

ORIGINAL RESEARCH

Open access

Why Active Learning Fails for Defect Formation Energies: When Uncertainty Sampling Misleads Discovery

Yousef Al-Qahtani^{1*}, Fahad Al-Salem¹, Abdullah Al-Harbi²

Abstract

Active learning is widely used in materials discovery to reduce the number of expensive density functional theory calculations required for surrogate-model training. Its dominant acquisition rule, uncertainty sampling, assumes that the most uncertain configuration will also yield the greatest information gain. Although this logic performs well for many bulk properties, it breaks down for defect formation energies. In this setting, uncertainty sampling repeatedly selects uninformative structures, misallocates computational budget, and fails to reach the defective configurations that govern vacancy, interstitial, and substitutional behavior in technologically important materials. This failure reflects a structural mismatch between the acquisition rule and defect physics. Defect configurations are rare in configuration space, energetics are highly localized, and DFT labels often contain substantial aleatoric noise from finite-size effects, charge corrections, and supercell artifacts. These conditions weaken the link between predictive uncertainty and useful learning signal, while representation bias in models trained mainly on perfect crystals further erodes selectivity. This study develops a failure-mode analysis of uncertainty sampling for defect formation energies, identifying recurrent breakdowns in sampling, representation, calibration, and budget use. It also outlines practical detection principles and defect-aware mitigation strategies, including initialization with defective structures, hybrid acquisition functions, epistemic-only scoring, and budget partitioning. The central implication is that active learning for defect discovery cannot rely on generic uncertainty-based querying. Effective acceleration in this domain requires acquisition strategies designed around the rarity, locality, and noise structure of defect energetics.

Keywords Materials discovery, Active learning, Failure mode analysis, Uncertainty sampling, Defect formation energies, Sampling bias

*Correspondence:

Yousef Al-Qahtani

yousef.qahtani@gmail.com

¹ Department of Computational Materials Analytics, Faculty of Engineering, King Saud University, Riyadh, Saudi Arabia

² Department of Smart Materials Systems, Faculty of Science and Technology, King Abdulaziz University, Jeddah, Saudi Arabia

Introduction

Active learning promises to accelerate materials discovery by intelligently selecting which calculations to perform next. The most common query strategy—uncertainty sampling—selects the structure the model is most uncertain about [1–3]. For many properties, this works well. But for defect formation energies—a critical property for semiconductors, batteries, and irradiation-resistant alloys—uncertainty sampling systematically fails [4, 5]. It selects the wrong structures, wastes DFT calculations, and misses the most important defects. This paper analyzes failure modes of uncertainty sampling for defect formation energies and proposes alternatives.

The appeal of active learning is obvious. First-principles calculations remain the gold standard for defect energetics, yet the configuration space is astronomically large. Even modest supercells yield thousands of possible defect arrangements,

and each DFT relaxation is expensive. Machine-learning interatomic potentials, whether based on deep neural networks or equivariant graph architectures, can dramatically reduce this cost once adequately trained [6, 7]. The bottleneck is training data. Random sampling is inefficient; active learning offers a principled way to prioritize the most informative points [8, 9].

Uncertainty sampling is the default choice because it is simple and model-agnostic. After training on an initial labeled set, the model predicts mean and uncertainty on a large unlabeled pool; the configuration with the largest uncertainty is queried, labeled by DFT, added to the training set, and the loop repeats [1, 10, 11]. In smooth, data-abundant regimes—such as bulk elastic constants or molecular conformer energies—this correlation between uncertainty and information gain holds. Model error decreases rapidly, and the surrogate quickly becomes reliable across the relevant domain [12].

Defect formation energies, however, violate the assumptions that make uncertainty sampling effective. Defects are rare: in a perfect crystal the vast majority of atomic sites are defect-free [13]. The formation energy is only meaningfully defined once a defect is present, yet the unlabeled pool is overwhelmingly populated by pristine configurations. Uncertainty sampling therefore spends early iterations labeling more perfect crystals, which add little new information once the bulk potential is learned [14]. Even when defects are eventually queried, their energies are determined by local environments within 5–10 Å of the defect core. Global uncertainty metrics, however, respond to long-range disorder that may be irrelevant to the local defect physics.

Compounding the problem is the high aleatoric noise inherent to defect calculations. Supercell convergence, image-charge corrections, and finite-size electrostatics introduce irreducible scatter that uncertainty sampling cannot filter [15, 16]. Graph neural networks, optimized for periodic symmetric structures, also suffer representation bias when symmetry is broken by defects [7, 17]. The model therefore assigns high uncertainty indiscriminately to all defective configurations, rendering selection non-selective. Researchers nevertheless assume uncertainty sampling should transfer to defects because the property is important and data are scarce. Prior success with perfect-crystal potentials and molecular systems creates overconfidence [6, 17]. The hidden assumption—that high uncertainty always signals informative defect configurations—fails precisely when the target property is rare, localized, and noisy.

What Makes Defect Formation Energies Hard for Active Learning

Defect formation energies destabilize the assumptions underpinning active learning workflows, exposing a set of tightly coupled mechanisms through which uncertainty sampling—otherwise effective for bulk properties—systematically fails. Extreme configurational rarity ensures that randomly generated supercells overwhelmingly correspond to pristine lattices, so early iterations concentrate on low-information variants where residual model uncertainty persists, while genuinely defective states remain statistically inaccessible within finite computational budgets [8, 13, 14]. This imbalance is compounded by the strongly localized nature of defect energetics: formation energies are governed by atomic rearrangements confined to a few coordination shells, whereas global descriptors and supercell-level uncertainty metrics respond to distant, energetically irrelevant perturbations, thereby redirecting sampling toward structurally disordered yet physically inconsequential configurations [4, 5, 15, 18]. Under these conditions, the presence of substantial aleatoric noise further distorts the acquisition process, as finite-size corrections and electrostatic artifacts in DFT calculations introduce irreducible variance that uncertainty sampling misinterprets as epistemic, prolonging exploration of regions where additional data yields negligible informational gain [10, 16, 19]. The problem is reinforced by the intrinsic sparsity of defect data: training sets dominated by pristine configurations induce overconfident representations of the bulk lattice alongside uniformly diffuse uncertainty across defective states, preventing meaningful discrimination during query selection [14, 20]. Beyond these constraints, the non-smoothness of the defect energy landscape introduces sharp, high-curvature regions where small atomic displacements produce substantial energy shifts, causing surrogate models trained on smoother bulk regimes to extrapolate unreliably and generate spurious uncertainty in physically inaccessible domains [12, 21]. Taken together, these interacting effects invert the expected alignment between uncertainty and information gain,

rendering standard uncertainty sampling both inefficient and systematically misdirected in the context of defect-driven materials discovery.

Table 1 clarifies that the collapse of uncertainty sampling does not arise from a single implementation flaw but from a systematic mismatch between defect physics, the hidden assumptions of the acquisition rule, and the mechanisms through which failure unfolds.

Table 1. Crosswalk between Defect-Specific Target Conditions, Broken Assumptions of Uncertainty Sampling, and Failure Mechanisms

Defect-specific condition	What standard uncertainty sampling implicitly assumes instead	Why the assumption fails for defect formation energies	Resulting failure mechanism(s)	Analytical implication for acquisition design
Rarity of defect-containing configurations	High-uncertainty regions will be encountered frequently enough through pool-based ranking	The unlabeled pool is dominated by pristine or near-pristine structures, so ranking remains concentrated in abundant low-value regions	Rare-event blindness; perfect-crystal trap	Acquisition must contain an explicit defect-presence prior, defect-aware filtering, or enforced exploration into low-density subspaces
Locality of defect energetics	Whole-structure uncertainty is a suitable proxy for local information gain	Global disorder can inflate uncertainty even when the local defect environment is unchanged or irrelevant to formation energy	Local-minimum trapping; budget-exhaustion failure	Query functions should incorporate local-environment descriptors or defect-core-sensitive uncertainty metrics
High aleatoric noise from supercell artifacts and charge corrections	Total predictive uncertainty is predominantly reducible by acquiring more labels	A substantial share of uncertainty reflects irreducible scatter, so additional labeling does not meaningfully improve the surrogate	Noise overestimation; high-noise sink	Epistemic and aleatoric components must be separated before acquisition scores are computed
Sparse initial defect data	Early model uncertainty over defect structures will still be selectively informative	When the model has seen almost no defects, all defect structures appear uniformly unfamiliar, destroying ranking resolution	Representation bias; representation collapse	Initialization must include a deliberately structured seed set of defect-containing configurations
Non-smooth local energy landscape	High-uncertainty points lie near learnable transitions that can be progressively refined	Sharp local basins and defect relaxations can place high uncertainty in physically marginal or inaccessible regions	Local-minimum trapping; miscalibration failure	Acquisition should be constrained by physical plausibility, structure-based screening, and diversity-aware search
Perfect-crystal-dominated training distribution	Good performance on bulk-like environments	The learned representation inherits periodicity-friendly	Representation bias; exploration-	Representation learning must be adapted to the

	transfers smoothly to symmetry-broken structures	inductive biases that do not rank defect states meaningfully	exploitation collapse	defect regime before or during active learning
Pure uncertainty as the sole acquisition criterion	Exploitation naturally transitions into sufficient exploration	Querying stays too close to already modeled regions unless novelty pressure is explicitly introduced	Exploration-exploitation collapse; budget-exhaustion failure	Hybrid acquisition functions should combine uncertainty with diversity, novelty, or region-level budget constraints

How Uncertainty Sampling Works (And Why it Should Work for Other Properties)

The standard uncertainty sampling loop is deceptively simple. An initial small dataset of DFT-labeled structures trains a surrogate model—commonly a graph neural network or deep potential [6, 7]. The model then predicts both mean property and uncertainty (via ensemble variance, dropout, or evidential methods) across a large pool of unlabeled candidate configurations. The configuration with the highest uncertainty is selected, its true value is obtained via DFT, the labeled set is updated, the model is retrained, and the process repeats [1, 10, 11, 22].

This loop succeeds for many properties because high uncertainty reliably flags regions of high information gain. In smooth, densely sampled landscapes the model quickly reduces epistemic uncertainty, and each new point meaningfully shrinks the prediction error across the domain [8, 12]. The exploration-exploitation balance tilts naturally toward exploitation once the bulk behavior is captured [23].

Researchers therefore assume the same logic should apply to defects. Defects are technologically critical, data are scarce, and active learning has succeeded for perfect-crystal potentials and molecular systems [6, 17]. The literature on active learning in materials repeatedly endorses uncertainty sampling as the default, citing its simplicity and empirical track record [2, 8, 9, 20, 22]. The assumption is that the same query strategy will automatically discover important defect configurations because they are the points the model “does not know.”

The hidden assumption that breaks is that high uncertainty always corresponds to informative defect configurations. For defects, uncertainty frequently reflects aleatoric noise, representation mismatch, or abundance of perfect-crystal variants rather than missing defect physics [19]. When the target property is rare and noisy, the correlation between uncertainty and information gain dissolves. The loop therefore wastes queries on uninformative points, never reaching the rare, localized, high-signal defect configurations that would actually improve the model for the property of interest.

Mechanisms of Active Learning Failure

Uncertainty sampling fails for defect formation energies through a set of mutually reinforcing mechanisms that disrupt its core assumptions. The extreme sparsity of defect configurations renders the algorithm effectively blind to the relevant subspace, as early iterations remain dominated by abundant pristine structures that retain marginal uncertainty, exhausting the sampling budget before meaningful defect exploration occurs [8, 13, 14]. As confidence over perfect crystals increases, the acquisition process narrows further, gravitating toward minor distortions of the bulk lattice that offer negligible informational gain, while the larger configurational shifts required to access defect basins remain systematically excluded due to locally reduced uncertainty estimates [12, 23, 24]. This dynamic is exacerbated by persistent aleatoric noise arising from charge corrections and finite-size effects, which sustains elevated uncertainty levels even after repeated sampling, leading the model to revisit intrinsically noisy regions under the false premise of reducible ignorance [10, 19]. At the representational level, equivariant graph networks trained on periodic systems struggle to encode symmetry-breaking defect structures, producing homogenized uncertainty estimates that eliminate discrimination among

candidate configurations and undermine query selectivity [7, 17]. The absence of an explicit exploratory component further compounds these limitations, as the inherently exploitative bias of uncertainty sampling confines the search to partially understood regions rather than promoting traversal into distant, high-value domains [2, 8, 24]. Through their interaction, these effects induce a self-reinforcing collapse in which blindness gives rise to trapping, noise distorts prioritization, and representational inadequacies suppress differentiation, ultimately yielding a workflow that appears adaptive yet fails to access the defect states central to the scientific objective [19, 20].

Figure 1 maps the full analytical progression by linking defect-specific problem conditions to the hidden failure of uncertainty sampling, the resulting mechanism-level breakdowns, their observable signatures, and the redesign principles required for defect-aware active learning.




Figure 1. Structural Architecture of Failure, Detection, and Mitigation in Active Learning for Defect Formation Energies

A Typology of Failure Modes

Failure in uncertainty-driven active learning for defect formation energies exhibits recurrent, diagnostically distinct patterns that reflect deeper breakdowns in the acquisition process. A persistent bias toward pristine configurations often dominates early iterations, with queried structures remaining defect-free and the fraction of defect-containing samples failing to exceed 10 % even after multiple rounds, indicating that the search remains confined to low-information regions of configuration space [8, 14]. As sampling eventually reaches defective states, progress may stall under conditions where prediction error on defect energetics ceases to improve despite continued querying, revealing the influence of irreducible noise that disrupts learning dynamics [10, 19]. This stagnation is frequently accompanied by a collapse in representational discrimination, where uncertainty estimates for defect configurations converge to uniformly high values, eliminating any meaningful ranking signal and undermining selection efficiency [7, 17]. Over time, the discovery rate of novel defect types may diminish prematurely, with cumulative counts flattening well before computational resources are

fully utilized, reflecting an inability to access unexplored regions of interest [20, 24]. In parallel, systematic miscalibration can emerge, as discrepancies between predicted uncertainty and observed error distort the acquisition signal, particularly when overconfidence or underconfidence skews query priorities [10]. The presence of any such pattern is sufficient to compromise effectiveness, while their co-occurrence intensifies both computational inefficiency and the loss of scientific insight.

Detection Principles

The onset of these breakdowns can be anticipated through operational indicators that map directly onto their underlying mechanisms. Monitoring the proportion of queried configurations that contain defects provides an immediate signal of whether the search has escaped the dominance of pristine structures, with persistently low values indicating entrapment in the bulk regime [8, 14]. A more granular perspective emerges from decomposing uncertainty into epistemic and aleatoric components, where dominance of irreducible noise signals diminishing returns from continued sampling and the emergence of noise-driven inefficiency [10, 25]. Representation quality can be probed through latent-space similarity measures, as uniformly distant and unstructured embeddings for defect configurations indicate a loss of internal organization and predictive resolution [7, 17]. Learning-curve trajectories further clarify system behavior, with early plateaus in defect-specific error revealing that the acquisition process has expended its budget without accessing informative regions [20, 24]. Calibration diagnostics, particularly deviations from ideal reliability alignment, expose distortions in the uncertainty signal that misguide selection [10]. These indicators require minimal overhead yet provide early warnings; once multiple signals align, incremental tuning becomes insufficient and more fundamental intervention is required.

Mitigation Principles

Restoring effectiveness requires targeted adjustments that reintroduce selectivity while preserving the active-learning framework. Introducing defect configurations at initialization alters the prior distribution seen by the model, ensuring that symmetry-broken states are represented from the outset rather than left to rare discovery events [14, 18]. This shift is reinforced by incorporating diversity-aware acquisition criteria, where latent-space dispersion complements uncertainty to promote exploration beyond locally familiar regions [8, 23, 24]. Additional control can be exerted by constraining the candidate pool to structurally defect-like configurations, thereby preventing the dominance of bulk-like environments during ranking [15, 21, 26]. Multi-fidelity schemes extend this logic by leveraging inexpensive models to pre-screen large candidate sets, reserving high-cost evaluations for configurations that are both informative and defect-relevant [22, 27, 28]. A sharper focus on epistemic uncertainty further mitigates noise-driven inefficiencies by explicitly excluding irreducible variance from acquisition decisions [10, 25, 29]. Improvements at the representation level, particularly through pre-training on defect structures, enhance the model's capacity to encode symmetry-breaking features and restore meaningful differentiation in latent space [7, 17]. Finally, explicit partitioning of the computational budget between bulk refinement and defect exploration prevents resource monopolization by low-value regions and enforces sustained attention on scientifically relevant configurations [20, 24]. When these interventions are integrated, the alignment between uncertainty and information gain is re-established, allowing the active-learning process to function as an effective mechanism for defect discovery rather than a source of redundant computation.

Table 2 converts the failure analysis into an operational decision matrix by linking each observable breakdown pattern to a concrete diagnostic interpretation, an appropriate first-line intervention, and a clear statement of what will not solve the problem.

Table 2. Diagnostic Matrix for Real-Time Detection and Strategy Selection in Active Learning for Defect Formation Energies

Observable failure mode	Primary operational indicator	What pattern confirms the failure	What the pattern means analytically	Most appropriate first-line mitigation	What should not be expected to fix it
Perfect-crystal trap	Fraction of queried structures containing true defects	Defect fraction remains below 10% after the first 10 queries	The acquisition rule is still operating inside the abundant pristine manifold rather than entering the target defect subspace	Defect-aware initialization; structure-based filtering; region-specific budget reservation	More iterations of the same uncertainty-only strategy
High-noise sink	Defect-specific test error versus number of queries; epistemic/aleatoric split	Error plateaus while aleatoric uncertainty dominates acquisition scores	The loop is spending budget on irreducible noise rather than reducible ignorance	Epistemic-only sampling; multi-fidelity screening; down-weighting noisy charged-defect regions	Blindly increasing label budget in the same noisy region
Representation collapse	Distribution of uncertainty scores or latent distances across defect candidates	Defect candidates all receive similarly high uncertainty with weak internal ranking	The representation does not support selective discrimination among defect states	Defect-oriented representation pre-training; auxiliary defect-structure exposure before acquisition	Fine-tuning only the acquisition threshold while keeping the same latent geometry
Budget-exhaustion failure	Cumulative unique defect types discovered per query; defect-learning curve	Discovery rate flattens early and defect error remains above target accuracy	Query diversity is too low and exploration is not reaching underrepresented defect families	Hybrid uncertainty-plus-diversity acquisition; explicit novelty term; region-based budget allocation	Continued local exploitation around already sampled pockets
Miscalibration failure	Reliability diagram comparing predicted uncertainty to observed error	Strong deviation from the ideal diagonal, especially across defect subsets	The acquisition signal itself is corrupted, so ranking cannot be trusted	Recalibration, uncertainty decomposition, defect-specific validation splits, and revised acquisition scoring	Assuming high uncertainty still corresponds to high information gain without recalibration
Mixed or coupled failure	Two or more indicators appear simultaneously	Low defect rate, plateaued defect error, narrow high defect uncertainty, and poor calibration occur together	The problem is structural and spans sampling, representation, and uncertainty treatment simultaneously	Combine defect-aware initialization, hybrid querying, epistemic-only scoring, and representation adaptation	Single-parameter tuning or isolated cosmetic workflow changes

Relation to Other Failure Modes

The failure modes documented here are not isolated; they share deep structural similarities with other well-recognized pathologies in machine-learning-driven materials science.

Relation to catastrophic forgetting [20]: both phenomena originate from sampling bias that over-represents one region of configuration space (perfect crystals here, earlier tasks in sequential learning). In catastrophic forgetting the model overwrites previously learned knowledge; here the model never acquires the missing defect knowledge in the first place because the query strategy never visits that subspace. The common root is the absence of an explicit mechanism to protect or explore underrepresented regions [13].

Relation to transfer collapse [17]: both involve a representation mismatch between source and target domains. Transfer collapse occurs when a model trained on one chemical space fails to generalize to another; the present failure occurs when a model trained predominantly on perfect crystals fails to generalize to symmetry-broken defect configurations within the same material. In both cases the graph neural network's inductive bias—optimized for periodicity and high symmetry—produces embeddings that cannot support reliable uncertainty ranking in the target regime [7].

Relation to aleatoric/epistemic separation [10]: this failure-mode analysis demonstrates why such separation is not merely a technical refinement but a fundamental requirement for defect energetics. Without it, uncertainty sampling cannot distinguish configurations whose high uncertainty is reducible (and therefore valuable) from those whose high uncertainty is irreducible noise. The collapse of the exploration-exploitation trade-off is therefore an inevitable downstream consequence of treating total uncertainty as a monolithic signal [19, 25].

Recognizing these connections situates defect-formation-energy failure within the broader landscape of active-learning limitations [2, 3]. It also suggests that solutions developed for one pathology—such as replay buffers for forgetting or domain-adversarial training for transfer—may inspire analogous remedies here, provided they respect the localized, noisy, and rare-event character of defect physics.

Implications for Active Learning in Defect Discovery

For materials scientists the central takeaway is unambiguous: standard uncertainty sampling cannot be applied off-the-shelf to defect formation energies. Any workflow that begins with a perfect-crystal-dominated training set and relies on pure uncertainty sampling will almost certainly fall into the perfect-crystal trap or the high-noise sink. Pre-seeding the training set with known defect configurations, enforcing structure-based filtering, and allocating budgets by region are now minimum requirements rather than optional enhancements [14, 15, 18].

For machine-learning researchers the challenge is to design defect-specific acquisition functions that explicitly incorporate locality and noise awareness. Global uncertainty scores must be replaced or augmented by local-environment uncertainty metrics that focus on the 5–10 Å vicinity of potential defect cores. Representations must be pre-conditioned or jointly trained to handle both periodic and symmetry-broken geometries without representation collapse [7, 17]. The exploration-exploitation trade-off must be re-balanced with an explicit diversity or novelty term tuned for rare events [8, 23, 24].

For benchmark designers the implication is equally clear. Existing active-learning benchmarks that report only aggregate error reduction on mixed bulk-plus-defect test sets mask the very failures documented here. Future benchmarks must report defect discovery rate (unique defect types found per query), defect-specific learning curves, and the fraction of budget wasted on perfect-crystal queries. Only then can the community quantitatively compare uncertainty sampling against hybrid or defect-aware alternatives [9, 20].

Collectively these changes shift active learning from a generic black-box optimizer to a domain-informed discovery tool. Defect formation energies are not just another regression target; they are the gatekeepers to understanding radiation

damage, ionic transport, and semiconductor doping. If active learning is to fulfill its promise in these areas, it must respect the physical peculiarities that make defects hard—rarity, locality, noise, sparsity, and non-smoothness—rather than assume that a strategy successful for perfect crystals will automatically succeed for everything else.

Conclusion

Uncertainty sampling fails for defect formation energies not because active learning is intrinsically unsuitable, but because its default assumptions do not survive contact with defect physics. When the target property is rare, localized, noisy, and sparsely represented, high predictive uncertainty no longer provides a reliable proxy for information gain. The result is a recurring pattern of misdirected queries, weak defect discovery, and inefficient use of DFT budget. The analysis developed here shows that this breakdown is systematic rather than incidental, emerging through interacting failures in sampling, representation, calibration, and exploration. It also shows that these failures are detectable before the budget is exhausted and, importantly, that they can be mitigated through defect-aware initialization, hybrid acquisition, epistemic-only uncertainty treatment, representation adaptation, and explicit budget control. The broader lesson extends beyond defect energetics. Any materials problem defined by rare events, strong locality, or irreducible noise is unlikely to benefit from off-the-shelf uncertainty sampling. Progress therefore depends on replacing generic acquisition rules with domain-informed strategies that reflect the structure of the underlying physics. Only under that shift can active learning become a reliable instrument for defect discovery and, more broadly, for scientifically meaningful materials design.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 16 Jul 2023

Revised: 29 Sep 2023

Accepted: 20 Dec 2023

Published online: 18 January 2024

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Lombardi D, Shipley TF, Astronomy Team, Biology Team, Chemistry Team, Engineering Team, et al. The curious construct of active learning. *Psychol Sci Public Interest*. 2021;22(1):8-43.
<https://doi.org/10.1177/1529100620973974>.
- Kumar P, Gupta A. Active learning query strategies for classification, regression, and clustering: A survey. *J Comput Sci Technol*. 2020;35(4):913-45.
<https://doi.org/10.1007/s11390-020-9487-4>.
- Doolittle P, Wojdak K, Walters A. Defining active learning: A restricted systematic review. *Teach Learn Inq*. 2023;11:25.
<https://doi.org/10.20343/teachlearningqu.11.25>.
- Liu Y, Wang H, Guo L, Yan Z, Zheng J, Zhou W, et al. Deep learning inter-atomic potential for irradiation damage in 3C-SiC. *Comput Mater Sci*. 2024;233:112693.
<https://doi.org/10.1016/j.commatsci.2023.112693>.
- Rahman MH, Biswas M, Mannodi-Kanakithodi A. Understanding defect-mediated ion migration in semiconductors using atomistic simulations and machine learning. *ACS Mater Au*. 2024;4(6):557-73.
<https://doi.org/10.1021/acsmaterialsau.4c00095>.
- Zhang L, Han J, Wang H, Car R, E W. Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics. *Phys Rev Lett*. 2018;120(14):143001.
<https://doi.org/10.1103/PhysRevLett.120.143001>.
- Batzner S, Musaelian A, Sun L, Geiger M, Mailoa JP, Kornbluth M, et al. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat Commun*. 2022;13(1):2453.
<https://doi.org/10.1038/s41467-022-29939-5>.
- Lookman T, Balachandran PV, Xue D, Yuan R. Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. *Npj Comput Mater*. 2019;5(1):21.
<https://doi.org/10.1038/s41524-019-0153-8>.
- Kusne AG, Yu H, Wu C, Zhang H, Hattrick-Simpers J, DeCost B, et al. On-the-fly closed-loop materials discovery via Bayesian active learning. *Nat Commun*. 2020;11(1):5966.
<https://doi.org/10.1038/s41467-020-19597-w>.
- Nguyen VL, Shaker MH, Hüllermeier E. How to measure uncertainty in uncertainty sampling for active learning. *Mach Learn*. 2022;111(1):89-122.
<https://doi.org/10.1007/s10994-021-06003-9>.
- Konyushkova K, Sznitman R, Fua P. Learning active learning from data. In: *Advances in Neural Information Processing Systems 30*; 2017 Dec 4-9; Long Beach, CA, USA; 2017. p. 4228-38.
<https://doi.org/10.5555/3294996.3295177>.
- Zhang L, Lin DY, Wang H, Car R, E W. Active learning of uniformly accurate interatomic potentials for materials simulation. *Phys Rev Mater*. 2019;3(2):023804.
<https://doi.org/10.1103/PhysRevMaterials.3.023804>.
- Kaur H, Pannu HS, Malhi AK. A systematic review on imbalanced data challenges in machine learning: Applications and solutions. *ACM Comput Surv*. 2019;52(4):79.
<https://doi.org/10.1145/3343440>.
- Nie S, Xiang Y, Wu L, Lin G, Liu Q, Chu S, et al. Active learning guided discovery of high entropy oxides featuring high H₂-production. *J Am Chem Soc*. 2024;146(43):29325-34.

<https://doi.org/10.1021/jacs.4c06272>.

Mannodi-Kanakkithodi A, Xiang X, Jacoby L, Biegaj R, Dunham ST, Gamelin DR, et al. Universal machine learning framework for defect predictions in zinc blende semiconductors. *Patterns (N Y)*. 2022;3(3):100450.

<https://doi.org/10.1016/j.patter.2022.100450>.

Dral PO. Quantum chemistry in the age of machine learning. *J Phys Chem Lett*. 2020;11(6):2336-47.

<https://doi.org/10.1021/acs.jpclett.9b03664>.

Bartók AP, De S, Poelking C, Bernstein N, Kermode JR, Csányi G, et al. Machine learning unifies the modeling of materials and molecules. *Sci Adv*. 2017;3(12):e1701816.

<https://doi.org/10.1126/sciadv.1701816>.

Arrigoni M, Madsen GK. Evolutionary computing and machine learning for discovering of low-energy defect configurations. *Npj Comput Mater*. 2021;7(1):71.

<https://doi.org/10.1038/s41524-021-00537-1>.

Sharma M, Bilgic M. Evidence-based uncertainty sampling for active learning. *Data Min Knowl Discov*. 2017;31(1):164-202.

<https://doi.org/10.1007/s10618-016-0460-3>.

Pandi A, Diehl C, Yazdizadeh Kharrazi A, Scholz SA, Bobkova E, Faure L, et al. A versatile active learning workflow for optimization of genetic and metabolic networks. *Nat Commun*. 2022;13(1):3876.

<https://doi.org/10.1038/s41467-022-31245-z>.

Morawietz T, Artrith N. Machine learning-accelerated quantum mechanics-based atomistic simulations for industrial applications. *J Comput Aided Mol Des*. 2021;35(4):557-86.

<https://doi.org/10.1007/s10822-020-00346-6>.

Koizumi A, Deffrennes G, Terayama K, Tamura R. Performance of uncertainty-based active learning for efficient approximation of black-box functions in materials science. *Sci Rep*. 2024;14(1):27019.

<https://doi.org/10.1038/s41598-024-76800-4>.

Elreedy D, Atiya AF, Shaheen SI. A novel active learning regression framework for balancing the exploration-exploitation trade-off. *Entropy (Basel)*. 2019;21(7):651.

<https://doi.org/10.3390/e21070651>.

Lew T, Sharma A, Harrison J, Byland A, Pavone M. Safe active dynamics learning and control: A sequential exploration–exploitation framework. *IEEE Trans Robot*. 2022;38(5):2888-907.

<https://doi.org/10.1109/TRO.2022.3154715>.

Li L, Chang J, Vakanski A, Wang Y, Yao T, Xian M. Uncertainty quantification in multivariable regression for material property prediction with Bayesian neural networks. *Sci Rep*. 2024;14(1):10543.

<https://doi.org/10.1038/s41598-024-61189-x>.

Lv X, Duan F, Jiang JJ, Fu X, Gan L. Deep active learning for surface defect detection. *Sensors (Basel)*. 2020;20(6):1650.

<https://doi.org/10.3390/s20061650>.

Hong C, Kim J, Kim J, Jung J, Ju S, Choi JM, et al. Applications and training sets of machine learning potentials. *Sci Technol Adv Mater Methods*. 2023;3(1):2269948.

<https://doi.org/10.1080/27660400.2023.2269948>.

Karandikar J, Chaudhuri A, Smith S, Schmitz T, Willcox K. Process window estimation in manufacturing through Entropy-Sigma active learning. *Manuf Lett*. 2022;34:87-92.

<https://doi.org/10.1016/j.mfglet.2022.09.001>.

Hino H, Eguchi S. Active learning by query by committee with robust divergences. *Inf Geom.* 2023;6(1):81-106.
<https://doi.org/10.1007/s41884-022-00081-x>.