Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Autonomous materials discovery through a tabula rasa active learning framework

Chih-Yang Cheng

mushroom2368@gmail.com

National Chung Hsing University

Hsin-Ying Ting

National Chung Hsing University

Yi-Huan Wu

R.O.C Military Academy

Feng-Yin Li

National Chung Hsing University

Perspective

Keywords:

Posted Date: May 27th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9750606/v1

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

Additional Declarations: No competing interests reported.

Autonomous materials discovery through a tabula rasa active learning framework

Chih-Yang Cheng[1*], Hsin-Ying Ting[1], Yi-Huan Wu[2*], Feng-Yin Li[1*]

[1]Department of Chemistry, National Chung Hsing University, Taichung 402, Taiwan.
[2]Department of Chemistry, R. O. C. Military Academy, Kaohsiung, Taiwan.
[*]corresponding author: Chih-Yang Cheng (mushroom2368@gmail.com), Yi-Huan Wu (m1090008@rocma.edu.tw), Feng-Yin Li (feng64@nchu.edu.tw)

Abstract

Developing high-dimensional multi-element alloy electrocatalysts is crucial for clean energy, yet the discovery process is fundamentally hindered by the combinatorial explosion of the vast chemical space. Traditional trial-and-error methodologies are prohibitively slow, while current active learning approaches heavily rely on human-engineered features and frequently fall into the exploitation trap of localized optima. Here, we show an autonomous materials discovery framework operating from a tabula rasa state by coupling a descriptor-free surrogate model (PC-BAN) with a neuroscience-inspired acquisition function (LCBDS) that enforces surprise-driven exploration. Within a design space of 6074 quaternary combinations, our framework autonomously identified the absolute global optimum catalyst exhibiting an exceptional overpotential of 206 mV using only 220 queries to an experimental oracle, substantially outperforming traditional baselines and human-guided searches. By mathematically formalizing curiosity without prior chemical biases, this work lays a scalable, descriptor-free algorithmic foundation for the next generation of prospective autonomous laboratories.

1 Introduction

Finding optimal multi-element alloys in the vast chemical space is fundamentally constrained by the combinatorial explosion of possibilities and high experimental costs[1]. While transitioning from unary or binary systems to high-dimensional compositions like quaternary alloys provides a wealth of potential active sites and synergistic effects[2], navigating this expanded design space is computationally and experimentally prohibitive[3]. Conventional trial-and-error approaches are inadequate for such high-dimensional optimization because synthesizing and characterizing every possible candidate demands resources that scale exponentially[4].

To accelerate this process, machine learning and active learning have been widely adopted to predict catalytic activity and guide experiments[3-6]. However, these methods suffer from critical methodological limitations. First, traditional predictive models heavily rely on human-engineered chemical features and domain-specific descriptors[7]. This dependence restricts the models to existing human biases and strictly limits the potential for discovering entirely novel knowledge outside established paradigms. Second, although active learning can reduce overall data requirements, conventional acquisition functions, when navigating high-dimensional spaces, often fall into an exploitation trap[8]. These traditional algorithms frequently succumb to over-exploiting known low-value regions and lack a robust mechanism to escape local optima[7, 8].

To resolve these fundamental bottlenecks, we propose an autonomous materials discovery framework utilizing a tabula rasa active learning approach (Fig. 1). Specifically, we introduce a descriptor-free surrogate model (PC-BAN), which operates entirely without a priori human-engineered features or chemical descriptors. We couple this with a neuroscience-inspired algorithm termed LCBDS, which incorporates a diversity penalty and a surprise-driven curiosity mechanism. Instead of greedily seeking minor improvements in already heavily sampled domains, the LCBDS algorithm converts prediction errors into a surprise signal. This signal actively guides the model to explore unknown anomalous regions within the vast chemical space, mathematically overcoming the exploitation trap while maintaining a balance between exploration and optimal solution identification.

The efficacy of this framework is demonstrated through the accelerated discovery of high-performance oxygen evolution reaction (OER) electrocatalysts[9, 10]. In a screening space of 6074 quaternary combinations, the framework autonomously locked onto the optimal

composition using only 220 physical experiments. The discovered catalyst achieves an exceptional, near-global optimum overpotential of 206 mV, substantially outperforming traditional baselines and human-expert-guided blind searches. By completely removing the reliance on predefined features and mitigating the local exploitation trap, this work does not merely identify an exceptional catalyst. Rather, it establishes a highly efficient and universally applicable paradigm for autonomous materials discovery and complex process optimization.

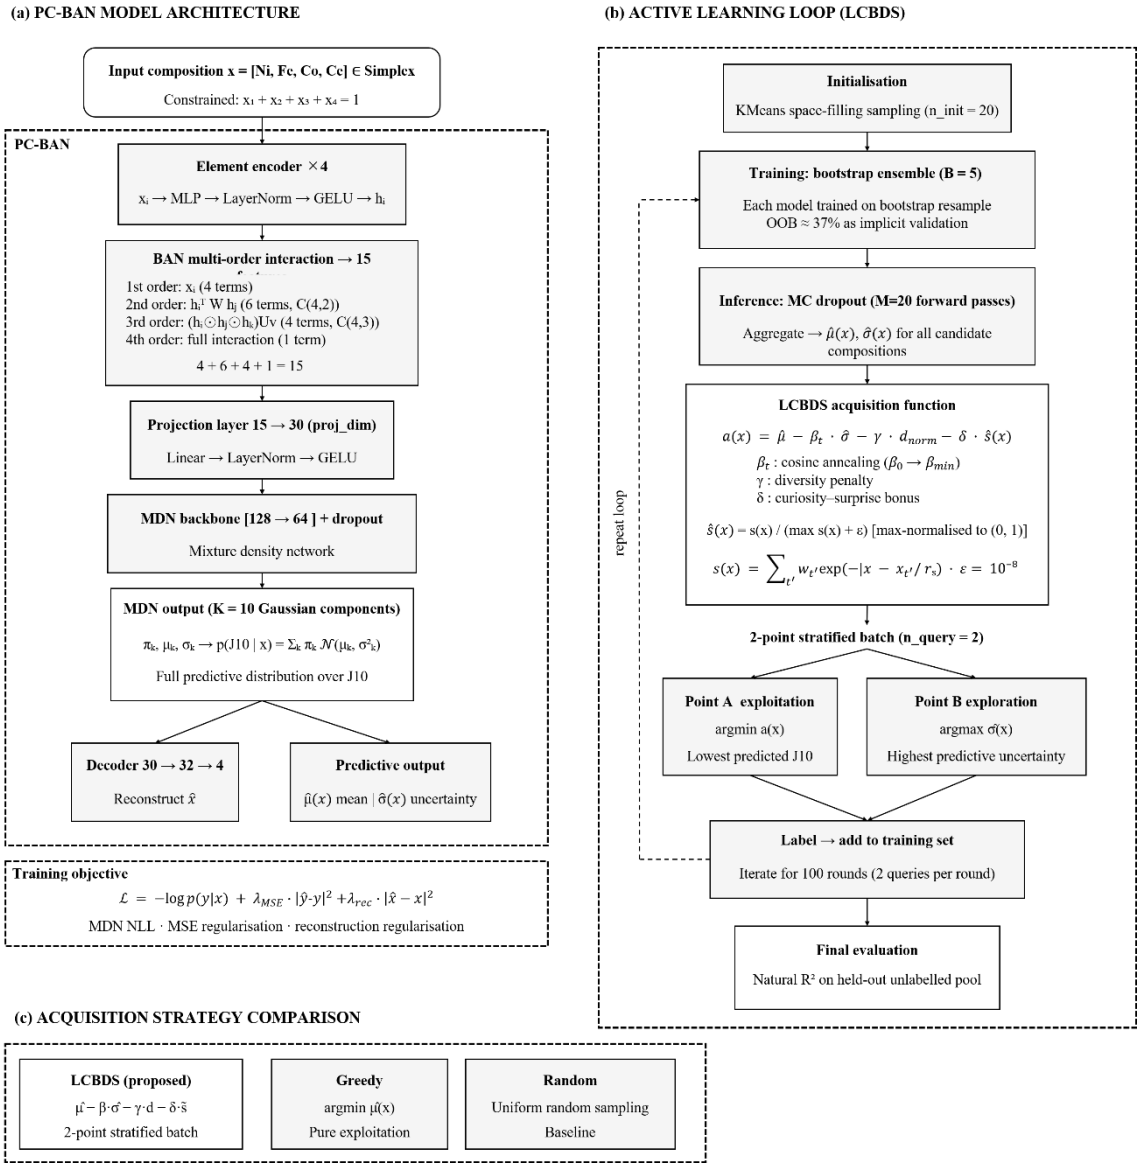


Fig. 1 | Architecture of the tabula rasa PC-BAN surrogate model and the closed-loop LCBDS active learning framework. The fully descriptor-free framework bridges direct compositional inputs to autonomous experimental queries without reliance on human-

engineered features. (a) PC-BAN model architecture. The network maps raw quaternary simplex fractions through independent element encoders and a multi-order interaction module, capturing up to 4-body synergistic effects. A Mixture Density Network (MDN) backbone ([128, 64]) outputs a predictive Gaussian mixture distribution (K=10), while an auxiliary reconstruction decoder regularizes the latent space. (b) The active learning loop. The cycle utilizes a dual uncertainty estimation mechanism combining a bootstrap ensemble (B=5) with MC Dropout (M=20). The LCBDS acquisition function evaluates the candidate pool by balancing exploitation ($\hat{\mu}$) and exploration ($\hat{\sigma}$), while enforcing spatial diversity (γ) and anomaly-driven curiosity (δ). The framework executes a two-point stratified batch query (Points A and B) over 100 iterations to dynamically navigate the complex catalytic landscape. (c) Acquisition strategy comparison, benchmarking the proposed LCBDS against purely exploitative (Greedy) and passive (Random) baseline strategies.

2 Methods

2.1 Chemical Space Formulation and Problem Definition

The exploration boundary for our autonomous materials discovery framework is defined by a quaternary alloy system comprising nickel (Ni), iron (Fe), cobalt (Co), and cerium (Ce). We represent each alloy composition as a molar fraction vector x, defined within the standard 3-simplex (Δ^3):

$$\mathrm{x} = [x_{Ni}, x_{Fe}, x_{Co}, x_{Ce}]^T \in \Delta^3, \ \ subject\ to \sum_{i=1}^{4} x_i = 1, \quad x_i \geq 0 \qquad Eq(1)$$

To rigorously evaluate the autonomous discovery efficiency of our framework, we utilized an exhaustive experimental database of (Ni-Fe-Co-Ce)Ox quaternary alloys as an experimental oracle. This dataset, originally reported by Gregoire and co-workers[11], comprises high-throughput electrochemical measurements (including J10 overpotentials) for 6074 unique compositions. By querying this pre-established oracle, we simulate the closed-loop feedback of a physical laboratory while maintaining a fixed ground-truth landscape, allowing for a reproducible and objective benchmarking of active learning strategies in a high-dimensional chemical space.

To systematically navigate this continuous compositional simplex, the domain was discretized to generate a comprehensive candidate pool U consisting of N=6074 unique, unlabeled compositional combinations. This discretization provides a sufficiently fine elemental resolution to capture complex synergistic effects between transition metals and rare-earth elements while remaining computationally and experimentally tractable for high-throughput screening.

The primary physical optimization target within this chemical space is to identify the specific composition that minimizes the catalytic overpotential at a current density of 10 mA cm^{-2}, denoted as $J10$. Therefore, the core machine learning problem is mathematically formulated as learning a probabilistic surrogate model $p(J10 \mid \mathbf{x})$ from an initially small labeled set. The ground-truth labeled data are physically acquired through a closed-loop experimental setup, where high-throughput synthesis techniques are coupled with standardized electrochemical measurements to accurately evaluate the $J10$ of queried candidates. By continuously updating the probabilistic surrogate model $p(J10 \mid \mathbf{x})$ with these newly acquired physical measurements, the AI framework can

autonomously evaluate the vast uncharted regions of the 6074 candidates and guide subsequent experimental queries toward the global optimum.

2.2 Descriptor-Free Surrogate Model (PC-BAN)

To completely eliminate the reliance on human-engineered features and domain-specific descriptors, we designed the Process Community Bilinear Attention Network (PC-BAN) as a *tabula rasa* surrogate model. The architecture processes the raw compositional data through a sequential pipeline of independent encoding, multi-order interaction, probabilistic output generation, and latent space regularization.

2.2.1 Element Encoder

The model directly ingests the molar fraction vector **x**. Instead of using pre-calculated physical or chemical properties, each scalar composition x_i is independently mapped to a latent embedding $\mathrm{h_i} \in \mathbb{R}^{\mathrm{D}}$. This is achieved through an encoder network that shares the same structure across all elements but is independently parameterized. This independent mapping mechanism ensures that the fundamental representation of each element is learned purely from the data:

$$h_i = Linear_2\left(GELU\left(LayerNorm\left(Linear_1(x_i)\right)\right)\right) \in \mathbb{R}^{\mathrm{D}} \qquad Eq(2)$$

where the embedding dimension is defined as D=32.

2.2.2 Multi-Order Interaction (BAN)

To capture the complex compositional synergies inherently present in quaternary alloys without any domain-specific feature engineering, we construct a Bilinear Attention Network (BAN). This module explicitly computes 15 interaction features spanning from the first to the fourth order:

$$f_i^{(1)} = x_i \ (4\ terms) \qquad Eq(3)$$

$$f_{ij}^{(2)} = h_i^T W_{ij} h_j + b_{ij} \ \left(\binom{4}{2} = 6\ terms\right) \qquad Eq(4)$$

$$f_{ijk}^{(3)} = \left(h_i \odot h_j \odot h_k\right)^T U_{ijk} v_{ijk} + b_{ijk} \ \left(\binom{4}{3} = 4\ terms\right) \qquad Eq(5)$$

$$f_{1234}^{(4)} = (h_1 \odot h_2 \odot h_3 \odot h_4)^T Uv + b \ (1 \ terms) \qquad Eq(6)$$

The concatenated feature vector $f = \left[f^{(1)}, f^{(2)}, f^{(3)}, f^{(4)}\right] \in \mathbb{R}^{15}$ fuses all single-element and synergistic effects. This vector is subsequently passed through a projection layer to obtain a condensed latent representation **z**:

$$z = GELU\left(LayerNorm\left(W_{proj} f + b_{proj}\right)\right) \in \mathbb{R}^{d_z} \qquad Eq(7)$$

where the projection dimension is set to $d_z = 30$.

2.2.3 MDN Output

To accurately quantify the predictive uncertainty required for active learning, the projected representation **z** is fed into a two-layer backbone network configured with LayerNorm, GELU activation, and Monte Carlo (MC) Dropout[12]. This backbone generates the parameter inputs for a *K*-component Mixture Density Network (MDN)[13]:

$$h_{bb} = Backbone(z; \theta) \qquad Eq(8)$$

The MDN constructs a mixture Gaussian distribution to model the complex target space probabilistically. The network outputs the mixture weights $\boldsymbol{\pi}$, means $\boldsymbol{\mu}$, and standard deviations $\boldsymbol{\sigma}$, defining the conditional probability $p(J10 \mid \mathbf{x})$:

$$\pi = Softmax(W_\pi h_{bb}), \mu = W_\mu h_{bb}, \sigma = Softplus(W_\sigma h_{bb}) + \epsilon$$

$$p(J_{10}|x) = \sum_{k=1}^{K} \pi_k \mathcal{N}(J_{10}; \mu_k, \sigma_k^2) \qquad Eq(9)$$

where K=10 components are utilized, and ϵ=10^{-4} ensures numerical stability. From this Gaussian mixture distribution, the overall predictive mean $\hat{\mu}(x)$ and predictive variance $\hat{\sigma}^2(x)$ are analytically computed as:

$$\hat{\mu}(x) = \sum_k \pi_k \mu_k, \quad \hat{\sigma}^2(x) \sum_k \pi_k (\sigma_k^2 + \mu_k^2) - \hat{\mu}^2(x) \qquad Eq(10)$$

2.2.4 Reconstruction Decoder

To regularize the latent space and ensure that the learned features remain compositionally interpretable, an auxiliary reconstruction decoder is integrated into the architecture. This mechanism acts as a critical regularizer by forcing the latent vector **z** to reconstruct the original input composition vector **x**, preventing the neural network from collapsing into unphysical or purely mathematical abstractions:

$$\hat{x} = Softmax\big(W_2 GELU(W_1 z)\big)\, \epsilon\, \Delta^3 \qquad Eq(11)$$

2.2.5 Training Objective

The PC-BAN model is optimized end-to-end by minimizing a composite loss function L. This objective comprehensively combines the negative log-likelihood (NLL) for probabilistic distribution calibration, the mean squared error (MSE) for prediction accuracy, and the reconstruction loss for latent space regularization:

$$\mathcal{L} = \underbrace{-logp(J10|x)}_{NLL} + \lambda_{MSE}\, \underbrace{|\hat{\mu}(x) - J10|^2}_{MSE} + \lambda_{rec}\ \underbrace{|\hat{x} - x|^2}_{Reconstruction} \qquad Eq(12)$$

The weighting hyperparameters are empirically set to $\lambda_{MSE} = 0.10$ and $\lambda_{rec} = 0.05$ to balance the gradient contributions across the three distinct optimization targets. The full model is trained end-to-end using the Adam optimizer[14].

2.2.6 Dynamic Capacity Scaling and Knowledge Distillation

It is important to note that the core neural network training protocols (e.g., AdamW optimization, CosineAnnealing learning rate schedule, and standard dropout) adhere strictly to standard deep learning practices rather than being heavily fine-tuned for this specific dataset. However, to explicitly prevent catastrophic overfitting in the extreme data-scarce regime of early active learning iterations (e.g., n=20), we implemented a progressive capacity scaling strategy (Adaptive Tiers, Supplementary Table 1). The network's degrees of freedom specifically the MDN components *K* and hidden dimensions are dynamically expanded only as the physical dataset geometrically grows. This ensures that the model capacity always remains commensurate with the available data volume, thereby producing well-calibrated epistemic uncertainties. Furthermore, during tier upgrades and warm-start iterations, a self-knowledge distillation mechanism ($\alpha = 0.5$) is applied. Instead of discarding the prior surrogate's understanding upon

retraining, this mechanism forces the updated model to retain the global physicochemical landscape learned from previous iterations, effectively mitigating catastrophic forgetting.

2.3 Epistemic Uncertainty Quantification

In the emerging paradigm of AI for Science, the success of an active learning framework depends not only on the predictive accuracy of the surrogate model but also fundamentally on its ability to quantify its own ignorance. Knowing what the model does not know, formally termed epistemic uncertainty, is the critical driver that prevents the system from falling into local exploitation traps and guides it toward uncharted, high-value chemical spaces. To rigorously quantify this epistemic uncertainty, we implemented a dual uncertainty estimation mechanism combining Bootstrap Ensembling[15] and Monte Carlo (MC) Dropout[12].

At each active learning iteration t, an ensemble of B=5 independent PC-BAN models is trained. Each model θ_b is optimized on a bootstrap resample $\mathcal{L}_t^{(b)}$ drawn randomly with replacement from the currently available labeled set $\mathcal{L}_t$. To further capture the uncertainty intrinsic to the neural network weights, each of the B bootstrap models is evaluated M=20 times with dropout layers enabled during the inference phase. This combined approach yields a total of $B \times M$=100 stochastic forward passes for every unlabeled candidate composition **x**. Let $\hat{y}_{b,m}(x)$ denote the output of the m-th stochastic pass from the b-th bootstrap model. The overall predictive mean $\hat{\mu}(x)$ and the predictive uncertainty $\hat{\sigma}(x)$ are subsequently aggregated as:

$$\hat{\mu}(x) = \frac{1}{BM}\sum_{b=1}^{B}\sum_{m=1}^{M}\hat{y}_{b,m}(x) \qquad Eq(13)$$

$$\hat{\sigma}(x) = std\{\hat{y}_{b,m}(x)\}_{b,m} \qquad Eq(14)$$

In traditional machine learning workflows, a fixed holdout set is typically used for performance validation. However, in the data-scarce environment of autonomous materials discovery, withholding expensive physical experimental data severely limits model training. The bootstrap ensemble naturally resolves this dilemma through the out-of-bag (OOB) samples. Statistically, approximately 37% of the data in $\mathcal{L}_t$ is left out of any single bootstrap resample $\mathcal{L}_t^{(b)}$. These OOB samples provide a robust mechanism for

implicit validation. We dynamically assess the generalization capability of the framework at each iteration by computing the OOB coefficient of determination (ROOB2):

$$R_{OOB}^2 = 1 - \frac{\sum_{i \in OOB} \left(y_i - \bar{y}_{OOB,i}\right)^2}{\sum_{i \in OOB} (y_i - \bar{y})^2} \qquad Eq(15)$$

where $\bar{y}_{OOB,i}$ represents the predictive mean for the i-th sample, calculated exclusively using the subset of ensemble models that did not include this specific sample in their respective training sets, and $\bar{y}$ is the mean of the true experimental labels. This implicit validation ensures that the model's predictive confidence is rigorously and continuously calibrated without sacrificing any physical experimental data.

2.4 Autonomous Discovery via LCBDS Acquisition

We adopt a pool-based active learning framework[16] in which a candidate pool $\mathcal{U}$ of N=6074 unlabeled compositions is evaluated sequentially to minimize the number of physical experiments required to identify high-performing catalysts. At any given iteration t, the set of physically measured compositions is denoted as $\mathcal{L}_t$. To establish an unbiased and spatially diverse foundation for the surrogate model, the initial labeled set $\mathcal{L}_0$ consisting of 20 points is constructed using a K-means space-filling technique. This method partitions the candidate pool into 20 distinct clusters and draws one representative composition from each cluster:

$$\mathcal{L}_0 = \left\{x^{(c)}\right\}_{c=1}^{20}, \quad x^{(c)} \sim Uniform(\{x : cluster(x) = c\}) \qquad Eq(16)$$

Conventional acquisition functions often succumb to the exploitation trap when navigating high-dimensional chemical spaces[17, 18]. To counteract this, we propose the Lower Confidence Bound[19] with Diversity and Surprise (LCBDS) acquisition function. This function acts as the central decision-making brain of the framework, strategically balancing exploitation of optimal values, exploration of uncertain regions, spatial diversity, and a neuroscience-inspired curiosity mechanism:

$$a(x) = \hat{\mu}(x) - \beta_t \hat{\sigma}(x) - \gamma d_{norm}(x) - \delta_s(x) \qquad Eq(17)$$

where $\hat{\mu}(x)$ and $\hat{\sigma}(x)$ are the predictive mean and epistemic uncertainty derived from the PC-BAN ensemble, respectively.

The diversity term, $d_{norm}(x)$, acts as a spatial penalty designed to prevent redundant sampling in heavily explored localized regions[20]. It normalizes the minimum distance between a candidate **x** and all currently labeled points in $\mathcal{L}_t$, scaled by the maximum possible distance within the pool:

$$d_{norm}(x) = \frac{min_{x' \in \mathcal{L}_t}|x - x'|}{max_{x'' \in \mathcal{U}_i} min_{x' \in \mathcal{L}_t}|x'' - x'|} \in [0,1] \qquad Eq(18)$$

This penalty inherently promotes a broader exploration of the compositional simplex and ensures structural diversity in the selected batches.

To mathematically formalize curiosity and avoid local optima, we introduce the surprise bonus, $s(\mathbf{x})$. During the active learning loop, if the model encounters a physical measurement that deviates from its prediction by more than 2σ, this event is recorded as a "surprise." The algorithm translates these substantial prediction errors into a surprise signal, actively guiding the model to investigate unknown anomalous regions that might harbor novel synergistic effects:

$$s(x) = \frac{\sum_{t'} w_{t'} \exp\left(-\frac{|x - x_{t'}|}{r_s}\right)}{max_{x''} \sum_{t'} w_{t'} \exp\left(-\frac{|x'' - x_{t'}|}{r_s}\right)} \in [0,1] \qquad Eq(19)$$

In this formulation, $w_{t'}$ represents the residual magnitude (in mV) at the surprise event t', and $r_s = 0.1$ defines the spatial decay radius.

The dynamic balance between exploring highly uncertain regions and exploiting known low-overpotential domains is governed by the trade-off parameter β_t. Rather than utilizing a static threshold, β_t is dynamically annealed from an exploration-heavy initial state to an exploitation-focused final state via a cosine decay schedule[21]:

$$\beta_t = \beta_{min} + (\beta_0 - \beta_{min}) \cos\left(\frac{\pi t}{2T}\right) \qquad Eq(20)$$

where the hyperparameters are empirically set to $\beta_0 = 2.5$, $\beta_{min} = 0.2$, and the total number of planned iterations is T=100. This scheduled annealing guarantees that the framework thoroughly maps the boundaries of the chemical space in early stages before aggressively converging on the global optimum in later iterations.

2.5 Active Learning Workflow and Baseline Evaluation

To implement the closed-loop autonomous discovery, we designed a highly efficient active learning workflow coupled with a stratified batch query strategy. At each active learning iteration t, the framework systematically selects exactly two distinct candidate compositions from the unlabeled pool $\mathcal{U}_t$ for physical synthesis and electrochemical characterization. This dual-point batch is strategically partitioned to fulfill complementary optimization objectives. Specifically, Point A focuses on robust exploitation by querying the candidate with the minimum LCBDS acquisition score, thereby aggressively refining the currently known high-performing domains. Conversely, Point B is dedicated entirely to exploration by selecting the candidate with the maximum epistemic uncertainty $\hat{\sigma}(x)$ among the remaining pool. This division ensures continuous expansion of the model's knowledge boundaries and actively prevents the framework from experiencing mode collapse into local optima. This iterative cycle proceeds until a strict stopping condition of 100 consecutive rounds is met. When combined with the 20 initial space-filling points, the entire workflow operates under a highly constrained experimental budget of exactly 220 physical data points.

To rigorously demonstrate the superiority of the proposed LCBDS-driven framework, its performance is benchmarked against two standard acquisition strategies evaluated under the identical experimental budget constraint. The first control is a Greedy baseline, which represents a purely exploitative approach. The Greedy strategy strictly selects candidates based solely on the lowest predicted mean $\hat{\mu}(x)$, deliberately ignoring predictive uncertainty, spatial diversity, and curiosity-driven exploration. The second control is a Random baseline, which uniformly selects candidates at random from the available pool, mathematically simulating a traditional, undirected high-throughput trial and error screening process.

Furthermore, to quantitatively assess the discovery efficiency and the ultimate optimization capability of these distinct strategies, we establish two core evaluation metrics. The first metric, designated as $break_{exp}$, measures the exact number of experimental queries required for a strategy to first successfully identify a catalyst

composition that breaches a predefined high-performance overpotential threshold. This metric explicitly evaluates the acceleration speed of the discovery phase. The second metric, $best_{mV}$, is defined as the lowest experimental overpotential achieved by a given strategy upon the complete exhaustion of the 220-experiment budget. Together, these dual metrics provide a comprehensive and objective evaluation of the framework, definitively proving its capacity to rapidly identify and reliably converge upon the global optimal catalyst within the vast chemical space.

3 Results

3.1 Autonomous Discovery Performance and Surrogate Evolution

The autonomous discovery performance of the proposed framework is first evaluated under a strictly constrained budget of 220 physical experiments. As illustrated in Fig. 2a, the optimization trajectories reveal a profound trade-off between discovery speed, stability, and surrogate model integrity. The Greedy strategy (red dashed line), which myopically prioritizes immediate performance gains, exhibits highly erratic and high-variance behavior. While its cross-seed mean appears to rapidly breach the 370 mV threshold and reach lower overpotentials, this is heavily driven by extreme "lucky" runs within specific random seeds, as indicated by its massively inflated standard deviation (shaded region).

Critically, this reckless exploitation comes at a catastrophic cost to the artificial intelligence's global understanding. As quantitatively tracked in Fig. 2b, the Greedy strategy causes the bootstrap out-of-bag R^2 to collapse dramatically to approximately 0.4. By exclusively exploiting localized minimums, the model creates severe blind spots and entirely loses its generalization capability.

In stark contrast, the baseline LCBDS strategy ($n_q = 2$, blue solid line) fundamentally prioritizes cognitive stability and consistent spatial mapping over reckless exploitation. While its optimization trajectory in Fig. 2a is more conservative steadily approaching the 370 mV threshold with minimal variance it successfully maintains an exceptionally robust global surrogate model, with its R^2 firmly stabilizing above 0.8 (Fig. 2b). This proves that the proposed diversity and curiosity mechanisms effectively protect the framework from premature model collapse. Furthermore, this stable baseline serves as the indispensable foundation for advanced querying tactics. When augmented with a third exploration query (LCBDS+Maxμ, purple line), the framework harnesses this stable cognitive foundation to trigger aggressive late-stage optimization, ultimately bypassing all baselines to reach deep optimal domains.

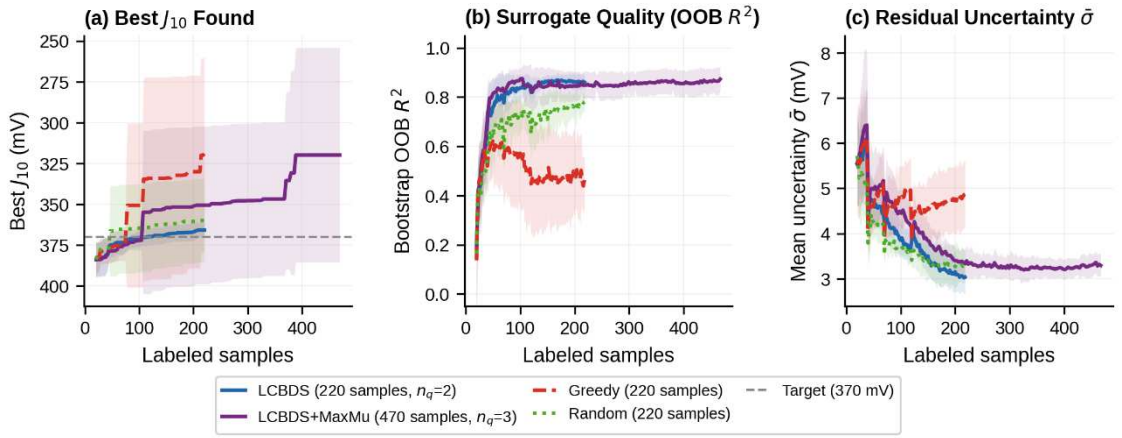


Fig. 2 | Autonomous discovery performance and surrogate evolution. Active-learning trajectories are evaluated across 10 independent random seeds. Solid and dashed lines represent the cross-seed mean, while shaded regions indicate ± 1 standard deviation. (a) Optimization trace showing the best overpotential (J10) discovered as a function of the experimental budget. The dashed gray line marks the 370 mV discovery threshold. (b) The evolution of the surrogate model's predictive quality, quantified by the bootstrap out-of-bag R^2. (c) Global epistemic uncertainty ($\bar{\sigma}$) decay across the entire unlabeled candidate pool. The pure Greedy strategy (red) exhibits highly erratic optimization with massive variance and causes a catastrophic collapse in model quality ($R^2 \approx 0.4$). Conversely, the baseline LCBDS strategy ($n_q = 2$, blue) conservatively navigates the landscape with minimal variance, maintaining exceptional surrogate quality ($R^2 > 0.8$). The augmented Bold LCBDS strategy ($n_q = 3$, purple) demonstrates the ultimate potential of the framework, leveraging this stable model to trigger deep, late-stage optimization breakthroughs.

3.2 Deconstructing LCBDS: Component Ablation Analysis

Having established the macroscopic optimization superiority of the LCBDS framework, it is imperative to deconstruct the algorithm and isolate the specific drivers of this performance. To demonstrate the rigorousness and necessity of the proposed system design, we performed a comprehensive component ablation analysis, with the quantitative results and performance trade-offs summarized in Fig. 3 and Table 1.

The analysis reveals a critical balance between discovery efficiency and optimization depth, fundamentally governed by the spatial diversity and curiosity mechanisms. As detailed in Table 1, when the diversity penalty is entirely removed (w/o γ), the discovery efficiency of the framework is severely compromised. Without this spatial regularization,

the algorithm wastes valuable physical experimental budgets on redundant localized sampling, causing the $break_{exp}$ metric to drastically increase by 53%, inflating from 113 to 173 required experiments. Conversely, eliminating the neuroscience-inspired surprise mechanism (w/o δ) exposes a completely different failure mode. While the ablated model can still rapidly breach the initial 370 mV performance barrier ($break_{exp} = 104$), it loses the critical curiosity-driven momentum required to map unknown anomalous regions. Consequently, it ultimately fails to navigate toward deeper chemical spaces, prematurely stalling before reaching the absolute best optimal solutions achieved by the full LCBDS framework.

Most importantly, the ablation study unequivocally proves the absolute necessity of integrating the dedicated exploration query, Point B. When the maximum epistemic uncertainty query is removed (w/o MaxΣ), the active learning framework is forced into a purely exploitative sampling regime. While this might intuitively seem beneficial for rapid optimization, Table 1 reveals that this greedy behavior leads to a catastrophic collapse of the surrogate model's generalization capability, with the out-of-bag R^2 plummeting from 0.872 down to a mere 0.696. This profound degradation demonstrates that forced exploration is not merely an accessory function for escaping local minima. Rather, actively enforcing contrastive landscape coverage is strictly essential for maintaining the overall predictive quality, calibration, and structural integrity of the surrogate model throughout the autonomous discovery process.

Finally, while certain ablated variants may appear to outperform the full LCBDS framework on single isolated metrics, these hollow victories come at unacceptable systemic costs. For instance, the Pure LCB strategy (Point A only) breaches the discovery threshold more rapidly (87 ± 41 queries) but completely sacrifices global generalization capability (R^2 collapses to 0.674). Similarly, omitting the β annealing schedule allows the model to occasionally stumble upon a deeper minimum (343 mV), but it induces severe instability and massive variance across independent runs (± 66 queries). Ultimately, Table 1 proves that the full LCBDS framework is not merely optimized for a single metric, but uniquely achieves a robust, optimal balance maintaining the highest global predictive quality ($R^2 = 0.872$) while simultaneously guaranteeing rapid discovery and deep optimization depth without erratic variance.

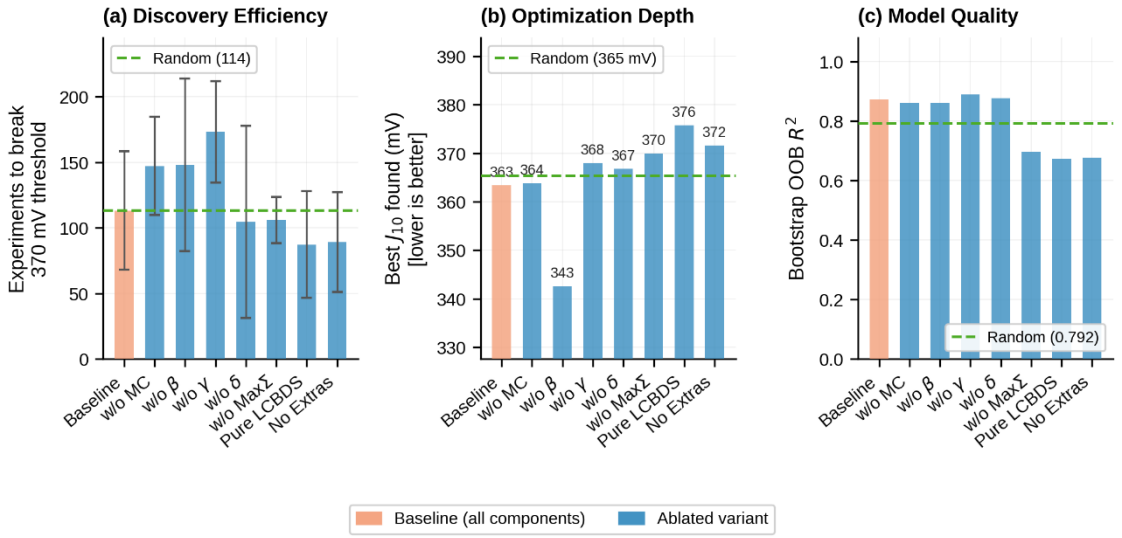


Fig. 3 | Deconstructing the LCBDS acquisition mechanisms. Component ablation analysis evaluated over n=5 independent seeds. The full proposed LCBDS framework (salmon bars) is compared against its ablated variants (blue bars) and a passive random baseline (green dashed lines). Error bars represent ± 1 standard deviation. (a) Discovery efficiency, defined as the required budget ($break_{exp}$) to breach the 370 mV threshold. (b) Optimization depth, denoting the absolute lowest J10 successfully synthesized. (c) Final surrogate generalization capability (OOB R^2). Notably, the removal of the diversity penalty (γ) severely degrades discovery efficiency (+53% queries), whereas the omission of the purely exploratory Point B ($Max\Sigma$) causes a catastrophic collapse in model quality ($\Delta R^2 \approx -0.18$), highlighting that contrastive landscape coverage is indispensable for accurate uncertainty calibration.

Table 1 | Quantitative ablation analysis. Performance metrics evaluated across n=5 independent seeds (mean $\pm$ standard deviation). $Break_{exp}$ indicates the number of experimental queries required to identify a composition with J10 < 370 mV. The full LCBDS configuration (bolded) provides the optimal balance across efficiency, depth, and surrogate generalization.

Method	$break_{exp} \downarrow$	Best $J_{10}(mV) \downarrow$	OOB$R^2 \uparrow$
LCBDS (Full, ours)	113 ± 45	363	0.872
w/o MC Dropout	147 ± 37	364	0.861
w/o β annealing	148 ± 66	**343**	0.860
w/o diversity (γ)	173 ± 39	368	**0.890**
w/o surprise (δ)	104 ± 73	367	0.878

w/o MaxΣ (Point B)	106 ± 18	370	0.696
Pure LCB (Point A only)	87 ± 41	376	0.674
No extras (Points A+B only)	89 ± 38	372	0.677
Random (passive baseline)	114 ± 58	365	0.792

3.3 Escaping Local Minima: The Impact of Bold Exploration

While traditional active learning frameworks excel at rapidly identifying high-performing candidates, they frequently converge prematurely on suboptimal local solutions, raising the fundamental question of whether they can truly and autonomously locate the absolute global optimum. This limitation arises because conventional surrogate models tend to exploit smoothly varying and heavily sampled regions of the chemical space, inherently struggling to navigate toward sharply bounded or isolated optimal zones.

To push the physical boundaries of autonomous discovery and evaluate the limits of our framework, we investigated its performance under two distinct sampling configurations. The first is an Aggressive strategy ($n_q = 2$), which utilizes the standard two-point batch query focusing on strict exploration and exploitation. The second is a Bold strategy ($n_q = 3$), which deliberately expands the batch by introducing a third query point targeting the maximum predicted mean (Maxμ). Paradoxically, this third point intentionally samples the theoretically worst-performing regions. Furthermore, this Bold strategy is synergistically coupled with a significantly amplified surprise weight (δ=25) to forcefully incentivize curiosity-driven sampling.

The comparative optimization trajectories, detailed in Fig. 4, reveal a nuanced but profound breakthrough in the framework's performance limits. The standard Aggressive LCBDS framework, prioritizing extreme cognitive stability, effectively bounds its variance but becomes strictly trapped above the 360 mV barrier, exhibiting a 0% probability of reaching deep optima ($\leq 260\, mV$). In contrast, the Bold LCBDS framework successfully breaches this deep performance threshold in 30% of independent runs, ultimately locking onto the absolute optimal composition (206 mV, $Ni_{0.302}Fe_{0.169}Co_{0.071}Ce_{0.471}$) with a superior 20% success rate. While the purely exploitative Greedy baseline does occasionally stumble upon these deep minima (achieving 20% and 10% rates for the respective thresholds) purely through high-variance,

stochastic luck as evidenced by its massive cross-seed scatter in Fig. 4a it lacks algorithmic reliability. This divergence provides compelling evidence that introducing a seemingly counter-intuitive exploration point, combined with an intensified surprise mechanism, provides a more systematic and higher-probability algorithmic momentum to carry the surrogate model across vast energy barriers.

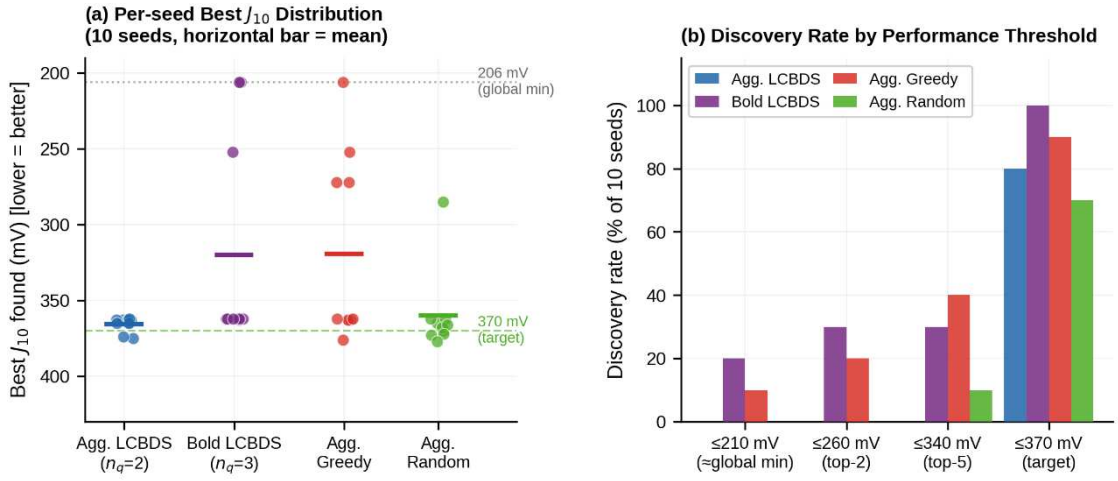


Fig. 4 | Overcoming energy barriers via bold exploration. Comparison between the Aggressive ($n_q = 2$, budget=220) and Bold ($n_q = 3$, budget=470, elevated surprise weighting $\delta = 25$) LCBDS variants across n=10 seeds. (a) Per-seed distribution of the optimal J10 discovered. Scatter markers denote individual trials, and horizontal solid bars indicate the cross-seed mean. The global optimum ($Ni_{0.302}Fe_{0.169}Co_{0.071}Ce_{0.471}$, 206 mV) is marked by the top dotted line. (b) Stepwise discovery success rate across progressively stringent performance thresholds. The Bold LCBDS variant successfully identifies the near-global minimum ($\leq 260\ mV$) in 30% of the independent runs, and the absolute optimum ($\leq 210\ mV$) in 20% of runs. While the Greedy strategy relies on high-variance stochastic luck to occasionally locate these deep minima (20% and 10% rates, respectively), Bold LCBDS achieves a superior discovery probability by systematically utilizing anomaly-driven surprise signals to escape local optima.

3.4 Hyperparameter Sensitivity and System Robustness

A common critique of complex autonomous discovery frameworks is their potential over-reliance on meticulously fine-tuned hyperparameters. To explicitly address this concern and establish the foundational stability of our approach, we conducted a systematic sensitivity analysis on the LCBDS algorithm. As depicted in Fig. 5, we evaluated the framework's performance across a comprehensive spectrum of values for three critical

parameters: the final exploration-exploitation boundary (β_{min}), the spatial diversity weight (γ), and the curiosity-driven surprise bonus (δ).

The systematic sweep results conclusively demonstrate that the algorithmic superiority of the LCBDS framework is not confined to a singular, cherry-picked configuration. Instead, the system maintains a robust and stable optimization advantage across a broad and physically reasonable parameter regime. We can map these stable functional regimes directly to algorithmic and physical intuitions. For the exploration boundary β_{min}, intermediate values ranging from 0.1 to 0.3 provide an ideal balance between exploring the chemical space and exploiting known optimal regions. However, setting this value too high ($\beta_{min} \geq 0.5$) forces the surrogate model into a state of perpetual over-exploration, which strictly prevents the framework from efficiently converging onto deep global minima.

Regarding the spatial diversity weight (γ), its presence is absolutely essential for efficient and stable discovery. Removing it entirely ($\gamma=0$) noticeably degrades the mean discovery efficiency and induces massive instability across independent runs. As shown in the expanded variance in Fig. 5b, without spatial regularization, worst-case runs can require nearly double the experimental budget (≈ 210 queries) compared to the optimized baseline. Conversely, applying an excessively high diversity penalty ($\gamma \geq 15$) overwhelms the predictive mean of the acquisition function. This extreme setting forces the model to select candidates purely based on their spatial isolation within the simplex, which leads to severe over-randomization and ironically re-introduces sampling redundancy by completely disregarding the predictive quality of the surrogate model. Finally, the surprise weight (δ) exhibits a highly synergistic behavior. A moderate application of this curiosity bonus (δ between 8 and 15) successfully drives the model to escape local traps and deepens the ultimate optimization depth, all without sacrificing the rapid discovery efficiency required in early iterations.

These consistent trends confirm that the LCBDS framework is inherently robust and does not suffer from brittle parameter dependence. Serving as the foundational cornerstone for our entire experimental architecture, the complete and optimized hyperparameter configuration utilized throughout the core evaluations of this study is comprehensively detailed in Supplementary Table 1 (Table S1).

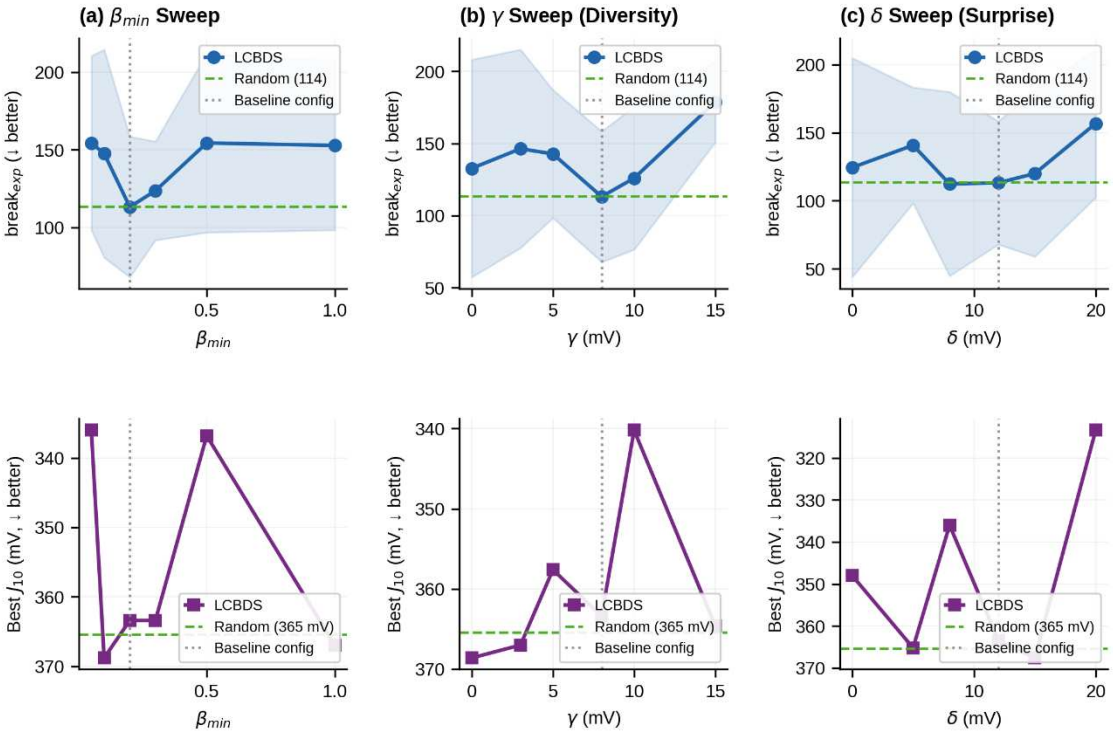


Fig. 5 | System robustness and hyperparameter sensitivity. Impact of key LCBDS parameters on discovery efficiency ($break_{exp}$, top row) and optimization depth (best J10, bottom row) over n=5 seeds. Markers represent the mean, and shaded areas denote ± 1 standard deviation. Dotted vertical lines highlight the default configuration used in the primary framework. (a) Annealing floor (β_{min}): intermediate values properly balance exploration and exploitation. (b) Diversity penalty (γ): moderate penalization effectively prevents redundant queries without causing over-dispersion. (c) Surprise bonus (δ): elevated curiosity weightings successfully push the model toward deeper optimization limits without penalizing early-stage discovery rates.

4 Discussion

4.1 Chemical Insights from Tabula Rasa Discovery

Although the PC-BAN surrogate model was initialized from a tabula rasa state completely devoid of a priori chemical rules or physical descriptors, it successfully navigated the high-dimensional chemical space to identify a striking global optimum at 206 mV with the specific composition of $Ni_{0.302}Fe_{0.169}Co_{0.071}Ce_{0.471}$. This optimal formulation is far from a random mathematical artifact. Instead, it indicates that the artificial intelligence autonomously recognized the profound synergistic effects existing between conventional earth-abundant transition metals (Ni, Fe, Co) and the rare-earth metal (Ce).

In typical transition metal oxide systems, optimizing the oxygen evolution reaction is inherently bottlenecked by the rigid scaling relations among the binding energies of the *OH, *O, and *OOH reaction intermediates[22-26]. The AI-driven selection of a Ce-rich composition provides a sophisticated chemical bypass to this fundamental limitation. From a physicochemical perspective, the incorporation of rare-earth elements like Ce heavily modulates the local electronic state and shifts the d-band center of the active transition metals through their unique 4f electronic structure[27, 28]. Additionally, the presence of Ce strongly facilitates the generation of abundant surface oxygen vacancies. These vacancies provide novel active sites that optimize the critical *OOH adsorption step, thereby fundamentally breaking the rigid scaling relations that typically restrict the performance ceiling of traditional Ni-Fe-Co ternary catalysts[29, 30].

The success of this specific discovery phase serves as compelling evidence regarding the true nature of the proposed learning algorithm. By converging on a complex and physically sound catalytic mechanism without being fed predefined structural features or density functional theory rules, the PC-BAN framework proves that it is not merely a high-dimensional mathematical fitting tool. Rather, it demonstrates a profound capability to autonomously rediscover and actively expand upon human chemical intuition purely through data-driven exploration.

4.2 The Role of "Surprise" in Navigating High-Dimensional Landscapes

In high-dimensional chemical spaces, the combinatorial explosion of possible atomic configurations creates a highly complex optimization landscape riddled with mediocre local optima. Traditional Bayesian Optimization frameworks, which heavily rely on conventional acquisition functions, frequently fail in these vast environments. These standard strategies typically prioritize safe exploitation or simple variance-based

exploration. Consequently, they cause the surrogate model to greedily converge upon and become irreversibly trapped within suboptimal performance regions. To overcome this fundamental mathematical limitation, our framework elevates the acquisition process by incorporating a surprise-driven curiosity mechanism directly inspired by cognitive neuroscience[31].

As demonstrated by the success of the Bold strategy, deliberately querying the Maxμ candidate presents a striking paradox. It intentionally allocates expensive physical experimental resources to synthesize the theoretically worst-performing materials within the candidate pool. However, this counter-intuitive strategy is physically and algorithmically necessary because it forces the surrogate model to confront its most pessimistic assumptions. In traditional models, heavily penalized regions are permanently ignored, creating vast blind spots in the AI's chemical intuition. By forcing the evaluation of these explicitly rejected domains, the framework rigorously tests the boundaries of its own predictive landscape.

When a physical experiment yields a result that significantly deviates from these bleak predictions, it generates a massive residual error. Instead of merely treating this as noise, the LCBDS algorithm systematically translates this error into a powerful surprise signal. Philosophically, this surprise bonus acts as a sudden surge of curiosity-driven algorithmic momentum. This localized momentum forcefully propels the surrogate model across vast energy barriers that separate different compositional clusters. By mathematically formalizing cognitive curiosity, the surprise mechanism enables the framework to escape the gravitational pull of barrier-enclosed local minima, ultimately navigating toward true breakthrough discoveries that traditional optimization would inherently deem impossible[7].

4.3 Bridging the Gap: From In Silico Validation to Prospective Autonomous Labs

While the proposed LCBDS framework demonstrates exceptional optimization capability, it is imperative to objectively acknowledge the boundaries of the current experimental setup. Specifically, the present study represents an in silico validation using an existing comprehensive database to act as an experimental oracle. This retrospective approach successfully proves the fundamental algorithmic superiority of the curiosity-driven search mechanism. However, transitioning from a static database oracle to a fully prospective, hardware-driven autonomous laboratory introduces practical complexities that cannot be fully captured by computational simulations alone.

If the current surrogate model were directly deployed to guide physical experiments in a real-world laboratory, the framework would immediately face three critical physical challenges. The first challenge is extrapolation errors. When the algorithm navigates entirely uncharted regions of the chemical space to discover novel element combinations, the surrogate model may inevitably struggle with out-of-domain predictions due to the lack of neighboring reference data. The second challenge involves distribution shifts. Real-world high-throughput synthesis intrinsically involves unavoidable variations in environmental and kinetic parameters, such as fluctuations in precursor mixing dynamics, local annealing temperatures, and subtle differences in electrochemical testing conditions. These physical distribution shifts between the idealized training assumptions and the actual experimental reality can severely degrade the accuracy of the surrogate model over time. The final challenge is model overconfidence. In data-scarce regimes typical of early-stage materials discovery, complex machine learning models frequently produce poorly calibrated and overconfident predictions, which can critically mislead the acquisition function and trap the active learning loop in suboptimal regions[32].

Despite these formidable physical challenges, the ultimate value of this research does not merely reside in the specific catalyst identified. Rather, this work provides a powerful initialization and navigation engine for future autonomous discovery. In the forthcoming era of prospective active learning campaigns integrated with intelligent robotic platforms, the tabula rasa PC-BAN model will serve as the foundational algorithmic starting point to direct actual automated hardware synthesis[33, 34]. By continuously feeding physical, dynamic measurement data back into the system for real-time iterative refinement, the framework will dynamically calibrate its network weights to actively correct for physical distribution shifts and prediction errors. This continuous closed-loop integration will systematically bridge the gap between computational in silico validation and the realization of fully self-driving materials laboratories[33-35].

4.4 Generality and Broader Impacts

The true potential of the LCBDS framework extends far beyond the specific optimization of OER electrocatalysts. Because the PC-BAN surrogate model operates entirely tabula rasa and is descriptor-free, it is fundamentally agnostic to the underlying physical or chemical domain. Traditional machine learning models often require extensive domain-specific feature engineering, which inherently restricts their transferability to new or distinctly different material systems. In contrast, our approach directly processes raw compositional proportions without relying on predefined physical rules or structural

heuristics. This inherent flexibility makes the framework universally adaptable to any scientific optimization problem characterized by a massive chemical design space.

This generalizability is particularly critical for tackling the combinatorial explosion of possibilities inherent in other cutting-edge scientific frontiers. For instance, the design of high-entropy alloys involves vast compositional spaces where the number of elemental combinations scales astronomically, rendering traditional optimization methods computationally and experimentally intractable[36-38]. Similarly, the search for optimal solid-state battery materials or the rapid navigation of the virtually infinite molecular configurations in modern drug discovery urgently requires intelligent sampling strategies capable of efficient exploration[39, 40]. The curiosity-driven surprise bonus and the spatial diversity penalty embedded in the LCBDS algorithm provide a robust, mathematically versatile mechanism to escape local optima across all these diverse and highly complex landscapes. By mathematically formalizing curiosity and spatial diversity, this framework lays a scalable foundation for the next generation of self-driving laboratories.

Conclusion

In summary, this work successfully demonstrates that an artificial intelligence framework, operating entirely from a tabula rasa state without relying on any a priori chemical features or domain descriptors, can autonomously uncover profound physicochemical principles within a massive high-dimensional design space. This achievement represents a fundamental methodological paradigm shift in active learning for materials discovery. The integration of the Maxμ query with the surprise bonus within the LCBDS algorithm proves a powerful philosophical and mathematical concept: to discover the absolute best results, the artificial intelligence must actively dare to confront its worst predictions. By deliberately sampling regions heavily penalized by conventional models and converting the resulting massive residual errors into curiosity-driven momentum, the framework mathematically overcomes the exploitation trap and successfully escapes deep local minima.

Having rigorously established this in silico validation using an exhaustive experimental database as an oracle, the critical next step is to bridge the gap toward practical application. The historical mandate moving forward is the direct deployment of this tabula rasa optimization engine onto fully integrated, closed-loop robotic hardware systems. In such prospective autonomous laboratories, the framework will continuously refine its network weights via real-time physical feedback, dynamically correcting for real-world hardware variations, kinetic uncertainties, and physical distribution shifts. Ultimately, this research marks a decisive and transformative step away from traditional human-guided materials screening, laying a robust and scalable foundation for the future of machine-autonomous chemical intuition.

Data availability

All data used in this work are from public resources. The experimental OER dataset used in this study is obtained from the publicly available supplementary material of Haber et al., Energy Environ. Sci. 2014, 7, 682–688 (DOI: 10.1039/C3EE43683G).

Code availability

The source code, implementation details, and precomputed results are freely available at the GitHub repository (https://github.com/CHIHX12/tabula-rasa-active-learning) and permanently archived on Zenodo (https://doi.org/10.5281/zenodo.20272736)[41].

Additional information

Competing interests: The authors declare no competing interests.

Funding

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

Reference

1. Ding, R. et al. Leveraging data mining, active learning, and domain adaptation for efficient discovery of advanced oxygen evolution electrocatalysts. Sci. Adv. 11, eadr9038 (2025).

2. Sun, Y. & Dai, S. High-entropy materials for catalysis: a new frontier. Sci. Adv. 7, eabg1600 (2021).

3. Kim, M. et al. Searching for an optimal multi-metallic alloy catalyst by active learning combined with experiments. Adv. Mater. 34, 2108900 (2022).

4. Tran, K. & Ulissi, Z. W. Active learning across intermetallics to guide discovery of electrocatalysts for CO2 reduction and H2 evolution. Nat. Catal. 1, 696–703 (2018).

5. Butler, K. T. et al. Machine learning for molecular and materials science. Nature 559, 547–555 (2018).

6. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. Nature 521, 436–444 (2015).

7. Lookman, T., Balachandran, P. V., Xue, D. & Yuan, R. Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. npj Comput. Mater. 5, 21 (2019).

8. Fromer, J. et al. Batched Bayesian optimization by maximizing the probability of including the optimum. arXiv preprint arXiv:2410.06333 (2024).

9. Turner, J. A. Sustainable hydrogen production. Science 305, 972–974 (2004).

10. Suen, N.-T. et al. Electrocatalysis for the oxygen evolution reaction: recent development and future perspectives. Chem. Soc. Rev. 46, 337–365 (2017).

11. Haber, J. A., Xiang, C., Guevarra, D., Jung, S., Jin, J. & Gregoire, J. M. High-throughput mapping of the electrochemical properties of (Ni–Fe–Co–Ce)Ox oxygen-evolution catalysts. ChemElectroChem 1, 1613–1621 (2014).

12. Gal, Y. & Ghahramani, Z. Dropout as a Bayesian approximation: representing model uncertainty in deep learning. Proc. Mach. Learn. Res. 48, 1050–1059 (2016).

13. Bishop, C. M. Mixture density networks. Tech. Rep. NCRG/94/004, Aston University (1994).

14. Kingma, D. P. & Ba, J. Adam: a method for stochastic optimization. Proc. Int. Conf. Learn. Represent. (ICLR, 2015).

15. Lakshminarayanan, B., Pritzel, A. & Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. Adv. Neural Inf. Process. Syst. 30, 5065–5076 (2017).

16. Settles, B. Active Learning. Morgan & Claypool Publishers (2012).

17. Shahriari, B. et al. Taking the human out of the loop: a review of Bayesian optimization. Proc. IEEE 104, 148–175 (2016).

18. Frazier, P. I. & Wang, J. Bayesian optimization for materials design. Inf. Sci. Mater. Discov. Design 45–75 (Springer, 2016).
19. Srinivas, N., Krause, A., Kakade, S. M. & Seeger, M. Gaussian process optimization in the bandit setting: no regret and experimental design. Proc. 27th Int. Conf. Mach. Learn. 1015–1022 (2010).
20. Nava, E., Mutny, M. & Krause, A. Diversified sampling for batched Bayesian optimization with determinantal point processes. Proc. Mach. Learn. Res. 151, 7031–7054 (2022).
21. Loshchilov, I. & Hutter, F. SGDR: stochastic gradient descent with warm restarts. Proc. Int. Conf. Learn. Represent. (ICLR, 2017).
22. Seh, Z. W. et al. Combining theory and experiment in electrocatalysis: insights into materials design. Science 355, eaad4998 (2017).
23. Craig, M. J. et al. Universal scaling relations for the rational design of molecular water oxidation catalysts with near-zero overpotential. Nat. Commun. 10, 4993 (2019).
24. Rossmeisl, J. et al. Electrolysis of water on oxide surfaces. J. Electroanal. Chem. 607, 83–89 (2007).
25. Nørskov, J. K. et al. Origin of the overpotential for oxygen reduction at a fuel-cell cathode. J. Phys. Chem. B 108, 17886–17892 (2004).
26. Man, I. C. et al. Universality in oxygen evolution electrocatalysis on oxide surfaces. ChemCatChem 3, 1159–1165 (2011).
27. Hu, J. et al. Mastering the D-band center of iron-series metal-based electrocatalysts for enhanced electrocatalytic water splitting. Int. J. Mol. Sci. 23, 15405 (2022).
28. Hammer, B. & Nørskov, J. K. Electronic factors determining the reactivity of metal surfaces. Surf. Sci. 343, 211–220 (1995).
29. Ovcharenko, R., Voloshina, E. & Sauer, J. Water adsorption and O-defect formation on Fe2O3(0001) surfaces. Phys. Chem. Chem. Phys. 18, 25560–25568 (2016).
30. McCrory, C. C. L. et al. Benchmarking heterogeneous electrocatalysts for the oxygen evolution reaction. J. Am. Chem. Soc. 135, 16977–16987 (2013).
31. Oudeyer, P.-Y. & Kaplan, F. What is intrinsic motivation? A typology of computational approaches. Front. Neurorobot. 1, 6 (2007).
32. Merchant, A. et al. Scaling deep learning for materials discovery. Nature 624, 80–85 (2023).
33. MacLeod, B. P. et al. Self-driving laboratory for accelerated discovery of thin-film materials. Sci. Adv. 6, eaaz8867 (2020).
34. Abolhasani, M. & Kumacheva, E. The rise of self-driving labs in chemical and materials sciences. Nat. Synth. 2, 483–492 (2023).

35. Szymanski, N. J. et al. An autonomous laboratory for the accelerated synthesis of novel materials. Nature 624, 86–91 (2023).
36. Yao, Y. et al. High-entropy nanoparticles: synthesis-structure-property relationships and data-driven discovery. Science 376, eabn3103 (2022).
37. George, E. P., Raabe, D. & Ritchie, R. O. High-entropy alloys. Nat. Rev. Mater. 4, 515–534 (2019).
38. Cantor, B., Chang, I. T. H., Knight, P. & Vincent, A. J. B. Microstructural development in equiatomic multicomponent alloys. Mater. Sci. Eng. A 375, 213–218 (2004).
39. Pyzer-Knapp, E. O. Bayesian optimization for accelerated drug discovery. IBM J. Res. Dev. 62, 2:1–2:7 (2018).
40. Schmidt, J. et al. Recent advances and applications of machine learning in solid-state materials science. npj Comput. Mater. 5, 83 (2019).
41. Cheng, C.-Y., Ting, H.-Y., Wu, Y.-H. & Li, F.-Y. Autonomous materials discovery through a tabula rasa active learning framework. Zenodo, DOI:10.5281/zenodo.20272736 (2026).

Supplementary Files

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

- Supplementarymaterial.docx