Larsson *et al. Journal of Artificial Intelligence for Materials Science* *(2026) 1:122*
https://doi.org/10.51847/b837828985

ORIGINAL RESEARCH **Open access**

Conceptual Foundations for Regret-Aware Materials AI Systems

Sven Larsson[1*], Erik Johansson[1], Anna Nilsson[2]

Abstract

In the rapidly advancing domain of artificial intelligence applied to materials science, systems are frequently called upon to make critical decisions under conditions of substantial uncertainty, such as selecting which candidate material to synthesize next, which experiment to prioritize for evaluation, or which property to measure in a given campaign. A fundamental aspect that current materials AI approaches largely ignore is the phenomenon of regret—the realization, after the fact, that a different choice would have produced a superior outcome, often carrying emotional, cognitive, and practical costs for the decision-maker. Regret theory, originating in decision theory and economics, provides a powerful alternative lens for understanding choice under uncertainty by incorporating not only expected utilities but also the anticipation and experience of post-decision disappointment or rejoicing. This paper proposes a conceptual framework for regret-aware materials AI systems that explicitly integrates regret quantification, theoretical regret bounds, regret minimization objectives, regret-aware acquisition functions, and regret communication mechanisms into the decision-making pipeline. The framework further delineates four primary types of regret encountered in materials contexts—synthesis regret, measurement regret, discovery regret, and resource regret—each arising from the irreversible, sequential, and high-stakes nature of experimental materials research. By embedding these elements, the proposed framework shifts materials AI from a narrow focus on reward maximization toward systems that more closely mirror the nuanced realities of scientific decision-making, where the avoidance of avoidable regret becomes a central design goal. Ultimately, embracing regret awareness promises more robust exploration of vast material spaces, better alignment between AI recommendations and laboratory constraints, and enhanced trust between human researchers and autonomous systems, thereby accelerating genuine discovery while mitigating the hidden costs of overlooked alternatives.

Keywords Bayesian optimization, Counterfactual reasoning, Regret theory, Regret-aware materials AI, Regret minimization, Regret-aware acquisition functions

*Correspondence:
Sven Larsson
sven.larsson@gmail.com
[1] Department of Computational Materials Analytics, Karolinska Institute, Stockholm, Sweden
[2] Department of AI Materials Engineering, Lund University, Lund, Sweden

Introduction

Materials AI systems have transformed the landscape of materials discovery by leveraging machine learning and optimization algorithms to guide the identification of novel compounds possessing targeted properties such as high conductivity, superior mechanical strength, or enhanced catalytic activity [1-6]. These systems routinely confront sequential decision-making tasks: determining which material candidate merits synthesis, which experimental protocol should be executed next, or which measurement technique will yield the most informative data given limited resources [7]. Each such decision carries tangible and often irreversible consequences because synthesis routes consume costly precursors and instrument time, experimental campaigns operate under strict budgetary and temporal constraints, and the outcomes directly influence subsequent research directions, publication prospects, and funding availability [8-15]. When an outcome proves suboptimal or when retrospective analysis reveals that a choice would have yielded superior results, both the human investigators and the supporting AI systems confront a form of regret. Despite its pervasive presence in real-world materials research, regret remains a largely unmodeled and unaddressed concept within contemporary materials AI architectures. Most existing frameworks prioritize metrics such as expected improvement or uncertainty reduction without explicitly accounting for the counterfactual nature of regret or its cumulative impact across a discovery campaign [8, 14, 16, 17].

This paper identifies the core problem as the systematic neglect of regret modeling in materials AI. Decisions in materials science are rarely cost-free; synthesizing a material via an ill-chosen route not only wastes immediate resources but may foreclose entire branches of the search space that could have proven fruitful [7, 18-23]. Early suboptimal selections create path dependencies that constrain later possibilities, amplifying the long-term effects of any single poor judgment [18, 24, 25]. High-cost evaluations, whether involving rare samples, specialized beamtime, or lengthy computational simulations, further elevate the stakes, as each erroneous prioritization delays breakthroughs and consumes finite laboratory capacity [9, 15]. Moreover, the broader scientific ecosystem—publication pressures, peer review, and competitive funding cycles—magnifies the practical and career-related consequences of directions that ultimately prove unproductive [22, 26]. Current Bayesian optimization and active learning approaches, while powerful, typically optimize for cumulative or instantaneous performance without considering how a sequence of choices might lead to avoidable regret [11, 13, 24].

The framework proposed in this article seeks to rectify this oversight by integrating regret awareness at the foundational level of materials AI design. By doing so, such systems can navigate the inherent trade-offs between exploration (to reduce future regret) and exploitation (to minimize immediate regret) in a manner that better reflects the practical realities of laboratory research [3, 12, 19]. The introduction of regret concepts promises to render materials AI more aligned with human decision-making processes, wherein the anticipation and minimization of regret often serve as powerful motivators and evaluative criteria [1, 2, 16]. This shift does not replace existing reward-based paradigms but augments them with a complementary lens that acknowledges the emotional and cognitive dimensions of scientific choice. The remainder of the paper proceeds systematically: first by tracing the theoretical origins of regret concepts, then by examining their specific manifestations in materials AI contexts, followed by the detailed articulation of a five-component regret-aware framework, and finally by elaborating each component with conceptual depth and nuance. Through this progression, the article establishes regret not as a peripheral psychological afterthought but as a central, actionable principle capable of elevating the sophistication and trustworthiness of materials AI systems.

Regret Theory: Origins and Concepts

Regret theory emerged in the late twentieth century as a psychologically grounded alternative to classical expected utility theory in the study of decision-making under uncertainty. Kochenderfer *et al.* [1] provided one of the earliest formal treatments, demonstrating that rational agents often evaluate prospects not solely based on probabilistic outcomes but also on the basis of the regret they would experience if a forgone alternative later proved superior. This insight challenged the assumption of purely forward-looking utility maximization by showing that anticipated regret can systematically alter choice behavior, leading decision-makers to prefer options that reduce the potential for post-decision disappointment even when those options carry lower expected value. Wang *et al.* [2] extended this line of reasoning into a comprehensive alternative theory of rational choice, arguing that the utility of any outcome is modified by the regret or rejoicing associated with the difference between the chosen action and the best possible action under the realized state of the world. Their model treats regret as a real and measurable component of decision utility, thereby offering a descriptively richer account of human and organizational behavior under uncertainty.

In the context of materials AI, regret is defined as the emotional and cognitive experience arising when a decision-maker (human or algorithmic) realizes that a different choice would have led to a superior outcome, measured either as the difference between the achieved result and the best attainable result given the same information or as the accumulated shortfall across a sequence of decisions.

Central to regret theory are the distinctions between simple regret and cumulative regret. Simple regret concerns the regret associated with the final recommendation after a fixed budget of evaluations has been exhausted; it focuses on the gap between the best possible material or property value discoverable within the campaign and the value actually recommended by the system [18, 25]. Cumulative regret, by contrast, aggregates the instantaneous regrets incurred at each step of the sequential decision process, capturing the total “cost” of suboptimal choices made along the path [8, 13]. Counterfactual reasoning lies at the heart of both forms, requiring the decision-maker to compare the actual observed outcome against hypothetical outcomes that would have arisen had an alternative action been selected [16, 22]. These concepts have been further refined in machine learning and optimization literature, where regret minimization has become a cornerstone of algorithms operating in bandit and Bayesian settings [10-12].

In materials AI, the theoretical apparatus of regret theory gains additional salience because decisions are not abstract but tied to concrete laboratory realities. The anticipation of synthesis regret, for instance, may lead a system to avoid overly speculative candidates even when their expected performance appears promising, precisely because the cost of realizing a poor outcome is high [20, 23]. Similarly, the framework of counterfactual regret minimization encourages algorithms to maintain an internal model of alternative decision paths, enabling more informed trade-offs between immediate gains and long-term robustness [27-29]. By grounding materials AI in these foundational ideas, the field can move beyond purely statistical performance criteria toward a decision-theoretic paradigm that explicitly acknowledges the human and practical dimensions of scientific choice. This section has thus laid the conceptual groundwork for applying regret concepts to the specific challenges of materials discovery, where uncertainty, irreversibility, and high stakes intersect in ways that traditional reward-centric models fail to capture.

Regret in Materials AI Contexts

Regret acquires particular importance in materials AI because the decisions involved are characterized by irreversibility, sequential dependencies, high financial and temporal costs, pronounced opportunity costs, and significant downstream consequences for scientific careers and institutional priorities. First, synthesis decisions are often irreversible in practice; once a material is fabricated using a chosen precursor or processing route, reverting to an alternative pathway typically requires restarting the entire campaign, consuming additional resources and time that cannot be recovered [20]. Second, the sequential nature of discovery campaigns means that early choices constrain the feasible region for later experiments, creating path dependencies that can lock the system into suboptimal trajectories if an initial selection proves regrettable [7, 17]. Third, experimental evaluations are inherently expensive, whether measured in terms of specialized equipment access, rare elemental samples, or lengthy characterization protocols, so selecting the wrong candidate imposes immediate and measurable economic penalties [15]. Fourth, opportunity costs are acute because laboratory capacity, computational budgets, and researcher attention are finite; devoting resources to one candidate necessarily precludes exploring another that might later emerge as superior [25]. Fifth, the broader ecosystem of scientific publication, peer review, and funding exerts pressure such that a direction that yields disappointing results can jeopardize grant renewals, publication acceptance, or career advancement, thereby layering institutional and personal stakes onto algorithmic choices [22, 26].

Consider, for example, a Bayesian optimization campaign aimed at discovering a new perovskite composition for photovoltaic applications. If the system selects a candidate whose synthesis fails due to phase instability, the resulting synthesis regret manifests not only as wasted precursors but as lost weeks of laboratory time that could have been allocated to a more stable alternative later identified through counterfactual analysis [19, 23]. Similarly, measurement regret arises when an inappropriate characterization technique is chosen—perhaps Raman spectroscopy instead of X-ray diffraction—yielding data that proves insufficiently informative and forcing costly repeat experiments [20]. Discovery regret occurs when a promising lead is deprioritized early in the campaign, only to be recognized as valuable after competing research groups publish superior results on the same composition [16, 27]. Resource regret, meanwhile, surfaces when beamtime at a synchrotron facility or high-performance computing cycles are allocated to low-value candidates, crowding out higher-potential explorations and generating cumulative shortfalls that compound over the project lifetime [21, 24].

These examples illustrate that regret in materials contexts is not merely a psychological construct but a quantifiable dimension of decision quality with direct implications for efficiency and innovation. Current materials AI pipelines, by focusing predominantly on expected improvement or upper confidence bounds, implicitly treat all suboptimal outcomes as equivalent in terms of their regret implications, thereby overlooking the differential costs associated with different failure modes [9, 14]. Integrating regret awareness addresses this limitation by providing a language and set of tools for explicitly modeling the asymmetries inherent in materials decisions. The following sections build upon this foundation by proposing a comprehensive framework and elaborating its components in detail.

A Framework for Regret-Aware Materials AI

This article proposes a conceptual framework for regret-aware materials AI systems organized around five interconnected components that together enable explicit consideration of regret throughout the decision-making lifecycle.

Figure 1 synthesizes the manuscript’s central argument by showing how materials AI decision contexts generate distinct regret forms that are translated into five design components and, ultimately, into more robust and trustworthy scientific decision support.

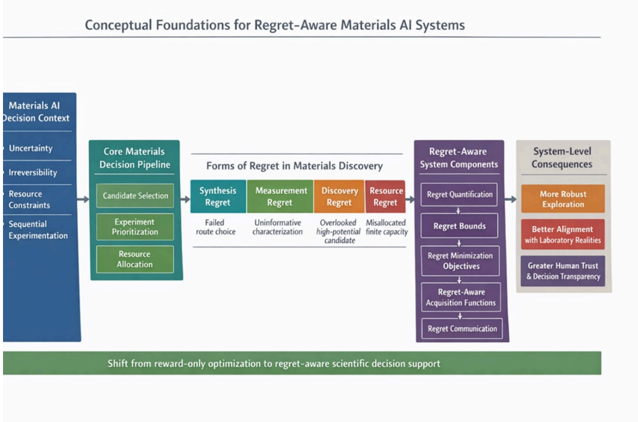

Figure 1. Conceptual foundations for regret-aware materials AI systems

The framework shifts the optimization objective from pure reward maximization to a balanced pursuit of performance gains alongside explicit minimization of avoidable regret, thereby aligning algorithmic behavior more closely with the practical and psychological realities of materials research [3, 18, 28].

The five components are: (1) regret quantification, which establishes rigorous metrics for measuring both simple and cumulative regret in materials-specific terms; (2) regret bounds, which provide theoretical guarantees on the maximum regret achievable under given search-space characteristics, evaluation costs, and measurement noise; (3) regret minimization objectives, which reformulate the optimization problem to trade off immediate performance against the reduction of future regret; (4) regret-aware acquisition functions, which modify standard Bayesian optimization heuristics to select experiments that explicitly account for the potential regret of choices; and (5) regret communication, which translates internal regret estimates and bounds into interpretable reports for human decision-makers.

Conceptually, these components form a closed-loop architecture. Regret quantification supplies the foundational metrics that feed into regret bounds, which in turn inform the setting of regret minimization objectives. These objectives guide the design of regret-aware acquisition functions, whose outputs are executed in the laboratory or simulation environment. Finally, regret communication closes the loop by conveying the realized and anticipated regret back to both the algorithm and the human user, enabling adaptive recalibration of priorities and trust calibration. A conceptual figure of the framework would depict these five components arranged in a pentagonal cycle with bidirectional arrows indicating information flow: quantification at the base feeding upward into bounds and minimization, acquisition functions positioned at the right-hand side driving experimental selection, and communication at the apex returning feedback that refines all prior stages. Dashed lines would connect each component to a central “materials decision core” representing the integration with existing autonomous laboratory infrastructure, emphasizing that regret awareness augments rather than replaces current pipelines [3, 19, 24].

By structuring materials AI around these elements, the framework ensures that regret is not an after-the-fact diagnostic but an active design principle embedded from the outset. It accommodates the unique features of materials search spaces—high dimensionality, expensive black-box evaluations, and noisy measurements—while preserving compatibility with established techniques such as Gaussian process surrogates and Bayesian optimization [9, 14, 21]. The subsequent section elaborates on each component with conceptual detail, highlighting its interdependencies and materials-specific instantiations.

Components of the Framework

Component 1: Regret quantification

This component establishes precise, context-sensitive metrics for measuring regret within materials AI campaigns. Simple regret is quantified as the difference between the best achievable performance (for example, the highest recorded photovoltaic efficiency or mechanical strength within the evaluated set) and the performance of the final recommended candidate after a fixed evaluation budget [18, 25]. Cumulative regret aggregates the instantaneous shortfalls at each decision step, thereby capturing the total opportunity cost incurred across the entire sequential process [8, 13]. In materials contexts, quantification further incorporates domain-specific costs such as synthesis failure probability, measurement noise variance, and resource expenditure, ensuring that regret scores reflect not only property deviations but also laboratory realities [20, 23].

Component 2: Regret bounds

Regret bounds supply theoretical upper limits on the regret that any algorithm can incur given the size of the candidate space, the cost of each evaluation, and the level of observation noise [11, 21]. Materials-specific factors—such as the combinatorial explosion of possible compositions, the heteroscedastic noise inherent in experimental measurements, and the non-stationarity of synthesis outcomes—necessitate tailored bound derivations that go beyond generic bandit analyses [10, 27]. These bounds serve both as performance guarantees and as diagnostic tools, allowing practitioners to assess whether observed regret exceeds theoretically achievable minima and to diagnose whether the surrogate model or acquisition strategy requires refinement [29].

Component 3: Regret minimization objectives

Rather than optimizing solely for expected reward, this component reformulates the objective function to penalize actions likely to produce high future regret while still permitting controlled exploration [12, 17]. The trade-off is explicit: aggressive exploitation may reduce immediate regret but leave the system vulnerable to discovery regret later in the campaign, whereas excessive exploration may inflate cumulative regret in the short term [8, 24]. Materials AI implementations, therefore, incorporate dynamic weighting that adapts to campaign stage, remaining budget, and observed noise levels, ensuring that minimization objectives remain responsive to evolving laboratory conditions [3, 23].

Component 4: Regret-aware acquisition functions

Standard acquisition functions such as expected improvement or upper confidence bound are augmented to penalize selections whose counterfactual regret—i.e., the regret that would arise if an alternative, currently less promising candidate later proves superior—is high [19, 24]. This modification encourages the system to balance the immediate expected gain against the risk of locking the campaign into a regrettable trajectory, thereby producing acquisition decisions that are both performance-oriented and regret-conscious [13, 28].

Component 5: Regret communication

This final component addresses the human-AI interface by translating internal regret estimates, bounds, and trade-off analyses into accessible visualizations and textual reports [3, 16, 22]. Effective communication might include regret heatmaps overlaid on the materials composition space, confidence intervals around simple-regret projections, or narrative summaries that highlight the most regret-prone decisions made thus far. By making regret transparent, the component fosters calibrated trust, enables human override when domain knowledge suggests alternative priorities, and supports post-campaign retrospectives that improve future algorithm design [29].

Each component is designed to operate interdependently, ensuring that regret awareness permeates every layer of the materials AI pipeline without introducing prohibitive computational overhead. The framework thus offers a comprehensive yet modular blueprint for transitioning from regret-ignorant to regret-aware systems.

Types of Regret in Materials Decisions

Regret in materials AI manifests in four distinct yet interrelated types that arise directly from the irreversible, sequential, and resource-constrained nature of discovery campaigns [3, 20, 23]. Each type carries unique implications for quantification and minimization, requiring tailored strategies within the regret-aware framework.

Table 1 consolidates the four regret types into an operational design schema by linking each type to its decision locus, counterfactual logic, measurable form, and most appropriate system response.

Table 1. Analytical mapping of regret types to decision points, metrics, and design responses in materials AI

Regret type	Typical decision point	Counterfactual comparison	Prir reg expre
Synthesis regret	Selecting a material candidate, precursor set, or synthesis route	Chosen synthesis path versus a feasible alternative route or candidate that would have produced a better material outcome	Reali tha selecte faile yieldec prop while a route likely succ
Measurement regret	Choosing which property or characterization technique to measure next	Chosen measurement versus the alternative measurement that would have reduced uncertainty more effectively	Reali tha collect w insuffi informa downs deci
Discovery regret	Prioritizing or discarding candidate materials during the search	Chosen search trajectory versus neglected candidate that later proves superior	Reali tha depric candid the valu discove
Resource regret	Allocating beamtime, compute, budgets, or researcher attention	Resources assigned to low-value candidates versus higher-value opportunities left unexplored	Reali that capac consu low- acti
Cross-type implication	Multi-stage campaign governance	The entire realized campaign versus the strongest feasible alternative decision sequence	Reco that comp acr interde deci

Synthesis regret

Synthesis regret captures the moment at which a chosen experimental pathway reveals itself as suboptimal relative to unrealized alternatives. This form of regret emerges when a selected synthesis route or precursor configuration yields materials with degraded properties—or fails entirely—despite the existence of viable pathways that would have produced superior outcomes [20, 23]. In materials practice, the contrast between a high-temperature solid-state route that introduces phase impurities and a sol–gel approach capable of stabilizing a pure perovskite phase exemplifies this dynamic. What is at stake is not simply experimental inefficiency, but a misalignment between decision criteria and latent process–structure relationships that only become visible after execution. Quantification, therefore, requires a counterfactual framing in which the realized figure of merit—such as bandgap or conductivity—is evaluated against the best attainable outcome within the same chemical space, adjusted for synthesis cost and failure likelihood [19]. This formulation foregrounds the asymmetry between observed and forgone trajectories, making regret sensitive to both performance and feasibility. Mitigation is followed by embedding regret-aware acquisition functions that incorporate predicted synthesis failure risks, often derived from counterfactual simulation, thereby discouraging pathways that appear promising in isolation but are fragile under realistic process conditions [24].

Measurement regret

A parallel but distinct challenge arises at the level of characterization, where measurement regret reflects the consequences of acquiring information that proves insufficient for downstream inference. Here, the issue is not that data are absent, but that the chosen measurement modality fails to resolve the properties most critical for decision-making [20]. For instance, prioritizing Raman spectroscopy in the analysis of a catalytic system may yield structural signatures, yet omit oxidation-state information that would have been accessible through X-ray absorption spectroscopy and ultimately decisive for mechanistic interpretation [16]. The resulting limitation becomes apparent only when subsequent decisions demand information that was never captured. Conceptually, this form of regret can be expressed through an information-theoretic lens, as the gap between the utility of observed data and the maximum attainable utility across alternative measurement strategies, operationalized as a reduction in posterior uncertainty [21]. Such a framing emphasizes that measurement choices shape not only what is known, but what can be known. Addressing this constraint requires integrating measurement-cost-adjusted regret bounds into experimental design, favoring strategies that either combine complementary modalities or sequence measurements in ways that preserve optionality for future inference [28].

Discovery regret

Discovery regret emerges more gradually, often becoming visible only in retrospect as the consequences of early prioritization decisions unfold. It arises when a candidate initially deemed unpromising is excluded from further exploration, only for later evidence—whether internal or external—to demonstrate its superiority [16, 27]. In materials discovery campaigns, this pattern is frequently observed when compositions dismissed during early screening phases later achieve breakthrough performance, as in the case of lead halide systems whose long-term stability was initially underestimated. Unlike synthesis or measurement regret, which are tied to specific experimental choices, discovery regret reflects the structure of the search process itself and its susceptibility to premature convergence. Its quantification is typically grounded in simple regret, defined as the performance gap between the final recommended candidate and the best material identified either during or after the campaign [18, 25]. This metric captures the cost of missed opportunities rather than inefficient execution. Mitigation, accordingly, depends on maintaining a calibrated balance between exploration and exploitation, with explicit exploration terms embedded in optimization objectives to ensure that high-uncertainty regions of the materials space remain accessible despite short-term performance trade-offs [12, 17].

Resource regret

At a broader scale, resource regret reflects the cumulative consequences of how finite experimental and computational capacities are allocated over the course of a discovery campaign. It arises when time, funding, or infrastructure—such as synchrotron beamtime or large-scale simulations—are disproportionately committed to candidates that ultimately yield limited value. At the same time, more promising directions remain underexplored [21, 24]. The significance of this form of regret lies in its compounding nature: each allocation decision constrains subsequent options, gradually shaping the trajectory of the entire research program. In practical terms, dedicating scarce experimental slots to marginal alloy systems at the expense of high-potential oxide materials exemplifies how opportunity costs materialize in real time. Quantification requires aggregating resource-weighted regret across all decisions, integrating both temporal and monetary expenditures into a unified measure of lost potential [8, 13]. This perspective shifts attention from isolated outcomes to the efficiency of the overall allocation strategy. Minimization can then be achieved through dynamic, budget-aware regret bounds that continuously re-evaluate remaining resources and redirect them toward candidates offering the greatest expected reduction in regret per unit cost, thereby preserving flexibility even in late-stage exploration [11, 29].

Taken together, these forms of regret reveal that decision-making in materials AI is inherently shaped by asymmetric and path-dependent consequences. Reward maximization alone fails to capture these dynamics, as it privileges immediate gains without accounting for unrealized alternatives or foreclosed opportunities. Incorporating regret as a central design consideration thus reorients the system toward a more reflective mode of optimization, one that remains sensitive to both what is achieved and what is inadvertently left behind [3, 19].

Relation to Existing Concepts

Regret is conceptually distinct from several related notions frequently invoked in materials informatics, yet it complements them by adding a counterfactual and ex-post evaluative dimension [1, 2, 16].

Table 2 clarifies why regret should be treated as a distinct analytical construct in materials AI rather than as a synonym for opportunity cost, loss, risk, or uncertainty.

Table 2. Distinguishing regret from adjacent decision concepts in materials AI

Concept	Core question answered	Temporal orientation	Comparis basis
Regret	Was another unchosen option better after the outcome became known?	Ex post, with counterfactual reflection	Chose action versus th best forg alternati under th realize state
Opportunity cost	What benefit was forgone by not choosing the next-best alternative?	Usually ex ante or static comparative	Chose option versus ne best availabl option
Loss function	How wrong was the prediction or decision relative to the truth?	Typically model-evaluation oriented	Predictio action versus observe outcom
Risk	How likely are undesirable outcomes before acting?	Ex ante	Probabil distributi of possib advers outcom
Uncertainty quantification	How uncertain is the model about what it knows?	Primarily ex ante	Estimat confiden or varian around predictio
Expected reward/expected improvement	Which action appears most beneficial on average?	Ex ante optimization	Expecte value un the curre model

Opportunity cost measures the forgone benefit of the next-best alternative in purely economic terms. In contrast, regret additionally incorporates the emotional and cognitive experience of that loss, making it a richer descriptor of decision quality under uncertainty [1, 25]. A loss function quantifies the factual discrepancy between predicted and observed outcomes; regret, by contrast, is inherently counterfactual, evaluating what would have occurred under an unchosen action [18]. Risk concerns the ex-ante probability distribution of adverse outcomes before a decision is made. At the same time, regret evaluates outcomes ex-post, after uncertainty has resolved and alternatives can be directly compared [1, 22]. Finally, even perfect uncertainty quantification does not eliminate regret, because regret can arise from the realization that a different choice—despite correctly modeled probabilities—would have been superior in the realized world [16, 29]. By distinguishing itself along these axes, regret theory enriches rather than replaces existing frameworks, providing a unified language for evaluating both algorithmic and human scientific choices [2, 3].

Implications for Materials AI Practice

Adoption of the regret-aware framework implies concrete changes across the materials AI ecosystem. For authors, manuscripts should routinely report both simple and cumulative regret bounds alongside performance metrics, explicitly discuss trade-offs between immediate reward and future regret reduction, and present counterfactual analyses of key decision points [3, 11, 28]. Reviewers, in turn, should query whether proposed algorithms incorporate regret minimization objectives or merely optimize expected improvement, demanding evidence that regret types relevant to the specific materials domain have been considered [19, 22]. For the broader community, the framework calls for the establishment of standardized regret benchmarks tailored to synthesis, measurement, and resource constraints, the development of open-source regret-communication toolkits, and empirical studies of how human researchers experience and respond to different regret types in autonomous laboratories [9, 15, 29]. Collectively, these shifts will foster systems that not only accelerate discovery but also earn greater trust by transparently acknowledging and mitigating the human cost of suboptimal decisions.

Conclusion

The regret-aware materials AI framework articulated herein integrates regret quantification, bounds, minimization objectives, acquisition functions, and communication into a cohesive conceptual architecture that explicitly acknowledges the irreversible and high-stakes nature of materials decisions. By distinguishing simple from cumulative regret, delineating four domain-specific regret types, and clarifying regret's unique relation to opportunity cost, loss, risk, and uncertainty, the framework offers a principled path beyond reward-centric optimization. Materials AI systems designed under this paradigm will better mirror the nuanced realities of laboratory science, minimizing avoidable regret while preserving exploratory power. The field is therefore urged to transition from regret-ignorant to regret-aware architectures, thereby aligning algorithmic intelligence more closely with the practical, cognitive, and emotional dimensions of scientific discovery.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 27 May 2023 Revised: 02 Jul 2023 Accepted: 07 Oct 2023

Rights and permissions

References

Kochenderfer MJ, Wheeler TA, Wray KH. Algorithms for decision making. Cambridge (MA): MIT Press; 2022.

Wang J, Ma X, Xu Z, Zhan J. Regret theory-based three-way decision model in hesitant fuzzy environments and its application to medical decision. IEEE Trans Fuzzy Syst. 2022;30(12):5361-75.

Zhang N, Zheng S, Tian L, Wei G. Study the supplier evaluation and selection in supply chain disruption risk based on regret theory and VIKOR method. Kybernetes. 2024;53(10):3848-74.

Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A. Machine learning for molecular and materials science. Nature. 2018;559(7715):547-55.

Schmidt J, Marques MR, Botti S, Marques MA. Recent advances and applications of machine learning in solid-state materials science. NPJ Comput Mater. 2019;5(1):83.

Zunger A. Inverse design in search of materials with target functionalities. Nat Rev Chem. 2018;2(4):0121.

Montoya JH, Aykol M, Anapolsky A, Gopal CB, Herring PK, Hummelshøj JS, et al. Toward autonomous materials research: Recent progress and future challenges. Appl Phys Rev. 2022;9(1).

Nguyen V, Gupta S, Rana S, Li C, Venkatesh S. Regret for expected improvement over the best-observed value and stopping condition. In: Asian Conference on Machine Learning. Seoul: PMLR; 2017. p. 279-94.

Liang Q, Gongora AE, Ren Z, Tiihonen A, Liu Z, Sun S, et al. Benchmarking the performance of Bayesian optimization across multiple experimental materials science domains. NPJ Comput Mater. 2021;7(1):188.

Iwazaki S, Takeno S, Tanabe T, Irie M. Failure-aware Gaussian process optimization with regret bounds. Adv Neural Inf Process Syst. 2023;36:24388-400.

Scarlett J. Tight regret bounds for Bayesian optimization in one dimension. In: International Conference on Machine Learning. Stockholm: PMLR; 2018. p. 4500-8.

Sim RH, Zhang Y, Low BK, Jaillet P. Collaborative Bayesian optimization with fair regret. In: International Conference on Machine Learning. PMLR; 2021. p. 9691-701.

Krishnamurthy SK, Zhan R, Athey S, Brunskill E. Proportional response: contextual bandits for simple and cumulative regret minimization. Adv Neural Inf Process Syst. 2023;36:30255-66.

Frazier PI. A tutorial on Bayesian optimization. arXiv [Preprint]. 2018:arXiv:1807.02811.

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.

Xu R, Yu Y, Zhang C, Ali MK, Ho JC, Yang C. Counterfactual and factual reasoning over hypergraphs for interpretable clinical predictions on EHR. In: Machine Learning for Health. New Orleans, LA: PMLR; 2022. p. 259-78.

Hur T, Allenby GM. A choice model of utility maximization and regret minimization. J Mark Res. 2022;59(6):1235-51.

Raschka S, Mirjalili V. Python machine learning: Machine learning and deep learning with Python, scikit-learn, and TensorFlow 2. Birmingham: Packt Publishing Ltd; 2019.

Deng Z, Zheng Y, Zhang J, Liu P, Zhu Z. Eco-environmental regret-aware optimization of networked multi-energy microgrids with fully carbon elimination and electric vehicles' promotion. Sustain Cities Soc. 2024;115:105807.

Ommer J, Kalas M, Neumann J, Blackburn S, Cloke HL. Turning regret into future disaster preparedness with no-regrets. EGUsphere. 2024;2024:1-5.

Azar MG, Osband I, Munos R. Minimax regret bounds for reinforcement learning. In: International Conference on Machine Learning. Sydney: PMLR; 2017. p. 263-72.

Wallis CJ, Zhao Z, Huang LC, Penson DF, Koyama T, Kaplan SH, et al. Association of treatment modality, functional outcomes, and baseline characteristics with treatment-related regret among men with localized prostate cancer. JAMA Oncol. 2022;8(1):50-9.

Braverman M, Mao J, Schneider J, Weinberg M. Selling to a no-regret buyer. In: Proceedings of the 2018 ACM Conference on Economics and Computation. New Orleans, LA: ACM; 2018. p. 523-38.

Wang X, Jin Y, Schmitt S, Olhofer M. Recent advances in Bayesian optimization. ACM Comput Surv. 2023;55(13s):1-36.

Farina G, Kroer C, Sandholm T. Regret circuits: composability of regret minimizers. In: International Conference on Machine Learning. Long Beach, CA: PMLR; 2019. p. 1863-72.

Mishra AK, Tsionas MG. A minimax regret approach to decision making under uncertainty. J Agric Econ. 2020;71(3):698-718.

Chehade M, Mccarthy MM, Squires A. Patient-related decisional regret: an evolutionary concept analysis. J Clin Nurs. 2024;33(11):4484-503.

Asas J, Hawkins M. Does spirituality affect your amount of regret? Available from: https://opus.govst.edu/research_day/2022/sessions/7/.

Brown N, Lerer A, Gross S, Sandholm T. Deep counterfactual regret minimization. In: International Conference on Machine Learning. Long Beach, CA: PMLR; 2019. p. 793-802.