**ORIGINAL RESEARCH** **Open access**

# Uncertainty-Conditioned Experiment Planning: A Conceptual Framework for AI-Guided Materials Exploration

Hiroshi Tanaka[1], Yuki Sato[1*], Kenji Mori[2], Rina Okabe[1]

## Abstract

Materials exploration faces persistent challenges stemming from vast chemical spaces, high experimental costs, and inherent uncertainties in predictive models. While machine learning has accelerated property prediction and guided candidate selection, conventional approaches often treat uncertainty as a uniform metric within fixed acquisition strategies. This conceptual paper introduces uncertainty-conditioned experiment planning (UCEP) as a novel theoretical framework for AI-guided materials discovery. UCEP reframes experiment planning as a dynamic process conditioned on the multidimensional character of uncertainty, integrating epistemic and aleatoric components, data-related biases, and model limitations into the steering logic. Rather than relying on static acquisition functions, the framework emphasizes adaptive interaction dynamics between uncertainty characterization and planning decisions, enabling context-sensitive trade-offs between exploration, exploitation, and bias mitigation. Drawing on interpretive insights from materials informatics and uncertainty quantification literature, UCEP highlights systems-level feedback structures that can enhance epistemic robustness and scientific efficiency without presupposing empirical outcomes. The framework offers analytical implications for rethinking how AI systems interpret and respond to uncertainty in iterative discovery cycles, contributing to more reflective and integrative AI-assisted materials research.

**Keywords** Epistemic uncertainty, Materials informatics, Uncertainty quantification, Active learning, Bayesian optimization, Experiment planning

*Correspondence:
Yuki Sato
yuki.sato@outlook.com
[1] Department of Intelligent Materials Engineering, Faculty of Engineering, University of Tokyo, Tokyo, Japan
[2] Department of AI-Driven Materials Discovery, Faculty of Information Science, Kyoto University, Kyoto, Japan

## Introduction

The discovery of advanced materials remains a cornerstone of technological progress, underpinning developments in energy storage, catalysis, electronics, structural engineering, and emerging sustainability-driven applications. Historically, materials innovation has relied on Edisonian approaches characterized by intuition-guided trial-and-error, incremental optimization, and localized experimental exploration. While such methods have yielded notable breakthroughs, they scale poorly in the face of the immense combinatorial design space defined by chemical composition, crystal structure, processing pathways, and operating conditions. Even with domain expertise, exhaustive exploration of this space is infeasible, rendering materials discovery inherently resource-intensive, time-consuming, and path-dependent [1].

Over the past decade, this challenge has catalyzed a paradigm shift toward data-driven materials science, enabled by advances in high-throughput computation, automated experimentation, and large-scale materials databases. Within this paradigm, artificial intelligence (AI) and machine learning (ML) have emerged as central instruments for navigating complex design spaces. Supervised and representation-learning models trained on curated datasets can predict key

materials properties—such as formation energies, band gaps, elastic moduli, and phase stability—with sufficient accuracy to inform pre-screening and candidate prioritization [2]. By compressing high-dimensional structure–property relationships into tractable surrogate models, these approaches promise to accelerate discovery by reallocating experimental effort toward regions of higher expected utility.

Beyond static prediction, AI has increasingly been embedded into closed-loop discovery workflows through active learning and Bayesian optimization. In such settings, models do not merely evaluate candidates but also propose subsequent experiments based on acquisition functions that balance predicted performance against uncertainty or information gain [3, 4]. Conceptually, this shift reframes materials discovery as a sequential decision-making process in which each experimental choice influences the evolving knowledge state of the system. These approaches have demonstrated proof-of-concept success in reducing sample complexity and accelerating convergence toward target properties, reinforcing the perception that uncertainty-aware AI can serve as an efficient experimental navigator.

However, despite their methodological sophistication, existing AI-guided experiment planning frameworks exhibit conceptual limitations that constrain their epistemic robustness. Most notably, uncertainty—although central to acquisition strategies such as expected improvement, probability of improvement, or upper confidence bound—is often operationalized as a scalar model output divorced from its epistemic origins [5]. Epistemic uncertainty arising from sparse data coverage, model misspecification, or representation mismatch is qualitatively distinct from aleatoric uncertainty rooted in irreducible physical variability or measurement noise. Treating these heterogeneous sources as interchangeable signals can obscure their implications for experimental design, leading to planning strategies that either over-exploit poorly understood regions or under-explore systematically biased domains.

These concerns are amplified by the structural properties of the materials datasets themselves. Public and proprietary datasets frequently reflect historical research priorities, computational convenience, or experimental accessibility, resulting in uneven coverage of chemical families, composition ranges, and property regimes [6]. Such biases are not merely statistical artifacts but epistemic constraints that shape what the model can learn and, by extension, what the discovery system can propose. When acquisition strategies operate atop biased representations, uncertainty estimates may misleadingly convey confidence or ignorance, reinforcing existing exploration pathways rather than challenging them. In this sense, AI-driven planning risks becoming self-referential, optimizing within inherited blind spots rather than expanding the epistemic horizon.

Philosophical analyses of scientific inference offer a valuable lens for interpreting these dynamics. Scientific decision-making under uncertainty has long been recognized as value-laden, particularly when evidence is incomplete, and errors carry asymmetric consequences [7]. Choices about acceptable risk, the breadth of exploration, and evidentiary thresholds implicitly encode non-epistemic values alongside epistemic aims such as accuracy or generalizability. In AI-assisted materials discovery, acquisition functions instantiate these value commitments algorithmically, privileging certain trade-offs—such as rapid optimization versus broad exploration—without making them explicit or adaptable to context. The absence of mechanisms to condition planning behavior on the type and structure of uncertainty thus represents a conceptual gap rather than a purely technical shortcoming.

This manuscript proposes uncertainty-conditioned experiment planning (UCEP) as a conceptual framework designed to address these limitations. Rather than treating experiment planning as optimization under a fixed objective, UCEP reconceptualizes planning as a conditioned process that responds dynamically to the evolving uncertainty landscape. By distinguishing uncertainty dimensions and linking them to differentiated planning logics, the framework foregrounds uncertainty as an active steering signal rather than a passive confidence measure. This perspective enables adaptive modulation of exploration strategies to align AI guidance with broader epistemic goals, including robustness, diversity of inquiry, and bias awareness.

Grounded in recent literature on materials informatics, uncertainty quantification, and data bias, the discussion emphasizes interaction dynamics and systems-level implications over algorithmic prescription. UCEP is presented not as

a replacement for existing optimization techniques but as an interpretive scaffold that clarifies how uncertainty, values, and planning decisions co-evolve in semi-autonomous discovery systems. Through this lens, the framework articulates a pathway toward more reflexive, transparent, and integrative human–AI collaboration in materials discovery—one that treats uncertainty not merely as noise to be minimized, but as structure to be understood and engaged.

# Theoretical Background and Literature Synthesis

## Materials informatics and data-driven discovery

Materials informatics integrates computational, statistical, and data-centric methods to accelerate materials discovery by leveraging large datasets and predictive modeling frameworks [8]. Large-scale high-throughput density functional theory (DFT) initiatives have played a foundational role in this transition, producing repositories such as the Materials Project and JARVIS-DFT, which collectively provide standardized property data across thousands of inorganic compounds [2]. These resources have enabled the systematic learning of structure–property relationships that were previously inaccessible through isolated experimental studies.

Building upon these datasets, machine learning models—including kernel-based regressors, ensemble methods, and increasingly graph neural networks—have demonstrated growing accuracy in mapping compositional and structural representations to target properties [9]. By encoding atomic environments and bonding topologies, these models facilitate rapid screening across vast chemical spaces, shifting discovery from sequential intuition-driven exploration toward algorithmically guided prioritization. Nevertheless, the epistemic reach of such models remains constrained by the quality, representativeness, and contextual assumptions embedded in their training data. As a result, predictive success within well-sampled regions often coexists with fragility under distributional shifts, underscoring that data-driven acceleration does not eliminate epistemic risk but rather redistributes it.

## Uncertainty quantification in machine learning for materials

Uncertainty quantification (UQ) has emerged as a critical component of machine learning applications in materials science, where overconfident predictions can misdirect experimental resources and obscure failure modes [5]. Recent methodological developments emphasize the importance of estimating instance-level uncertainty rather than relying solely on aggregate performance metrics. Approaches such as quantile regression, direct error modeling, and Gaussian process regression have been advocated to capture predictive dispersion at the level of individual material candidates [1]. Notably, direct error modeling has demonstrated balanced calibration across heterogeneous property classes, including elastic moduli and electronic band gaps, suggesting improved robustness in uncertainty estimation [1].

Bayesian paradigms further provide principled mechanisms for propagating uncertainty through probabilistic model formulations, enabling posterior distributions over predictions rather than point estimates [10]. These developments reinforce a growing consensus that effective UQ must disentangle aleatoric uncertainty—stemming from irreducible physical variability or measurement noise—from epistemic uncertainty arising from limited data, model inadequacy, or representational mismatch. Crucially, however, most existing frameworks treat uncertainty as an informational descriptor rather than a structuring variable for decision-making, leaving its implications for experiment planning largely implicit.

## Active learning and bayesian optimization in materials exploration

Active learning and Bayesian optimization formalize experiment selection as an iterative decision process under uncertainty and budget constraints [3, 4]. Within these paradigms, acquisition functions balance exploration—sampling regions of high uncertainty—and exploitation—refining candidates with high predicted performance—often implemented using Gaussian process surrogates [11]. Extensions to multi-objective optimization and noisy experimental settings have broadened applicability to realistic synthesis and characterization scenarios [12]. More recent formulations, such as

Bayesian algorithm execution, allow complex, user-defined criteria to be encoded into adaptive acquisition strategies, improving targeting of relevant material subsets [4].

Despite these advances, most implementations rely on static acquisition logics that presuppose a fixed relationship between uncertainty magnitude and sampling behavior. As discovery progresses and the uncertainty landscape evolves—through data accumulation, bias reinforcement, or representational shifts—acquisition strategies typically lack mechanisms to adapt their decision rationale accordingly. This rigidity limits the capacity of active learning systems to respond reflexively to changing epistemic conditions, particularly when uncertainty signals reflect structural blind spots rather than mere data sparsity.

## Data bias and its implications in materials informatics

Materials datasets frequently exhibit structural biases rooted in historical research trajectories, computational feasibility, and application-driven priorities. Overrepresentation of specific material families—such as thermoelectrics or semiconductors—produces uneven coverage of chemical space, shaping what models can reliably infer [6]. Empirical analyses using large experimental repositories, including Starrydata2, demonstrate that prediction accuracy degrades systematically for underrepresented material classes, revealing sharp boundaries in the model's applicability domain [6].

These biases are not merely statistical irregularities but epistemic constraints that influence downstream planning decisions. When biased datasets inform uncertainty estimates, models may exhibit unwarranted confidence in familiar regimes while masking ignorance elsewhere. Addressing such effects requires explicit attention to data provenance, diversity metrics, and bias-aware evaluation strategies, as well as mechanisms to prevent the reinforcement of inherited exploration pathways [13]. In this sense, bias functions as a latent steering force within data-driven discovery systems, shaping trajectories of inquiry in ways that often remain unexamined.

## Epistemic considerations and values in scientific inference

Uncertainty in scientific inference is inseparable from epistemic values, particularly in contexts where inductive risk necessitates judgments about the consequences of error [7, 14]. Decisions about acceptable uncertainty, evidentiary thresholds, and exploration breadth reflect not only technical considerations but also normative commitments regarding robustness, generalizability, and resource allocation. In AI-assisted discovery, these commitments are instantiated through acquisition strategies, model evaluation criteria, and stopping conditions, embedding value judgments into ostensibly objective workflows.

Recent philosophical reflections on the opacity of machine learning and inductive risk emphasize that non-epistemic values shape how uncertainty is interpreted and acted upon [15-19]. In materials discovery, this implies that different uncertainty profiles warrant different planning responses: high epistemic uncertainty may justify broader exploration to challenge existing assumptions, whereas dominance of aleatoric uncertainty may support focused exploitation within stable regimes. These interpretive dynamics highlight the absence of frameworks that explicitly condition experiment planning on the structure of uncertainty rather than its magnitude alone. Addressing this gap motivates the need for uncertainty-conditioned planning paradigms that align AI guidance with both epistemic aims and scientific values.

## Proposed conceptual framework: Uncertainty-Conditioned Experiment Planning (UCEP)

The proposed uncertainty-conditioned experiment planning (UCEP) framework reconceptualizes AI-guided materials exploration as a context-sensitive epistemic steering process rather than an optimization problem governed by fixed acquisition rules. At its core, UCEP advances the premise that experiment selection should be conditioned on the structure and composition of uncertainty, not merely on its magnitude. This shift reframes uncertainty from a passive confidence descriptor into an active organizing signal that shapes how discovery trajectories unfold.

UCEP distinguishes three analytically distinct yet interacting uncertainty dimensions: epistemic uncertainty, arising from data sparsity, representation limits, or model misspecification; aleatoric uncertainty, reflecting irreducible physical variability or measurement noise; and bias-related uncertainty, originating from structural imbalances in data coverage, historical sampling priorities, or inherited research focus. Rather than collapsing these dimensions into a single scalar, the framework treats them as a multidimensional uncertainty profile whose relative composition carries actionable epistemic meaning.

The framework operates through a closed-loop architecture in which the current knowledge state—comprising the surrogate model, accumulated dataset, and inferred material space—feeds into an uncertainty assessment module. This module decomposes predictive uncertainty into its constituent dimensions and augments them with indicators of domain coverage and bias exposure. Importantly, uncertainty is not evaluated solely in isolation but interpreted in relation to the representational completeness of the explored space.

The resulting uncertainty profile is passed to a conditioning layer that serves as an interpretive mediator rather than a deterministic controller. Instead of selecting actions via rigid thresholds or predefined objectives, the conditioning layer translates uncertainty compositions into relative planning tendencies. For example, dominance of epistemic uncertainty—particularly when aligned with underrepresented regions of chemical space—elicits diversity-oriented sampling strategies to expand representational support. Conversely, the prevalence of aleatoric uncertainty within well-characterized regimes justifies exploitation-oriented refinement that respects intrinsic variability rather than pursuing diminishing epistemic returns. Bias-related signals trigger compensatory adjustments, such as rebalancing sampling emphasis or introducing corrective diversification to counteract overrepresentation.

These conditioned interactions generate adaptive feedback structures across the discovery lifecycle. Early phases, characterized by high epistemic uncertainty and uneven coverage, naturally favor exploratory and corrective dynamics. As epistemic uncertainty contracts through targeted data acquisition, the framework transitions toward exploitation while remaining responsive to emergent bias signals or shifts in uncertainty composition. Crucially, this evolution is not governed by stage-based heuristics but emerges from continuous reinterpretation of the uncertainty landscape.

## Conceptual formalization of uncertainty conditioning

To conceptually formalize the conditioning logic without prescribing algorithmic implementation, UCEP represents predictive uncertainty as a multidimensional epistemic profile rather than a scalar estimate:

$$f\{U\}(x) = \left\{ U_{epi}(X), \quad ; \quad U_{ale}(X), \quad ; \quad U_{bias}(X) \right\} \tag{1}$$

where $Uepi$ denotes epistemic uncertainty arising from data sparsity or model limitations, $Uale$ captures irreducible aleatoric variability, and $Ubias$ reflects structural imbalances in domain coverage. Experiment planning is then conditioned on this composite uncertainty structure in relation to the evolving system knowledge state:

$$P(a \mid x) = C\left(f\{U\}(x), \quad ; \quad K_t\ big\right) \tag{2}$$

where $P(a \mid x)$ represents the conditioned planning tendency associated with candidate action a, $K_t$ denotes the discovery system's knowledge state at iteration $t$, and $C(\cdot)$ functions as an interpretive conditioning operator translating uncertainty composition into planning orientation. These expressions are conceptual rather than prescriptive, formalizing the adaptive steering dynamics embedded in the UCEP framework.

By embedding uncertainty interpretation directly into planning logic, UCEP positions experiment selection as a reflexive process aligned with epistemic goals of robustness, diversity, and resilience. The framework does not prescribe specific algorithms or metrics; rather, it provides a conceptual scaffold for understanding how AI-guided discovery systems can

adapt their exploratory behavior in response to evolving uncertainty conditions without presupposing optimality or completeness. As shown in the proposed conditioning architecture for uncertainty-aware planning **(Figure 1)**, the knowledge state of a materials discovery system is dynamically updated not by direct optimization, but through an adaptive loop where quantified uncertainty profiles condition and weight the selection of planning strategies.

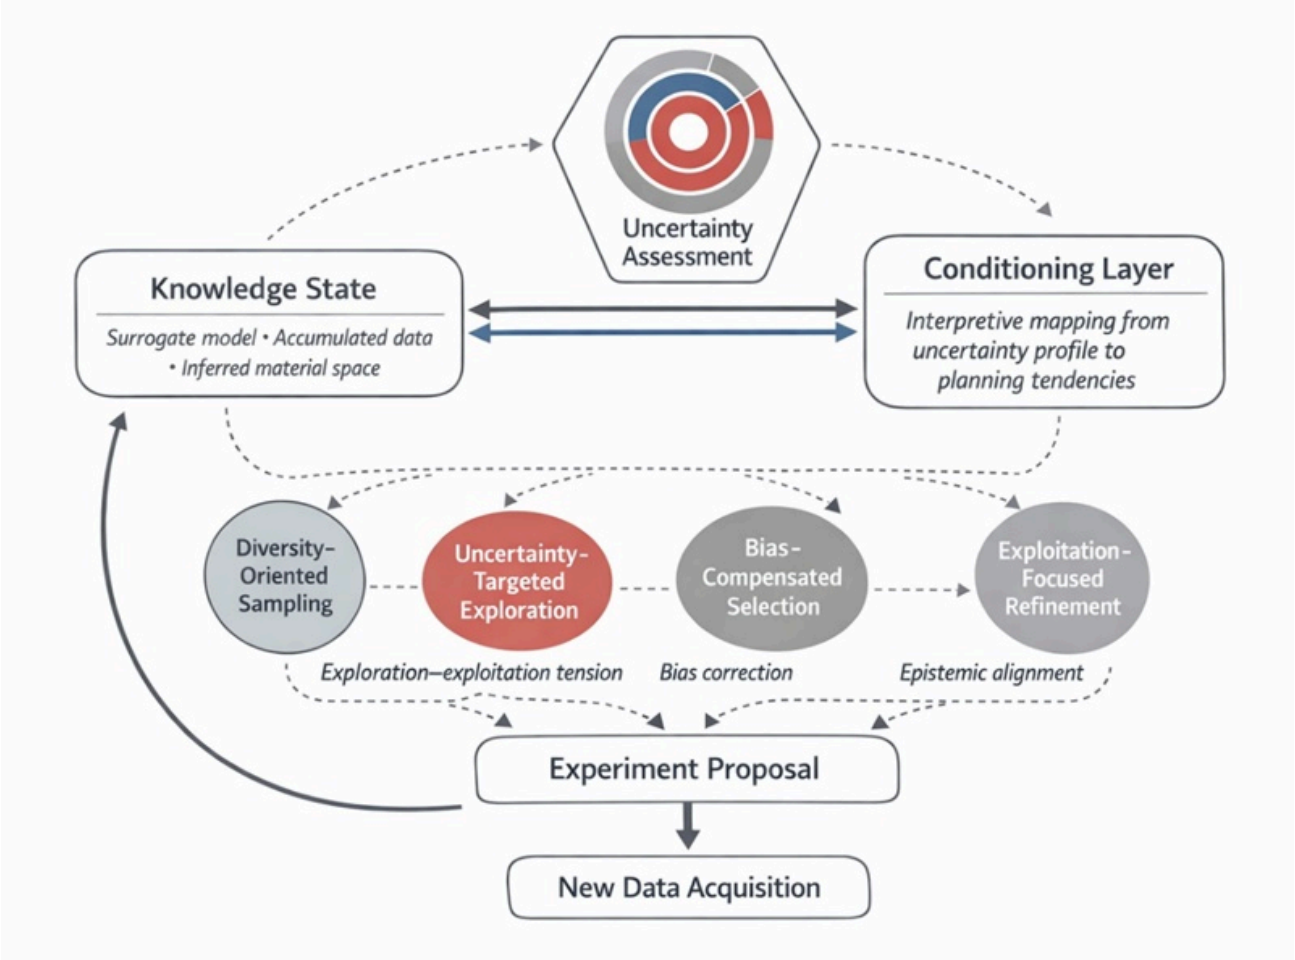


**Figure 1.** Conditioning architecture for uncertainty-aware planning in materials discovery.

This conceptual architecture fosters systems-level insights into how uncertainty can steer discovery toward greater epistemic resilience and inclusivity, without presupposing specific implementations or outcomes.

## Analytical implications

The Uncertainty-Conditioned Experiment Planning (UCEP) framework carries substantive analytical implications for how AI systems participate in materials discovery as epistemic agents rather than mere optimization engines. By conditioning experiment selection on a decomposed uncertainty profile, UCEP shifts interpretive emphasis away from maximizing a single scalar objective—such as expected improvement or information gain—toward navigating multidimensional trade-offs embedded within the uncertainty landscape. As outlined in **Table 1**, UCEP interprets dominant uncertainty regimes as signals for distinct planning responses, enabling adaptive transitions between exploration, exploitation, and bias compensation across iterative discovery cycles.

**Table 1.** Uncertainty regimes and conditioned planning responses in the UCEP framework

| Dominant uncertainty regime | Primary source characteristics | Epistemic risk if ignored | Conditioned planning response in UCEP | Systems-level effect on discovery trajectory |
|---|---|---|---|---|
| Epistemic uncertainty | Sparse data coverage; model misspecification; representation mismatch; extrapolative regimes | Overconfidence in poorly supported predictions; premature convergence; exclusion of underexplored chemical spaces | Diversity-oriented sampling; coverage expansion; exploratory experimentation targeting representational gaps | Broadens chemical space exploration; reduces path dependence; strengthens epistemic foundations |
| Aleatoric uncertainty | Irreducible physical variability; measurement noise; stochastic synthesis outcomes | Inefficient exploration attempts to reduce irreducible variance; wasted experimental resources | Exploitation-focused refinement; localized sampling to improve precision without over-expansion | Improves estimation accuracy; respects intrinsic variability; stabilizes prediction confidence |
| Bias-related uncertainty | Structural dataset imbalance; historical overrepresentation; skewed property regimes | Reinforcement of inherited blind spots; misleading uncertainty estimates; inequitable exploration | Bias-compensated selection; reweighting strategies; active diversification of underrepresented domains | Enhances inclusivity; improves generalizability; mitigates epistemic blind spots |
| Mixed/transitional regimes | Concurrent epistemic, aleatoric, and bias signals evolving over iterations | Misaligned planning if treated with static acquisition logic | Adaptive weighting across strategies via conditioning layer; context-sensitive trade-offs | Enables smooth transitions across discovery phases; maintains epistemic resilience |

This reframing exposes interaction dynamics that are obscured in conventional active learning formulations. Epistemic uncertainty, commonly associated with sparse sampling, representational gaps, or extrapolative regimes, is reinterpreted as a signal for structural exploration rather than as localized uncertainty-chasing. Under UCEP, high epistemic uncertainty motivates sampling strategies that deliberately expand chemical space coverage, challenge inherited data distributions, and reduce path dependence in discovery trajectories [1, 6]. In contrast, aleatoric uncertainty—when dominant within well-supported domains—justifies exploitation-oriented refinement that acknowledges irreducible variability without expending resources on futile epistemic reduction [5, 10, 20-24].

The framework also introduces adaptive feedback structures that evolve over iterative cycles. As epistemic uncertainty contracts through targeted exploration, conditioned planning naturally transitions toward exploitation. Importantly, this transition is neither monotonic nor irreversible: re-emergence of bias indicators or model discrepancies can recalibrate planning behavior, preventing premature convergence. This contrasts with static acquisition functions, which may unintentionally reinforce biases by repeatedly sampling high-uncertainty regions within already overrepresented domains, thereby amplifying epistemic blind spots [6, 25-29].

At a systems level, UCEP foregrounds bias mitigation as an endogenous component of experiment planning rather than a downstream corrective step. By embedding bias exposure metrics directly into the conditioning layer, the framework enables compensatory adjustments that counteract overrepresentation and uneven domain coverage [6, 13, 30-32]. Bias thus becomes an interpretable uncertainty signal that informs steering decisions, rather than a latent defect discovered post hoc. This integration positions uncertainty characterization as a proxy for epistemic health, guiding discovery toward robustness and generalizability without assuming ideal data distributions or complete representations.

Collectively, these analytical implications reposition AI-guided materials discovery as a reflexive process in which uncertainty, values, and planning co-evolve. UCEP does not seek to optimize discovery speed alone, but to cultivate epistemic resilience—supporting exploration strategies that are adaptive, inclusive, and aligned with the long-term scientific objective of producing reliable and transferable materials knowledge.

# Results and Discussion

The uncertainty-conditioned experiment planning (UCEP) framework reframes the role of AI in materials exploration by shifting emphasis from prescriptive optimization toward contextual interpretation and epistemic steering. Conventional active learning and Bayesian optimization approaches typically operationalize uncertainty through acquisition functions that assume a uniform semantic meaning across discovery stages [3, 4, 11]. While effective under controlled assumptions, such formulations implicitly treat uncertainty as a homogeneous quantity whose influence on planning remains invariant over time. In practice, however, uncertainty in materials discovery is heterogeneous in its origins, consequences, and epistemic significance. When these distinctions are ignored, acquisition strategies risk misalignment between planning behavior and the actual knowledge gaps they are intended to address.

UCEP responds to this limitation by treating uncertainty as a conditioned input rather than a scalar objective modifier. By explicitly distinguishing epistemic, aleatoric, and bias-related components, the framework enables planning logic to adapt to the evolving structure of ignorance rather than solely to its magnitude [5, 6]. This reinterpretation supports more nuanced trade-offs across the discovery lifecycle. In early stages, when epistemic uncertainty dominates due to sparse sampling or representational incompleteness, UCEP prioritizes exploratory dynamics to expand coverage and challenge inherited assumptions. As discovery progresses and epistemic uncertainty contracts, the framework permits a gradual reorientation toward exploitation when aleatoric uncertainty becomes the primary limiting factor, aligning refinement efforts with intrinsic physical variability rather than reducible knowledge deficits [10, 12, 33-35].

Beyond methodological adaptation, UCEP engages directly with epistemic questions concerning how scientific inference is guided under uncertainty. Experiment selection is not a neutral act: it distributes epistemic risk, allocates finite resources, and shapes which hypotheses become testable [7, 14]. Acquisition choices, therefore, embed value judgments regarding exploration breadth, acceptable error, and tolerance for uncertainty. Conventional frameworks often obscure these judgments behind mathematically optimized criteria, rendering their normative commitments implicit. In contrast, UCEP surfaces these commitments by embedding them within uncertainty-conditioned mappings, making explicit how preferences for robustness, diversity, or bias mitigation influence planning behavior.

This transparency has important implications for the epistemic integrity of AI-assisted discovery workflows. When uncertainty is interpreted in terms of its source and consequences, planning decisions become sensitive to the differential costs of error across regimes. For instance, high epistemic uncertainty in underrepresented chemical spaces carries risks of systematic exclusion and overconfidence if ignored. In contrast, aleatoric uncertainty in well-characterized domains primarily affects precision rather than validity [7, 15]. By conditioning planning responses accordingly, UCEP aligns AI guidance with scientific values that prioritize generalizability, reflexivity, and long-term knowledge accumulation rather than short-term optimization gains.

At a systems level, the framework also reconceptualizes bias mitigation as an endogenous feature of experiment planning rather than a downstream corrective exercise. Structural biases in materials datasets—arising from historical research focus, computational tractability, or application-driven priorities—shape both model predictions and uncertainty estimates [6]. When left unaddressed, these biases can be amplified through iterative sampling, reinforcing path dependence, and narrowing the design space explored. UCEP incorporates bias exposure and domain coverage as integral components of the uncertainty profile, allowing planning strategies to compensate proactively through diversification or rebalancing mechanisms [6, 13]. This integration highlights how uncertainty characterization can function as a diagnostic of epistemic health, signaling when discovery trajectories risk becoming insular or brittle.

Nonetheless, the framework introduces its own conceptual challenges. The interpretive mappings within the conditioning layer require careful design to avoid unintended distortions, such as overreacting to noisy uncertainty decompositions or privileging certain dimensions at the expense of others. Moreover, the practical realization of UCEP depends on the availability of meaningful indicators for epistemic uncertainty, aleatoric variability, and bias exposure—areas that remain active and unresolved within materials informatics [6, 13]. These limitations underscore that UCEP is not a turnkey solution but a conceptual scaffold whose effectiveness depends on thoughtful instantiation and critical oversight.

Despite these challenges, UCEP offers clear advantages over rigid acquisition paradigms by foregrounding adaptability, context-awareness, and epistemic reflexivity. Rather than prescribing how AI systems should optimize discovery, the framework clarifies how they can reason about uncertainty in ways that respect its heterogeneous nature. In doing so, UCEP contributes to a broader rethinking of AI-guided materials discovery as a co-evolving human–machine enterprise, where planning decisions are shaped as much by epistemic judgment as by predictive performance.

# Conclusion

Uncertainty-conditioned experiment planning (UCEP) introduces a novel conceptual lens for AI-guided materials exploration by reconceptualizing experiment selection as a process dynamically shaped by multidimensional uncertainty profiles. By distinguishing epistemic, aleatoric, and bias-related dimensions—and conditioning planning strategies on their relative composition—UCEP moves beyond fixed acquisition rules toward adaptive steering that is responsive to the evolving knowledge state of the discovery system.

This reframing yields interpretive insights into how uncertainty structures the dynamics of interaction among data, models, and experimental choices. By aligning exploratory behavior with epistemic gaps, exploitation with irreducible variability, and bias mitigation with coverage imbalances, the framework promotes discovery trajectories that are more robust, inclusive, and resilient. Importantly, UCEP does not position uncertainty merely as a technical nuisance to be minimized, but as a meaningful signal that guides how scientific inquiry unfolds over time.

Grounded in advances in materials informatics, uncertainty quantification, active learning, and bias awareness, the framework extends existing paradigms by emphasizing conditioned, context-sensitive logic rather than optimization under static assumptions. In doing so, it foregrounds the role of values, inductive risk, and interpretive judgment in AI-assisted discovery, contributing to more transparent and reflexive human–AI collaboration.

While conceptual in scope, UCEP provides a foundation for future theoretical refinement and methodological exploration. Its value lies not in prescribing specific algorithms, but in clarifying how uncertainty, bias, and planning can be coherently integrated into discovery workflows that respect the epistemic complexity of materials systems. As AI continues to shape the pace and direction of materials research, frameworks such as UCEP offer critical guidance for ensuring that acceleration is accompanied by epistemic responsibility and scientific integrity.

# Acknowledgements

None

# Conflict of interest

None

# Financial support

None

# Ethics statement

None

Received: 27 Apr 2025 Revised: 21 Jun 2025 Accepted: 13 Jul 2025
Published online: 18 January 2026



# References

Chen C, Ong SP. Uncertainty quantification for machine learning interatomic potentials in materials. npj Comput Mater. 2024;10:83.

Merchant A, Batzner S, Schoenholz SS, Aykol M, Cheon G, Cubuk ED. Scaling deep learning for materials discovery. Nature. 2023;624(7990):80-5.

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.

Zhong M, Tran K, Min Y, et al. Accelerating materials discovery with bayesian optimization. Nat Rev Mater. 2022;7:358-70.

Chen C, Ong SP. Direct error modeling for improved uncertainty quantification in regression tasks: Application to battery materials. Joule. 2022;6:1234-51.

Kumagai Y, Ando K, Tanaka I. Effects of data bias on machine-learning-based material discovery using experimental property data. Sci Technol Adv Mater. 2022;23:1-12.

Henschen T. How strong is the argument from inductive risk? Eur J Philos Sci. 2021;11(3):92.

Butler KT, Davies DW, Cartwright H, et al. Materials informatics. J Phys Condens Matter. 2021;33:203001.

Park CW, Wolverton C. Developing an improved crystal graph convolutional neural network framework for accelerated materials discovery. Phys Rev Mater. 2020;4(6):063801.

Heid E, Green WH. Characterizing uncertainty in machine learning for chemistry. J Chem Inf Model. 2023;63:4012-25.

Rao Z, Bajpai A, Zhang H. Active learning strategies for the design of sustainable alloys. Philos Trans A. 2024;382(2284):20230242.

Yang Y, Jiménez-Negrón OA, Kitchin JR. Machine-learning accelerated geometry optimization in molecular simulation. J Chem Phys. 2021;154(23).

Oviedo F, Ferres JL, Buonassisi T, Butler KT. Interpretable and explainable machine learning for materials science and chemistry. Acc Mater Res. 2022;3:597-607.

Rudner R, Toner J. Values and inductive risk in machine learning modelling: The case of binary classification models. Eur J Philos Sci. 2021;11:78.

Lam H, Rescorla M. On predicting recidivism: Epistemic risk, tradeoffs, and values in machine learning. Can J Philos. 2022;52:321-41.

Morgan JN. Calibration after bootstrap for accurate uncertainty quantification in regression models. npj Comput Mater. 2022;8:112.

Peerless JS, et al. Uncertainty quantification in battery materials discovery. J Phys Chem C. 2020;124:14567-78.

Mosavi A, et al. Materials informatics and sustainability. Data Centric Eng. 2021;2:e12.

Wang X, Zhao Y, Pourpanah F. Recent advances in deep learning. Int J Mach Learn Cybern. 2020;11(4):747-50.

Agrawal A, Choudhary A. Perspective: Materials informatics and big data. APL Mater. 2020;8:080401.

Hutchinson ML, Antono E, Gibbons BM, Paradiso S, Ling J, Meredig B. Overcoming data scarcity in materials science with transfer learning. Comput Mater Sci. 2021;186:110012.

Childs CM, Washburn NR. Integrating domain knowledge into machine learning for materials. MRS Bull. 2021;46:765-71.

Mosavi A, et al. State of the art survey of deep learning for uncertainty quantification in materials. Eng Appl Artif Intell. 2022;112:104845.

Oviedo F, et al. Explainable machine learning in materials science. npj Comput Mater. 2022;8:204.

Hüllermeier E, Waegeman W. Aleatoric and epistemic uncertainty in machine learning. Mach Learn. 2021;110:457-506.

Choudhary K, et al. Materials property prediction with uncertainty quantification: A benchmark study. Appl Phys Rev. 2023;10:021409.

Wen M, Tadmor EB. Uncertainty quantification for neural network potential foundation models. npj Comput Mater. 2020;6:175.

Liu Z, Singh A, Li Y. Feature importance and uncertainty quantification of machine learning model in materials science. Ind Eng Chem Res. 2023;62:14567-78.

Furat O, et al. Uncertainty quantification using generative adversarial networks. npj Comput Mater. 2022;8:68.

Yin P, Panapitiya G. Evaluating uncertainty-based active learning for accelerating the generalization of molecular property prediction. J Chem Inf Model. 2023;63:5123-34.

Pyzer-Knapp EO, Pitera JW, Staar PW, Takeda S, Laino T, Sanders DP, et al. Accelerating materials discovery using artificial intelligence, high performance computing and robotics. npj Comput Mater. 2022;8(1):84.

Honarmandi P, Arróyave R. Uncertainty quantification and propagation in computational materials science and simulation-assisted materials design. Integr Mater Manuf Innov. 2020;9(1):103-43.

Liu L, Zhou W, Guan K, Peng B, Xu S, Tang J, et al. Knowledge-guided machine learning can improve carbon cycle quantification in agroecosystems. Nat Commun. 2024;15(1):357.

Morgan D, Jacobs R. Opportunities and challenges for machine learning in materials science. Annu Rev Mater Res. 2020;50:71-103.

Lookman T, et al. Benchmarking active learning strategies for materials optimization and discovery. Oxf Open Mater Sci. 2022;2:itac006.