

ORIGINAL RESEARCH

Open access

A Conceptual Framework for Uncertainty-Aware Active Learning in Compositional Design Spaces with Multiple Objectives

Isabella Garcia^{1*}, Diego Herrera¹

Abstract

Compositional design spaces in materials engineering, particularly for alloys and multi-component systems, present exceptionally high-dimensional and sparsely populated landscapes that challenge conventional discovery workflows. A typical five-element alloy system sampled at 10% concentration increments can encompass millions of possible compositions, rendering exhaustive evaluation infeasible. Active learning has emerged as an essential paradigm for navigating these spaces efficiently, yet standard implementations often fail to address the simultaneous demands of multiple conflicting objectives—such as balancing mechanical strength against ductility, electrical conductivity against thermal stability, or performance against material cost—while properly accounting for uncertainty and practical constraints. This conceptual framework introduces an uncertainty-aware active learning approach tailored specifically for compositional design spaces with multiple objectives. The framework comprises four core components: robust uncertainty quantification that distinguishes epistemic from aleatoric contributions across heterogeneous composition regions, multi-objective acquisition functions grounded in Pareto optimization to explicitly explore trade-offs, systematic constraint handling for both hard thermodynamic limits and soft economic or manufacturability requirements, and adaptive sequential sampling strategies that prioritize informative candidates under limited evaluation budgets. The key innovation lies in a Pareto-based acquisition mechanism that integrates uncertainty estimates to balance exploration of uncertain regions, exploitation of promising trade-offs, and navigation of sparse high-dimensional manifolds. By embedding these elements into a unified conceptual structure, the framework provides design principles that enable more reliable, interpretable, and scalable materials discovery. It shifts the focus from single-objective optimization to holistic decision-making under uncertainty, ensuring that active learning not only accelerates identification of high-performing compositions but also respects real-world engineering constraints inherent to multi-component systems. The proposed approach is expected to guide both computational and experimental campaigns in alloy design, high-entropy materials, and mixed ionic conductors, ultimately fostering more sustainable and high-performance materials.

Keywords Uncertainty quantification, Active learning, Bayesian optimization, Multi-objective optimization, Compositional design spaces, Pareto frontier

*Correspondence:

Isabella Garcia
isabella.garcia@gmail.com

¹ Department of Intelligent Materials Engineering, Faculty of Engineering, University of Buenos Aires, Buenos Aires, Argentina

Introduction

Compositional design spaces—encompassing alloys, multi-component ceramics, and mixed ionic conductors—are vast. A 5-element system with 10% concentration steps has millions of compositions. Active learning is essential for navigating these spaces. But compositional active learning faces unique challenges: multiple objectives (strength, conductivity, cost, stability), high-dimensional sparsity, and heteroscedastic uncertainty. This paper proposes a conceptual framework for

uncertainty-aware active learning in compositional design spaces with multiple objectives, integrating uncertainty quantification, multi-objective acquisition, constraint handling, and sequential sampling [1-4].

The scale of these design spaces arises directly from the combinatorial nature of elemental mixing. Even modest systems with three to six principal elements, each allowing continuous concentration variations within the composition simplex, generate an astronomical number of candidate points. Traditional trial-and-error or grid-based screening quickly becomes impractical, consuming prohibitive computational or experimental resources. Recent studies on aluminum-based and high-entropy alloys have repeatedly demonstrated that conventional methods leave the majority of the space unexplored. Active learning addresses this by iteratively selecting the most informative compositions for evaluation, thereby focusing limited budgets on regions likely to yield breakthroughs [5-8].

Yet standard active learning pipelines, developed primarily for low-dimensional molecular or discrete spaces, encounter fundamental limitations when applied to compositional materials. Multiple objectives rarely align; maximizing one property frequently degrades another, necessitating explicit trade-off analysis rather than scalar optimization. High-dimensional sparsity exacerbates the issue because most compositions lie far from any previously evaluated point, leading to unreliable surrogate predictions. Uncertainty in these regions is not uniform: near-equiatomic or well-sampled zones exhibit lower epistemic uncertainty, while dilute or extreme compositions display markedly higher uncertainty due to data scarcity and model extrapolation demands [2, 9, 10].

Heteroscedastic uncertainty further complicates decision-making. Predictive error varies systematically across the composition landscape, with certain regions exhibiting larger aleatoric noise stemming from inherent simulation or measurement variability. Ignoring this heteroscedasticity can produce poorly calibrated acquisition decisions that either waste evaluations on already well-characterized areas or overlook promising but uncertain candidates. Practical constraints compound the difficulty. Thermodynamic stability, synthesizability, elemental availability, cost, and toxicity impose both hard filters and soft preferences that must be respected throughout the search [11-14].

The proposed conceptual framework directly tackles these intertwined challenges through a structured integration of four interlocking components. Uncertainty quantification supplies calibrated estimates that separate reducible epistemic uncertainty from irreducible aleatoric noise, ensuring acquisition functions remain reliable even in sparsely sampled territories. Multi-objective acquisition functions operate on the Pareto frontier, scoring candidates according to their potential to expand the set of non-dominated solutions while explicitly incorporating uncertainty. Constraint handling mechanisms enforce hard feasibility rules while folding soft preferences into the decision process. Sequential sampling orchestrates the iterative loop, beginning with space-filling initialization and proceeding through model updating until convergence or budget exhaustion [15-19].

A central innovation of the framework is its Pareto-based acquisition strategy that simultaneously balances exploration, exploitation, and objective trade-offs. Rather than collapsing multiple goals into a single weighted scalar—an approach prone to arbitrary weighting choices—the framework maintains the full Pareto front and quantifies improvement in hypervolume or frontier coverage. This preserves transparency about inherent compromises and prevents premature convergence to suboptimal regions. The framework also articulates explicit design principles that practitioners can adopt when tailoring active learning campaigns to specific material families. These principles emphasize calibration of uncertainty models to compositional manifolds, adoption of Pareto-aware scoring, systematic constraint propagation, and diversity-preserving batch selection for parallel workflows [20-23].

By grounding the entire process in uncertainty awareness, the framework mitigates the risk of overconfident selections in high-dimensional spaces and promotes more robust materials discovery pipelines. It aligns computational surrogates with experimental realities and provides a conceptual scaffold that future algorithmic developments can build upon. Ultimately, the framework aims to transform active learning from a heuristic search tool into a principled decision engine capable of handling the full complexity of modern compositional materials engineering [3, 24-26].

Figure 1 presents the conceptual architecture through which uncertainty quantification, Pareto-based acquisition, constraint handling, and sequential sampling are integrated to guide decision-making in high-dimensional compositional design spaces.




Figure 1. Conceptual architecture of uncertainty-aware multi-objective active learning in compositional design spaces

Challenges of Compositional Design Spaces

Compositional design spaces introduce six interlocking challenges that distinguish them from simpler optimization settings and necessitate a specialized active learning framework. High dimensionality arises because technologically relevant alloys and multi-component systems typically involve three to six or more principal elements, each with continuous concentration variables constrained to sum to unity; the resulting curse of dimensionality expands the searchable volume exponentially, rendering uniform coverage impossible and forcing sparse sampling that leaves vast regions without nearby training data [2-4]. This sparsity compounds when multiple conflicting objectives—mechanical strength, ductility, conductivity, phase stability, cost, and environmental impact—must be optimized simultaneously, as improvements in one property typically degrade another and produce a Pareto frontier of non-dominated trade-offs rather than a single optimum that scalarized single-objective methods can adequately reveal [14-18].

A related implication is the emergence of heteroscedastic uncertainty, where predictive variance is not uniform but systematically higher in dilute limits and extreme compositions that suffer from data scarcity, while aleatoric noise from simulations or experiments also varies with composition and undermines homoscedastic assumptions in acquisition functions [9, 10, 25]. Beyond this immediate concern lie hard and soft constraints—thermodynamic stability thresholds, synthesizability limits, exclusion of toxic elements, cost targets, and manufacturability preferences—that must be embedded directly into candidate selection, since violating hard constraints renders compositions unusable irrespective of predicted performance [11-13, 23]. Under these conditions, the most promising solutions often require deliberate extrapolation far from initial data clusters, yet standard mechanisms favor local exploitation and risk trapping the search away from globally superior regions of the simplex [2, 7, 27].

The prohibitive cost of evaluation further intensifies these tensions, as density functional theory calculations or experimental campaigns consume hours to weeks per point within budgets typically limited to 10–100 evaluations, demanding that each acquisition maximize information gain toward the Pareto frontier [1, 6, 24, 28]. In practice, these interlocking features explain why off-the-shelf active learning algorithms underperform in compositional materials design, prompting the proposed framework to address them through targeted conceptual components that transform active learning into a decision process attuned to dimensionality, objective multiplicity, uncertainty heterogeneity, constraints, extrapolation demands, and evaluation budgets [5, 19–22].

Table 1 analytically maps each challenge of compositional design spaces onto the specific framework mechanism required to prevent failure of conventional active learning strategies.

Table 1. Analytical mapping between compositional design-space challenges and the framework mechanisms required to address them

Challenge in compositional design spaces	Why conventional active learning fails here	Required framework mechanism	Decision consequence if mechanism is absent	Conceptual role in the proposed framework
High dimensionality of the composition simplex	Local neighborhoods become unrepresentative, uniform screening becomes impossible, and most candidate regions remain effectively invisible to the surrogate	Space-aware uncertainty quantification combined with adaptive sequential sampling	The search collapses into narrow, already sampled regions and cannot interrogate the broader composition manifold	Converts combinatorial scale into a navigable decision problem by prioritizing informative rather than exhaustive evaluations
Multiple conflicting objectives	Single-score optimization hides trade-offs and forces arbitrary weighting choices that can distort materials selection	Pareto-based acquisition using non-dominated frontier expansion	Apparent optima are produced that reflect scalarization choices rather than genuine engineering compromise structures	Preserves the full geometry of trade-offs among strength, ductility, conductivity, cost, stability, and sustainability
Heteroscedastic uncertainty	Homoscedastic assumptions miscalibrate acquisition, causing overconfidence in noisy regions and under-exploration in sparsely sampled ones	Explicit separation of epistemic and aleatoric uncertainty through calibrated uncertainty models	Candidate ranking becomes statistically unreliable because expected gains are not distinguished from noise-dominated predictions	Establishes uncertainty as a structured property of the compositional manifold rather than a generic model output
Hard and soft constraints	Post-hoc filtering wastes evaluations on infeasible candidates and obscures the difference between non-negotiable feasibility and negotiable preference	Hard-constraint filtering plus soft-constraint integration into ranking or objective structure	The system selects high-performing but unusable candidates or ignores economically meaningful trade-offs	Embeds engineering realism directly into the search logic
Need for deliberate extrapolation	Pure exploitation remains trapped near known data clusters and fails to probe	Uncertainty-weighted acquisition that	Discovery remains conservative and	Makes controlled extrapolation a legitimate and

	sparsely sampled but potentially superior regions	rewards informative frontier expansion beyond current data density	misses globally valuable compositions	deliberate part of active learning rather than an accidental byproduct
Severe evaluation-cost limits	Expensive DFT or experimental budgets make random or redundant sampling prohibitive	Budget-aware sequential or batch sampling with diversity preservation	Resources are exhausted before meaningful Pareto-front advancement occurs	Forces every evaluation to justify itself in terms of information gain, trade-off improvement, or feasibility discovery

Uncertainty Quantification for Compositional Spaces

Uncertainty-aware active learning is defined as an iterative selection process where acquisition decisions are driven by well-calibrated uncertainty estimates that explicitly separate epistemic (reducible with more data) from aleatoric (inherently irreducible) components. In compositional spaces this separation is indispensable because epistemic uncertainty dominates in unexplored regions while aleatoric contributions vary with local chemistry and processing conditions. Without calibrated uncertainty that respects the geometry of the composition simplex, acquisition functions risk selecting candidates whose predicted improvements are statistically meaningless [9, 10].

Effective uncertainty quantification for compositional spaces must satisfy three core requirements. First, estimates must remain calibrated across the entire manifold, avoiding overconfidence in data-rich zones or excessive pessimism in sparse zones. Second, the method must sharply distinguish near-training compositions (low epistemic uncertainty) from far-from-training compositions (high epistemic uncertainty). Third, it must accommodate heteroscedastic noise, recognizing that simulation or measurement variance differs systematically between equiatomic and dilute regimes [1, 25].

Several uncertainty quantification methods offer partial solutions, each with distinct advantages and limitations. Deep ensembles provide simplicity and straightforward parallel training yet do not naturally separate epistemic from aleatoric uncertainty. Conformal prediction supplies finite-sample coverage guarantees but typically yields only marginal intervals that ignore composition-specific heteroscedasticity. Bayesian neural networks can, in principle, deliver principled separation of the two uncertainty types, yet their training remains computationally demanding and prone to instability in high-dimensional inputs. Distance-based approaches are computationally inexpensive and geometrically interpretable on the simplex, but they lack probabilistic calibration and therefore cannot be used directly in acquisition functions that require expected improvement calculations [9, 10].

The framework therefore recommends a hybrid strategy: an ensemble of surrogate models to capture epistemic uncertainty combined with a learned noise model for aleatoric uncertainty. Total predictive uncertainty is obtained by adding the variances in quadrature, yielding a single scalar that faithfully represents overall risk at any composition. This construction preserves calibration across both dense and sparse regions and naturally reflects the heteroscedastic character of compositional data. The resulting uncertainty estimates feed directly into downstream acquisition functions, ensuring that candidates are selected not merely for predicted performance but for their capacity to reduce epistemic uncertainty while respecting inherent noise levels [11-13].

By embedding this uncertainty quantification strategy at the heart of the framework, acquisition decisions become robust to the unique statistical structure of compositional design spaces. Practitioners gain confidence that selected candidates genuinely advance knowledge rather than merely confirming already well-understood regimes [2-4].

Multi-Objective Acquisition Functions

A multi-objective acquisition function is defined as a scoring mechanism that evaluates candidate compositions according to their joint potential to improve multiple objectives and their associated uncertainties, thereby guiding the next evaluation toward the most valuable addition to the Pareto frontier. Single-objective acquisition functions such as expected improvement, upper confidence bound, or probability of improvement are insufficient because they collapse trade-offs into a scalar and cannot reveal the structure of conflicting goals [14, 15, 17].

Several multi-objective strategies have been proposed, each offering different trade-offs between expressiveness and computational cost. Scalarization combines individual objective improvements through weighted sums. While conceptually simple, this approach requires arbitrary weight choices that obscure underlying trade-offs and demand repeated re-optimization when stakeholder priorities shift. Pareto-based acquisition avoids scalarization entirely by maintaining the current non-dominated front and scoring candidates according to their ability to expand that front [16, 18–20].

A particularly effective Pareto-based metric is hypervolume improvement, which quantifies the additional volume dominated by the Pareto front after hypothetically adding a candidate. This measure naturally accounts for all objectives simultaneously and provides an intuitive geometric interpretation of progress. Uncertainty-weighted Pareto strategies further refine the selection by using lower confidence bounds for exploitation and upper confidence bounds for exploration, producing a Pareto-UCB variant that favors candidates offering the greatest potential frontier advancement under uncertainty. Multi-objective Thompson sampling draws samples from the joint posterior, computes the Pareto front for each sample, and selects the candidate that most frequently improves the sampled fronts; the stochastic nature automatically balances exploration and exploitation without manual tuning [21–23].

Expected hypervolume improvement represents the gold-standard approach by computing the expected increase in dominated volume while fully propagating predictive uncertainty. Although computationally intensive in high dimensions, recent approximations render it tractable for compositional spaces. For practical deployment the framework recommends either expected hypervolume improvement with suitable approximations for speed or Pareto-UCB for simplicity and interpretability. Both choices ensure that acquisition respects the multi-objective nature of materials design without collapsing trade-offs [14, 15, 17].

These acquisition strategies integrate seamlessly with the uncertainty quantification component described earlier. Epistemic and aleatoric estimates modulate the expected improvement calculations, directing the search toward compositions that simultaneously promise performance gains and knowledge expansion. The result is a principled decision engine capable of navigating the Pareto surface of high-dimensional compositional spaces [11, 12, 16].

Constraint Handling and Sequential Sampling

Constraint handling in compositional design distinguishes between hard and soft categories. Hard constraints include thermodynamic stability thresholds, synthesizability temperature limits, and exclusion of toxic elements; any candidate violating a hard constraint receives zero acquisition score and is immediately discarded. Soft constraints—such as cost ceilings or elemental availability preferences—are incorporated either as additional objectives to be minimized or as penalty terms within the acquisition score, allowing flexible trade-offs [13, 23].

The sequential sampling strategy proceeds in clearly defined steps that ensure systematic progress under budget limitations. The process begins with an initial space-filling design, typically a Latin hypercube sampling projected onto the composition simplex, to establish broad coverage with minimal evaluations. In each subsequent iteration the workflow trains surrogate models for every objective, computes epistemic and aleatoric uncertainty estimates, applies hard constraints as filters, folds soft constraints into the objective set or acquisition score, evaluates the chosen multi-objective acquisition function, and selects the top candidate (or batch of candidates) for evaluation. After obtaining the new

observation—whether from density functional theory or experiment—the surrogates are retrained and the loop repeats until the evaluation budget is exhausted or the Pareto frontier exhibits convergence [1, 6, 24, 26].

For parallel experimental or computational campaigns, batch sampling extends the single-candidate loop by selecting the top k candidates while imposing a diversity constraint that penalizes compositional similarity. This prevents redundant evaluations in nearby regions and maximizes collective information gain within each parallel round [3, 5, 8].

Throughout the process, constraint handling and acquisition scoring operate in tight coordination with uncertainty quantification. Hard constraints act as pre-filters before acquisition scoring, ensuring computational effort focuses exclusively on feasible candidates. Soft constraints influence ranking without eliminating options, preserving the ability to discover performance-cost trade-offs. The overall sequential strategy therefore maintains feasibility, respects budgets, and drives the search toward practically relevant regions of the Pareto surface [11, 12, 14, 17].

Relation to Existing Methods

Bayesian optimization forms the foundational paradigm upon which the proposed framework is constructed. Bayesian optimization provides a systematic way to select candidates by balancing predicted performance gains against uncertainty in surrogate models. This framework extends Bayesian optimization by embedding multi-objective Pareto handling, explicit constraint enforcement, and composition-specific adaptations that respect the simplex geometry and sparsity of alloy design spaces [4, 7, 8, 11, 12, 24].

Active learning studies focused on high-entropy alloys have revealed clear shortcomings of single-objective formulations. Research on high-entropy and multi-component systems has shown that conventional active learning often converges to local optima without exploring objective trade-offs. The present framework supplies a multi-objective solution that overcomes these limitations by continuously maintaining and expanding the Pareto frontier across all conflicting performance criteria [2, 6-8, 23, 26].

Sparsity within compositional manifolds demands acquisition strategies that prioritize broader exploration rather than purely local exploitation. Investigations of ternary alloys and high-dimensional multi-component systems have demonstrated that data scarcity in most regions leads to unreliable surrogate predictions unless acquisition functions explicitly account for distance from existing observations. This framework incorporates sparsity-aware exploration directly into the Pareto-based scoring mechanism, ensuring that candidates in underrepresented areas receive appropriate consideration [2-4].

Multi-fidelity strategies integrate naturally with the overall structure. Low-fidelity models serve as rapid screening tools to identify candidate regions, while high-fidelity evaluations refine the Pareto front for the most promising compositions. This hierarchical integration preserves the uncertainty-aware character of every selection step and extends the evaluation budget without compromising the multi-objective or constraint-handling components [29].

The framework also builds upon earlier multi-objective active learning developments originally demonstrated in molecular and reaction optimization contexts. Those approaches established the value of Pareto-aware scoring but operated in lower-dimensional or discrete spaces. By adapting these concepts to the continuous, constrained simplex of compositional design, the framework addresses the unique statistical structure of materials systems while retaining the core advantages of hypervolume-based improvement metrics [15-19, 22].

Constraint-handling techniques previously explored in shape memory alloys and architected materials provide additional conceptual grounding. Those studies emphasized the necessity of separating hard feasibility rules from soft preferences, yet typically applied them in single-objective settings. The current framework generalizes these ideas into a unified multi-objective loop, ensuring that constraint satisfaction remains an intrinsic part of every acquisition decision rather than an afterthought [13, 14, 23].

Overall, the proposed conceptual framework synthesizes and advances these related methods into a cohesive structure that is purpose-built for the high-dimensional, multi-objective, and uncertainty-heterogeneous nature of compositional materials discovery. It does not replace existing techniques but instead provides the integrative scaffold that allows practitioners to combine Bayesian optimization foundations, sparsity-aware exploration, multi-fidelity screening, and constraint logic within a single operational pipeline [1, 5, 20, 21, 25].

Design Principles

The framework rests on an explicit separation of epistemic and aleatoric uncertainty within the acquisition process. Epistemic uncertainty, which reflects reducible knowledge gaps, guides sampling toward genuinely informative regions, whereas aleatoric uncertainty—capturing irreducible noise inherent to the composition or evaluation method—prevents wasteful allocation of evaluations to inherently variable areas [9, 10, 25]. This distinction enables Pareto-based acquisition functions, such as expected hypervolume improvement or Pareto upper confidence bound, to evaluate candidates according to their potential to expand the non-dominated front across all objectives simultaneously, thereby preserving the richness of trade-offs that scalarization inevitably obscures [14-18].

A related mechanism incorporates constraints directly into selection: hard constraints are enforced through immediate filtering to ensure only thermodynamically or synthetically feasible compositions are considered, while soft constraints enter as additional objectives or penalty terms that shape ranking without prematurely discarding valuable compromises [13, 23]. Under these conditions, uncertainty-weighted scoring balances exploration and exploitation by modulating predicted improvements with calibrated epistemic estimates, naturally directing effort toward uncertain yet promising regions of the simplex while still capitalizing on well-characterized high performers [9, 11, 12]. When parallel evaluations are possible, batch sampling with diversity constraints further augments collective information gain by penalizing compositional redundancy and avoiding localized clustering [3, 5, 8].

Finally, retrospective validation of the acquisition strategy against historical datasets confirms its capacity to recover superior Pareto-optimal compositions within equivalent evaluation budgets, thereby establishing practitioner confidence before live deployment [1, 6, 24, 26]. Together these elements translate uncertainty quantification, multi-objective optimization, constraint handling, and sequential sampling into a coherent, implementable decision process that remains robust across surrogate models and material systems [2, 19-22].

Table 2 consolidates the manuscript’s design principles into an implementation-oriented matrix that clarifies what each principle governs, what error it prevents, and how it contributes to robust compositional discovery.

Table 2. Design-principle matrix for implementing uncertainty-aware multi-objective active learning in compositional materials discovery

Design principle	What it governs in practice	Failure mode prevented	Expected benefit for discovery campaigns	Primary manuscript component reinforced
Explicit separation of epistemic and aleatoric uncertainty	Distinguishes learnable knowledge gaps from irreducible noise before acquisition decisions are made	Wasteful sampling in intrinsically noisy regions and false confidence in sparse regions	More statistically meaningful candidate selection and stronger calibration across the composition landscape	Uncertainty quantification
Use of Pareto-based acquisition	Determines how multiple objectives are scored and	Hidden trade-offs, arbitrary weighting bias,	Clearer identification of non-dominated	Multi-objective acquisition

rather than scalarized optimization	compared during selection	and premature convergence to scalar optima	compositions and more transparent engineering compromise analysis	
Immediate filtering of hard constraints	Defines the boundary of admissible candidates before any ranking occurs	Selection of thermodynamically unstable, unsynthesizable, or toxic candidates	Prevents wasted evaluations and ensures the search remains anchored to feasible materials space	Constraint handling
Integration of soft constraints as objectives or penalties	Governs how cost, scarcity, manufacturability, or sustainability are incorporated without over-restricting the search	Overly idealized optimization that ignores real deployment priorities	Enables flexible negotiation between performance and practical engineering value	Constraint handling plus acquisition design
Uncertainty-weighted balancing of exploration and exploitation	Controls how the framework chooses between promising known regions and uncertain underexplored regions	Over-exploitation of local basins or undirected exploration with poor return on budget	More efficient frontier expansion under severe evaluation limits	Uncertainty quantification plus acquisition
Diversity-preserving batch selection for parallel campaigns	Governs candidate grouping when multiple evaluations are run simultaneously	Redundant batch composition choices concentrated in one region of the simplex	Greater collective information gain per iteration and better manifold coverage	Sequential sampling
Retrospective validation before live deployment	Tests whether the acquisition policy would have improved prior campaigns under the same budget	Uncritical deployment of a conceptually appealing but operationally weak selection strategy	Builds empirical confidence, improves trust, and reveals tuning weaknesses before real expenditure	Framework-wide operational governance

Implications for Materials Discovery

For practitioners engaged in alloy or multi-component system development, the framework recommends deploying expected hypervolume improvement as the default acquisition strategy whenever multiple performance criteria must be balanced simultaneously. This choice allows direct optimization of strength, conductivity, cost, and stability within a single campaign while automatically respecting thermodynamic and synthesizability constraints. Evaluation budgets of 10 to 100 calculations per iteration have proven sufficient in related compositional studies to reach practically useful regions of the Pareto frontier. Final candidate validation through targeted experiments remains essential to confirm that computational predictions translate to measurable improvements [11, 14, 15, 17, 23].

For model developers, the framework highlights the need to implement efficient approximations of expected hypervolume improvement that scale gracefully with the number of objectives and the dimensionality of the composition simplex. Composition-specific acquisition functions that incorporate the geometric constraints of the simplex offer another promising direction for future enhancement. Such developments will reduce computational overhead during each iteration and broaden accessibility for research groups without massive high-performance computing resources [4, 9, 10, 12, 29].

For benchmark designers, the framework calls for the creation of standardized multi-objective active learning test suites that incorporate realistic constraints, heteroscedastic noise models, and high-dimensional sparse composition spaces. These benchmarks should include both synthetic and experimentally grounded problems so that competing algorithms can be compared on criteria that mirror actual materials discovery workflows. The availability of such benchmarks will accelerate progress in uncertainty-aware methods and provide clear performance baselines for the community [2, 3, 16, 18–20].

Collectively, these implications position the framework as a practical guide that bridges conceptual advances with day-to-day discovery activities. Practitioners gain a structured decision process that reduces wasted evaluations, model developers obtain clear targets for algorithmic improvement, and benchmark designers receive a roadmap for creating more representative evaluation environments. The end result is a more efficient, transparent, and constraint-aware pathway from conceptual composition space to deployable high-performance materials [1, 6, 24, 26].

Open Questions

Extending the framework to hundreds of simultaneous objectives—spanning mechanical, thermal, electrical, magnetic, and environmental properties—demands new conceptual mechanisms, since Pareto-based methods lose both interpretability and tractability at high dimensionality. This scaling challenge is compounded by the need to incorporate qualitative or semi-quantitative goals such as synthesizability and processability, which resist direct numerical encoding; integrating expert judgment or auxiliary qualitative models into the Pareto acquisition loop while preserving uncertainty calibration thus remains a core conceptual tension. A related implication concerns the dynamic allocation of exploration effort across sparse and dense regions of composition space, where adaptive rules modulating the exploration–exploitation balance according to local data density could markedly improve efficiency, yet the precise triggering criteria and functional forms await rigorous definition. Beyond these immediate concerns lies the possibility of learning acquisition functions themselves through meta-learning across materials families, an approach that could yield policies more robust than hand-crafted ones, although the representation of compositional spaces and the construction of suitable meta-training tasks require careful theoretical development. Finally, propagating second-order epistemic uncertainty from constraints such as predicted thermodynamic stability through the acquisition and filtering stages—without inducing excessive conservatism or instability—emerges as a particularly subtle open problem. Addressing these interlocking issues will necessitate sustained collaboration among materials scientists, uncertainty quantification experts, and optimization theorists, thereby broadening the reach of uncertainty-aware active learning in complex compositional design [9, 10, 14, 15, 17].

Conclusion

Compositional design spaces require uncertainty-aware, multi-objective active learning. The central challenges include high dimensionality, multiple conflicting objectives, heteroscedastic uncertainty, practical constraints, the need for extrapolation, and severe limits on evaluation budgets. The proposed conceptual framework addresses these challenges through four integrated components: uncertainty quantification that separates epistemic from aleatoric contributions, multi-objective acquisition functions centered on Pareto frontier expansion, systematic constraint handling that distinguishes hard and soft requirements, and sequential sampling that proceeds from space-filling initialization to budget-aware iteration.

The design principles distilled from this structure—explicit separation of uncertainty types, adoption of Pareto-based scoring, enforcement of hard constraints via filtering, incorporation of soft constraints as objectives, uncertainty-weighted balancing of exploration and exploitation, diversity-preserving batch selection, and pre-deployment validation—provide a clear operational roadmap for implementation.

The framework shifts active learning from a heuristic search tool into a principled decision engine that respects the full complexity of modern materials engineering. It enables more reliable navigation of high-dimensional sparse spaces, clearer articulation of performance trade-offs, and tighter alignment between computational predictions and practical engineering requirements. Future progress will benefit from efficient approximations of hypervolume-based acquisition, standardized multi-objective benchmarks that incorporate realistic constraints, and continued integration of multi-fidelity and meta-learning concepts.

Ultimately, the conceptual framework presented here offers a unified foundation for accelerating the discovery of next-generation alloys, multi-component ceramics, and mixed ionic conductors. By embedding uncertainty awareness and multi-objective reasoning at every step, it promises to deliver materials that better satisfy simultaneous demands for performance, sustainability, and manufacturability within realistic development timelines.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 01 May 2025

Revised: 02 Aug 2025

Accepted: 26 Oct 2025

Published online: 18 January 2026

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

- Mo W, Xiao Y, Huang Y, Sun P, Li Y, Zheng X, et al. Active learning-based alloy design strategy for improving the strength-ductility balance of Al-Mg-Zn alloys. *Mater Des.* 2025;252:113772.
<https://doi.org/10.1016/j.matdes.2025.113772>.
- Deshmukh G, Wichrowski NJ, Evangelou N, Ghanekar PG, Deshpande S, Kevrekidis IG, et al. Active learning of ternary alloy structures and energies. *NPJ Comput Mater.* 2024;10(1):116.

<https://doi.org/10.1038/s41524-024-01256-z>.

Hu M, Tan Q, Knibbe R, Xu M, Liang G, Zhou J, et al. Designing unique and high-performance Al alloys via machine learning: Mitigating data bias through active learning. *Comput Mater Sci.* 2024;244:113204.

<https://doi.org/10.1016/j.commatsci.2024.113204>.

Sundar A, Tan X, Hu S, Gao MC. CALPHAD-based Bayesian optimization to accelerate alloy discovery for high-temperature applications. *J Mater Res.* 2025;40(1):112-22.

<https://doi.org/10.1557/s43578-024-01489-0>.

Cai J, Han M, Yan X, Chen Y, Li D, Zhao K, et al. A process-synergistic active learning framework for high-strength Al-Si alloys design. *NPJ Comput Mater.* 2025;11(1):228.

<https://doi.org/10.1038/s41524-025-01721-3>.

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

<https://doi.org/10.1098/rsta.2023.0242>.

Kurunczi-Papp D, Laurson L. Bayesian optimization of 7-component (AlVCrFeCoNiMo) single crystal alloy's compositional space to optimize elasto-plastic properties from molecular dynamics simulations. *Model Simul Mater Sci Eng.* 2024;32(8):085013.

<https://doi.org/10.1088/1361-651X/ad89b3>.

Torsti V, Mäkinen T, Bonfanti S, Koivisto J, Alava MJ. Improving the mechanical properties of Cantor-like alloys with Bayesian optimization. *APL Mach Learn.* 2024;2(1):016119.

<https://doi.org/10.1063/5.0179844>.

Zhang H, Chen W, Iyer A, Apley DW, Chen W. Uncertainty-aware mixed-variable machine learning for materials design. *Sci Rep.* 2022;12(1):19760.

<https://doi.org/10.1038/s41598-022-23431-2>.

Alvi SMAA, Mulukutla M, Flores N, Khatamsaz D, Janssen J, Perez D, et al. Accurate and uncertainty-aware multi-task prediction of HEA properties using prior-guided deep Gaussian processes. *NPJ Comput Mater.* 2025;11(1):306.

<https://doi.org/10.1038/s41524-025-01811-2>.

Wen D, Tucker V, Titus MS. Bayesian optimization acquisition functions for accelerated search of cluster expansion convex hull of multi-component alloys. *NPJ Comput Mater.* 2024;10(1):210.

<https://doi.org/10.1038/s41524-024-01391-7>.

Mints VA, Pedersen JK, Bagger A, Quinson J, Anker AS, Jensen KMØ, et al. Exploring the composition space of high-entropy alloy nanoparticles for the electrocatalytic H₂/CO oxidation with Bayesian optimization. *ACS Catal.* 2022;12(18):11263-71.

<https://doi.org/10.1021/acscatal.2c02563>.

Wang Y, Dang P, Xian Y, Zhou Y, Ding X, Sun J, et al. Uncertainty-aware multi-objective optimization for high work output and low hysteresis in TiNiCuHfCo shape memory alloys. *J Mater Sci Technol.* 2026;243:220-7.

<https://doi.org/10.1016/j.jmst.2025.03.095>.

Wei Q, Wang Y, Yang G, Li T, Yu S, Dong Z, et al. Discovering novel lead-free solder alloy by multi-objective Bayesian active learning with experimental uncertainty. *NPJ Comput Mater.* 2025;11(1):10.

<https://doi.org/10.1038/s41524-024-01480-7>.

Jablonka KM, Jothiappan GM, Wang S, Smit B, Yoo B. Bias free multiobjective active learning for materials design and discovery. *Nat Commun.* 2021;12(1):2312.

<https://doi.org/10.1038/s41467-021-22437-0>.

Kim M, Kim Y, Ha MY, Shin E, Kwak SJ, Park M, et al. Exploring optimal water splitting bifunctional alloy catalyst by Pareto active learning. *Adv Mater.* 2023;35(17):e2211497.
<https://doi.org/10.1002/adma.202211497>.

Park T, Kim E, Sun J, Kim M, Hong E, Min K. Rapid discovery of promising materials via active learning with multi-objective optimization. *Mater Today Commun.* 2023;37:107245.
<https://doi.org/10.1016/j.mtcomm.2023.107245>.

Du J, Guo J, Liu W, Li Z, Huang G, Zhang X. A universal strategy of multi-objective active learning to accelerate the discovery of organic electrode molecules. *Sci China Chem.* 2024;67(11):3681-7.
<https://doi.org/10.1007/s11426-024-2163-1>.

Peng B, Wei Y, Qin Y, Dai J, Li Y, Liu A, et al. Machine learning-enabled constrained multi-objective design of architected materials. *Nat Commun.* 2023;14(1):6630.
<https://doi.org/10.1038/s41467-023-42415-y>.

Antoniuk ER, Li P, Keilbart N, Weitzner S, Kaikhura B, Hiszpanski AM. Active learning enables generation of molecules that advance the known Pareto front. *NPJ Comput Mater.* 2026;12(1):81.
<https://doi.org/10.1038/s41524-025-01924-8>.

Wu X, Zhang F, Huang Y, Wang B, Xinye Y. Multiobjective design of bone scaffolds based on active learning. *J Mech Mater Struct.* 2026;21(1):29-48.
<https://doi.org/10.2140/jomms.2026.21.29>.

Torres JAG, Lau SH, Anchuri P, Stevens JM, Tabora JE, Li J, et al. A multi-objective active learning platform and web app for reaction optimization. *J Am Chem Soc.* 2022;144(43):19999-20007.
<https://doi.org/10.1021/jacs.2c08592>.

Wang Z, Chen X, Zhang D, Wu G, Ma Y, Zeng X, et al. Active learning design of bcc solid solution alloys with gigapascal strength and elemental metal-level ductility. *Proc Natl Acad Sci U S A.* 2026;123(5):e2530922123.
<https://doi.org/10.1073/pnas.2530922123>.

Sun Y, Zhou Y, Tao W, Wu J, Xiu Y, Ji H, et al. Inverse design of electrical conductivity in AlSi8 alloy using Bayesian optimization. *J Mater Sci.* 2026;61(14):9443-56.
<https://doi.org/10.1007/s10853-026-12394-8>.

Cao C, Kim HJ, Carbone MR, Reyes K, Yoo S, Lu F, et al. Uncertainty-aware machine learning for small-angle X-ray scattering analysis in autonomous experimentation. *Photon Sci.* 2026.
<https://doi.org/10.1021/photonsci.5c00053>.

Li Y, Jiang E, Hu K, Peng Y, Ni Z, Liu F, et al. Active learning-enabled the discovery of ultra-high saturation magnetization soft magnetic alloys. *Scr Mater.* 2025;257:116485.
<https://doi.org/10.1016/j.scriptamat.2024.116485>.

Chen J, Ou P, Chang Y, Zhang H, Li XY, Sargent EH, et al. Materials discovery using uncertainty-aware constrained Bayesian optimization with representation learning of high-dimensional inputs. *J Mech Des.* 2026;148(2):021707.
<https://doi.org/10.1115/1.4070206>.

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.
<https://doi.org/10.1007/s40192-020-00168-2>.

Aqib M, Ravikiran K, Li L, Prasad V. A multi-fidelity active learning framework for accelerated alloy design. *Mater Des.* 2026;263:115520.
<https://doi.org/10.1016/j.matdes.2026.115520>.