025.125578.
96. Grasselli, F.; Chong, S.; Kapil, V.; et al. Uncertainty in the Era of Machine Learning for Atomistic Modeling. *Digit. Discov.* **2025**, *4*, 2654–2675. https://doi.org/10.1039/D5DD00102A.
97. Bilbrey, J.A.; Firoz, J.S.; Lee, M.-S.; et al. Uncertainty Quantification for Neural Network Potential Foundation Models. *NPJ Comput. Mater.* **2025**, *11*, 109. https://doi.org/10.1038/s41524-025-01572-y.
98. Nair, A.S.; Foppa, L.; Scheffler, M. Materials-Discovery Workflow Guided by Symbolic Regression for Identifying Acid-Stable Oxides for Electrocatalysis. *NPJ Comput. Mater.* **2025**, *11*, 150. https://doi.org/10.1038/s41524-025-01596-4.