ORIGINAL RESEARCH Open access

The Problem of Scientific Friction: When AI Smooths Away Important Conceptual Difficulties

Emily Johnson[1*], Robert Smith[1], Laura Brown[2], Kevin Miller[1]

Abstract

In the rapidly advancing field of artificial intelligence for materials science, systems are routinely engineered to eliminate every form of resistance—tedious calculations, anomalous data points, conceptual dead-ends, and even the slow, uncomfortable process of reconciling contradictory physical intuitions. Yet not all resistance is wasteful. Scientific friction—the productive resistance, conceptual struggle, and material recalcitrance that forces deeper engagement—functions as an overlooked epistemic resource that sharpens understanding, exposes hidden assumptions, and drives genuine discovery. This failure mode analysis articulates scientific friction as the generative difficulty that AI increasingly erases through automation, smoothing, abstraction, and black-boxing. It distinguishes four epistemically valuable types of friction (conceptual, methodological, material, and social) that have historically propelled progress in materials design, inverse problems, and autonomous discovery pipelines. The paper then identifies four mechanisms by which contemporary AI systems remove these frictions and presents a typology of four resulting failure modes—conceptual smoothing, anomaly suppression, methodological opacity, and premature closure—that undermine the very understanding AI claims to accelerate. Drawing exclusively on the peer-reviewed literature in artificial intelligence for materials science, science and technology studies, and philosophy of science, the analysis proposes detection principles and mitigation strategies that allow researchers to preserve productive friction without sacrificing efficiency. By reframing friction not as a bug to be fixed but as a feature to be curated, this work offers a framework for designing friction-aware AI tools that sustain the epistemic depth required for transformative materials innovation rather than delivering only superficial, frictionless outputs.

Keywords Scientific friction, Productive difficulty, Epistemic engagement, AI-driven materials science, Conceptual struggle, Automation in discovery

*Correspondence:
Emily Johnson
emily.johnson@gmail.com
[1] Department of Materials Informatics and AI, University of Toronto, Toronto, Canada
[2] Department of Smart Materials Analytics, McGill University, Montreal, Canada

Introduction

Artificial intelligence systems deployed in materials science are explicitly designed to smooth away difficulty. They automate the generation of candidate structures, predict properties with startling speed, and close the loop between simulation and experiment in ways that once required months of painstaking human labor. Yet the very success of these tools raises a deeper epistemic concern. When every obstacle is removed, something essential to scientific understanding may also disappear. Star argued that ill-structured problems require friction, but contemporary AI architectures are optimized to eliminate precisely that resistance. Kuhn showed that resistance to anomalies drives scientific progress. Yet, machine learning pipelines in solid-state materials science are engineered to minimize such resistance through data augmentation and loss-function regularization [1-5]. The problem is not that AI accelerates discovery; the problem is that it accelerates discovery by systematically erasing the very forms of difficulty that have historically produced robust conceptual insight.

This paper analyzes the problem of scientific friction in artificial intelligence for materials science. Scientific friction is defined here as the productive resistance, conceptual struggle, and material recalcitrance that generate understanding rather than merely impede progress. The analysis begins by establishing a precise definition that distinguishes productive friction from mere inefficiency. It then delineates four types of friction that are epistemically valuable in materials research. Next, it examines the four primary mechanisms by which AI systems remove friction. A typology of four specific failure modes follows, each illustrated through materials-relevant scenarios drawn from the literature. The discussion stops at this point to allow evaluation before proceeding. Throughout, the argument remains conceptual and analytic, identifying patterns rather than claiming empirical outcomes [3].

The motivation for this analysis is urgent. As Montoya and colleagues note in their review of autonomous materials research, the field is moving toward fully closed-loop systems in which human conceptual engagement is minimized [6, 7]. Zunger's foundational work on inverse design already envisioned workflows in which target functionalities dictate material structures with minimal human intervention [6]. While such acceleration is undeniably powerful, the literature increasingly reveals unintended epistemic costs. Recent contributions in machine learning for materials highlight how automated pipelines can bypass the very conceptual difficulties that once forced researchers to refine their physical intuitions. Similarly, studies of AI-driven materiality document cases in which generative models produce plausible but conceptually shallow outputs because the underlying friction of reconciling model predictions with thermodynamic constraints has been engineered away [8-18].

The consequences extend beyond any single experiment. When friction is removed at scale, entire research programs risk shifting from deep understanding to pattern matching. Butler and co-authors, while celebrating machine learning's successes in molecular and materials science, implicitly acknowledge that the speed of prediction can outpace the development of explanatory frameworks [4]. Schmidt *et al.* similarly document how recent advances in solid-state applications rely on ever-larger datasets that smooth over the messy, resistant details where true insight often resides [5]. The present work, therefore, treats scientific friction not as an obstacle to be overcome but as an epistemic resource whose systematic removal constitutes a distinct failure mode in AI-assisted materials discovery. By mapping this failure mode, the paper provides the conceptual foundation required for future friction-aware design.

Defining Scientific Friction

Scientific friction must be defined with precision if it is to serve as an analytic category rather than a vague complaint about slow progress. Definition 1: Scientific friction is the productive resistance, difficulty, or conceptual struggle that generates understanding, reveals hidden assumptions, or forces creative problem-solving within the practice of scientific inquiry. This definition emphasizes productivity: the resistance must yield epistemic gain rather than mere delay. Star's foundational analysis of ill-structured problems already hinted at this distinction when she described boundary objects as sites where friction enables heterogeneous collaboration rather than paralysis [1]. Kuhn likewise located the engine of scientific revolutions in the productive resistance that anomalies create against prevailing paradigms [2].

It is essential to distinguish scientific friction from three related but distinct concepts. First, scientific friction is not inefficiency. Inefficiency is unproductive difficulty—redundant computation, poorly designed interfaces, or bureaucratic delays that consume resources without advancing understanding. Friction, by contrast, is the very difficulty that advances understanding. Second, scientific friction is not an obstruction. Obstructions are external barriers—funding shortfalls, equipment failures, or institutional gatekeeping—that block progress altogether. Friction operates inside the research process itself, shaping the trajectory of thought. Third, scientific friction is not noise. Noise is random interference that obscures the signal. Friction is structured resistance that sharpens the signal by compelling the researcher to confront and resolve underlying tensions.

A conceptual diagram of the friction spectrum helps clarify these distinctions. Imagine a horizontal axis running from unproductive friction (left) to productive friction (right), with a neutral zone of zero friction at the center. On the far left lie cases of pure inefficiency: repeated failed convergence of an optimization routine due to poor hyperparameter choice with no conceptual payoff. Near the center lies the zero-friction ideal promoted by many AI systems—seamless prediction pipelines that return results without requiring the user to grapple with why those results are obtained. On the right lie productive friction zones: the stubborn mismatch between a density-functional-theory band-structure calculation and experimental photoemission data that forces the researcher to reconsider exchange-correlation functionals; the social friction of peer reviewers demanding mechanistic justification for a machine-learned stability prediction; or the material friction of a synthesized compound refusing to adopt the predicted crystal structure, thereby revealing overlooked entropic contributions. The spectrum is not binary but graduated, and the location of any given difficulty depends on context, researcher expertise, and research goals.

The epistemic value of friction becomes especially clear in materials science. When AI systems are deployed to accelerate inverse design, they can eliminate the very iterative struggle that once led Zunger to emphasize the importance of human-guided target specification [6]. Similarly, in the autonomous discovery pipelines reviewed by Montoya *et al.*, the drive toward full automation risks collapsing the friction that historically forced researchers to articulate explicit hypotheses about structure–property relationships [7]. Recent analyses of conceptual difficulties in AI-assisted materials curricula further illustrate how removing struggle can leave trainees with accurate predictions but impoverished mechanistic understanding [18]. Scientific friction, properly understood, is therefore not an enemy of progress but a necessary condition for the depth of insight that distinguishes genuine scientific advance from mere computational efficiency.

Types of Scientific Friction

Scientific friction in materials research does not arise as a uniform impediment. Still, it manifests through qualitatively distinct forms of resistance, each embedded in different layers of the scientific process and each contributing in specific ways to the production of robust knowledge. What becomes apparent is that these forms of friction are not obstacles to be eliminated, but generative constraints that force refinement, correction, and reorientation. Their epistemic value lies precisely in their capacity to interrupt smooth computational or experimental progression, thereby exposing hidden assumptions and stabilizing the relationship between model and reality.

One of the most consequential forms emerges at the level of conceptual coherence, where tensions between competing representations or theoretical commitments compel deeper scrutiny. In materials science, such tension becomes visible when machine-learned models produce predictions that conflict with established thermodynamic databases or long-standing physical principles. Rather than indicating simple model failure, these moments often reveal misalignments between data-driven inference and the approximations embedded in existing frameworks. The resulting need to reconcile discrepancies drives theoretical refinement, prompting reconsideration of boundary conditions, parameterizations, or even foundational assumptions. Empirical accounts, including those documented by Butler *et al.*, show that such conceptual resistance has historically catalyzed advances in molecular and materials modeling that would not have emerged through data accumulation alone [4].

A related but operationally distinct form of resistance arises within the methods themselves. Computational and experimental techniques impose constraints that limit what can be straightforwardly computed, measured, or simulated, and it is precisely these limitations that often guide methodological innovation. High-throughput screening workflows, for instance, frequently encounter convergence failures when applied to systems characterized by strong electronic correlations. These interruptions are not merely technical inconveniences; they signal the boundaries of current approximations and compel the development of more stable algorithms or the adoption of uncertainty-aware practices. As emphasized by Schmidt and colleagues, such methodological resistance has played a central role in validating and contextualizing machine-learning predictions in solid-state materials research, ensuring that computational outputs remain anchored in physically meaningful regimes [5].

Beyond conceptual and methodological domains, resistance is also encountered in the behavior of matter itself. Physical systems rarely conform perfectly to the abstractions imposed by models, and it is through this mismatch that critical insights often emerge. In synthesis-driven contexts, materials may exhibit structural features—such as unexpected octahedral tilting in perovskites—that deviate from predictions optimized through AI-guided design. These deviations expose gaps in the assumed relationships between composition, processing conditions, and emergent structure, thereby forcing a reconsideration of how these relationships are encoded. The recalcitrance of matter thus acts as a corrective force, grounding theoretical constructs in empirical reality and preventing the uncritical extension of idealized models. This grounding has been particularly significant in the development of inverse design strategies, where iterative confrontation with experimental outcomes refines both predictive models and synthesis protocols [6].

A further layer of friction operates within the social organization of scientific practice. Disagreement, critique, and iterative evaluation are not peripheral to knowledge production but constitute a core mechanism through which standards are maintained and improved. When, for example, the interpretability of a graph neural network applied to polymer design is questioned during peer review, the ensuing exchange does more than refine a single model; it contributes to the calibration of what counts as acceptable explanation within the field. Such interactions sustain a collective epistemic discipline, ensuring that claims are subjected to scrutiny that extends beyond individual judgment. Recent analyses of AI-driven materials design highlight the continuing importance of this form of resistance, particularly as automated systems begin to influence research agendas and evaluative norms [17].

Taken together, these forms of scientific friction delineate a distributed system of constraints through which understanding is progressively stabilized. Tensions at the conceptual level sharpen theoretical commitments, methodological limitations drive the evolution of tools, the resistance of material systems anchors abstraction in empirical observation, and social critique maintains the integrity of collective judgment. When AI systems are optimized primarily for efficiency and throughput, these sources of productive resistance risk being diminished, not through explicit removal but through systematic attenuation. The consequence is not merely an acceleration of discovery, but a transformation in the conditions under which knowledge is generated, with potential implications for the depth, reliability, and interpretability of scientific outcomes.

Table 1 analytically links each type of scientific friction to its epistemic function, the AI processes that most readily attenuate it, and the characteristic failure risk that follows from its removal.

Table 1. Analytical matrix linking types of scientific friction to epistemic functions, AI removal mechanisms, and resulting failure risks

Type of scientific friction	Core definition in materials science	Primary epistemic function	Typical AI process that attenuates it	Immediate risk when removed	Dominant downstream failure mode
Conceptual friction	Tension between competing theories, assumptions, or explanatory models	Forces theoretical refinement and explicit reconciliation of conflicting interpretations	Automated prediction pipelines, latent abstraction, optimization without interpretive checkpoints	Researchers accept outputs without confronting underlying physical contradictions	Conceptual smoothing

Methodological friction	Resistance arising from computational, modeling, or experimental limitations	Improves robustness, uncertainty awareness, and procedural scrutiny	Workflow automation, convergence masking, interface simplification, and black-box model deployment	Procedural assumptions become invisible, and methods are treated as self-validating	Methodological opacity
Material friction	Physical recalcitrance of matter relative to idealized predictions	Anchors models in empirical reality and exposes overlooked synthesis–structure–property relations	Generative smoothing, surrogate modeling, and abstraction away from laboratory variability	Real materials appear more compliant than they are, reducing sensitivity to mismatch and constraint	Premature closure
Social friction	Productive disagreement, peer critique, interpretive challenge, and community scrutiny	Maintains collective standards of explanation, justification, and evidential sufficiency	Seamless decision-support tools, persuasive automated outputs, and reduced contestability of model claims	Critique is displaced by acceptance of polished predictions, and consensus is reached too quickly	Conceptual smoothing/methodological opacity
Cross-friction interaction	Overlap among conceptual, methodological, material, and social resistance within real workflows	Sustains multi-layered epistemic depth rather than isolated correctness	Fully closed-loop autonomous systems that compress many stages into a frictionless pipeline	Multiple forms of resistance disappear simultaneously, and the narrowing of scientific reasoning	Compound failure: anomaly suppression, opacity, and premature closure

Mechanisms of Friction Loss

Contemporary AI systems reshape the epistemic landscape of materials science not only by accelerating discovery but by systematically reconfiguring the conditions under which understanding is produced. A central feature of this transformation lies in the attenuation of scientific friction through intertwined mechanisms that operate across cognitive, methodological, and representational levels. As human effort is displaced by automation, tasks that once required iterative hypothesis construction and falsification are increasingly executed by trained models, removing the very processes through which intermediate reasoning was formed and refined. What is lost is not simply time spent on routine prediction, but the generative labor through which conceptual structures were historically built [7].

This reconfiguration is further reinforced by the tendency of modern learning systems to privilege smoothness over irregularity. Training strategies that incorporate regularization or data augmentation are designed to stabilize predictions, yet in doing so, they suppress precisely those anomalous cases that often carry disproportionate scientific significance. Outliers that might once have prompted re-examination of underlying assumptions are absorbed into generalized patterns, effectively erasing signals of potential novelty or theoretical disruption [5]. The consequence is a subtle but consequential shift in what counts as relevant data, where the statistical objective of minimizing error competes with, and often overrides, the epistemic value of exception.

Alongside this smoothing process, abstraction introduces an additional layer of distance between the researcher and the phenomena under investigation. Latent representations condense complex atomic and electronic information into high-dimensional embeddings that are computationally efficient but conceptually opaque. While such representations enable powerful generalization, they also remove the necessity of engaging directly with the structural details that historically grounded insight. The researcher is no longer compelled to confront the full complexity of the system, and with that loss comes a diminished opportunity for discovering new relationships embedded within that complexity [4].

This distancing is compounded when systems are treated as black boxes, where the internal pathways linking input to output are neither accessible nor interrogated. Once deployed as predictive or generative oracles, these models no longer invite scrutiny of their internal reasoning processes. The absence of such engagement removes a critical site of epistemic tension, where assumptions would otherwise be questioned and refined. In this sense, opacity does not merely obscure technical detail; it forecloses the iterative dialogue between model and scientist that underpins conceptual advancement [19].

These mechanisms rarely operate in isolation. When automation, smoothing, abstraction, and opacity converge within a single workflow, they produce an environment in which multiple forms of friction are simultaneously attenuated. The resulting systems deliver outputs with remarkable speed and apparent reliability, yet the epistemic processes that would render those outputs intelligible are progressively diminished. What emerges is not simply an acceleration of scientific practice, but a reconfiguration of its underlying logic, where efficiency is gained at the potential expense of depth.

A Typology of Friction Failure Modes

Within this transformed landscape, distinct patterns of failure begin to surface as the absence of friction reshapes how knowledge is generated. One such pattern becomes evident when conceptual difficulty is effectively bypassed, allowing researchers to arrive at accurate predictions without engaging the underlying theoretical tensions. In these cases, understanding remains superficial, as the predictive success of the model substitutes for the explanatory work that would otherwise have been required. The result is a form of conceptual smoothing in which the appearance of clarity masks the absence of genuine insight [18].

A related dynamic unfolds when anomalous cases are systematically excluded from consideration. By privileging average behavior within the training process, models become less sensitive to rare configurations that might signal new physical regimes or unrecognized mechanisms. The suppression of these cases removes a critical source of epistemic disruption, reducing the likelihood that existing frameworks will be challenged or expanded. Historically, such anomalies have played a central role in catalyzing paradigm shifts, and their absence represents a significant loss for exploratory science [2].

Opacity at the methodological level introduces another layer of vulnerability. As models grow more complex and less interpretable, the ability of the research community to interrogate their internal assumptions diminishes. This loss is not merely technical but epistemic, as it constrains the capacity to evaluate how conclusions are produced and whether they rest on stable foundations. Without access to the internal structure of these systems, critical examination becomes increasingly difficult, and methodological opacity emerges as a defining feature of contemporary practice [17].

Finally, the rapid resolution of problems by AI systems can lead to a form of premature closure in which alternative conceptual pathways are never explored. When solutions are generated with minimal delay, the iterative process of questioning, revising, and reframing that characterizes scientific inquiry is curtailed. The system delivers an answer before the space of possible explanations has been meaningfully traversed, limiting the opportunity for deeper insight to emerge. Under such conditions, efficiency becomes a constraint rather than an advantage, as the speed of resolution forecloses the very processes that would enrich understanding [6].

Figure 1 illustrates how epistemically valuable forms of scientific friction in materials research are attenuated by AI mechanisms of automation, smoothing, abstraction, and black-boxing, thereby generating distinct failure modes and downstream epistemic consequences unless friction-aware design interventions are introduced.

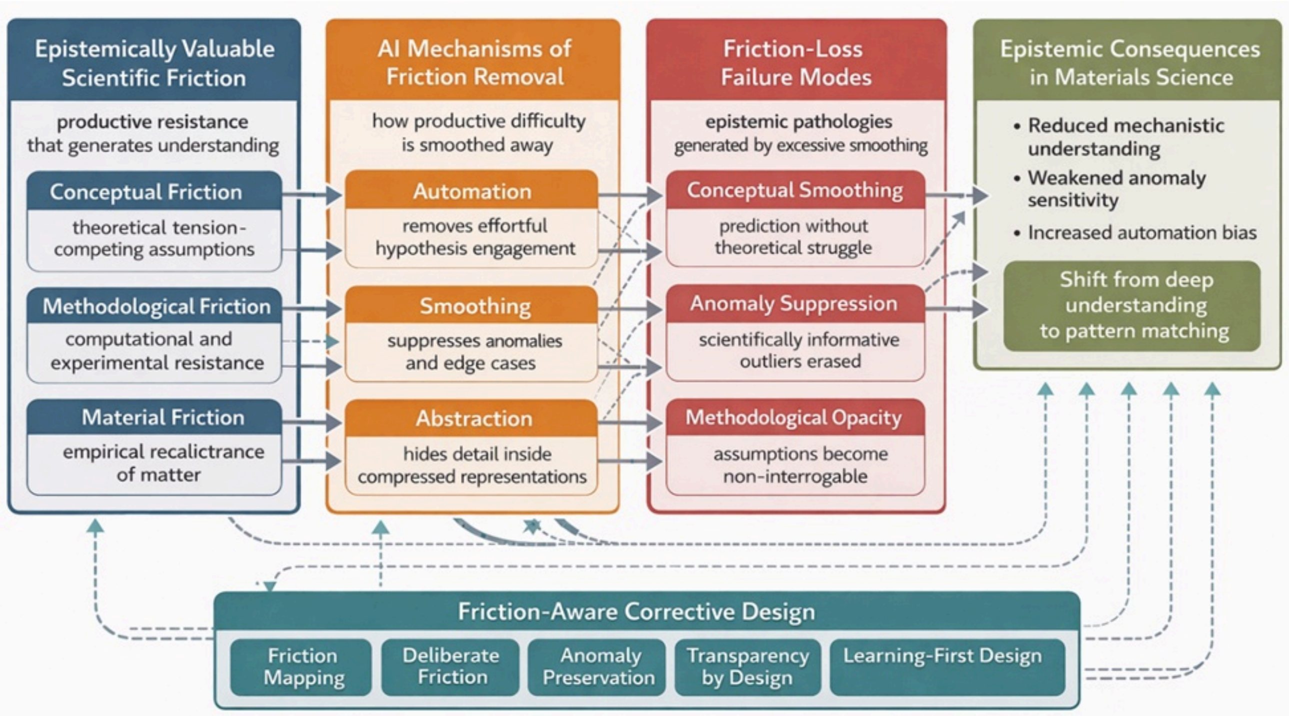

Figure 1. Architecture of scientific friction loss in AI-driven materials science.

Detection Principles

Detecting the unintended loss of scientific friction requires systematic, principle-based scrutiny rather than ad-hoc intuition. Four core principles provide researchers in artificial intelligence for materials science with practical tools to identify when valuable resistance is being eroded. These principles do not demand new experiments or metrics; instead, they prompt reflective audits of existing workflows.

Difficulty audit asks whether important conceptual difficulties are being bypassed rather than resolved. In practice, this means pausing at each major output of an AI pipeline—be it a predicted crystal structure or a generated phase diagram—and explicitly asking what conceptual tension the system has quietly resolved. As Montoya and colleagues observe in their roadmap for autonomous materials research [7], closed-loop systems frequently collapse iterative hypothesis refinement into a single optimized prediction. A difficulty audit would therefore reconstruct the chain of assumptions that the AI has smoothed away and determine whether the researcher has been deprived of the epistemic work that once accompanied manual iteration. In materials inverse design, for example, an audit might reveal that a generative model has eliminated the friction of reconciling predicted band gaps with known Mott-insulator physics, leaving the user with a plausible candidate but no deepened understanding of correlation effects [6].

Anomaly tracking requires deliberate preservation and interrogation of outliers that the model has down-weighted or suppressed. Schmidt *et al.* document how machine-learning pipelines in solid-state materials science routinely apply regularization that favors average behavior [5], thereby erasing the very anomalies that have historically signaled new physical regimes. An anomaly-tracking protocol involves extracting the top 5%–10% of highest-residual cases after each training epoch and forcing human review before they are discarded. When applied to polymer property prediction, this principle might surface a handful of structures whose unexpected thermal stability stems from overlooked conformational entropy—resistance that a frictionless model would have smoothed into the mean.

The transparency check evaluates whether the internal reasoning steps of the AI remain examinable. Malica and co-authors emphasize that advanced functional materials discovery increasingly relies on models whose decision pathways are opaque [19], preventing the critical examination that once accompanied every density-functional-theory run. A transparency check, therefore, demands that every deployed model expose at least one intermediate representation—atomic descriptors, attention weights, or uncertainty maps—that a domain expert can interrogate. If the only output is a final property value with a confidence score, the check fails and signals methodological opacity.

Learning Assessment measures whether interaction with the AI system demonstrably increases the researcher's conceptual grasp. Krotenko [18] has shown in engineering-education contexts that AI-assisted materials curricula can produce accurate predictions while leaving trainees with diminished mechanistic insight. A learning assessment, therefore, requires the researcher to articulate, before and after using the tool, the physical reasoning behind a given prediction. If the post-use explanation is no richer than the pre-use explanation, valuable friction has been lost. Collectively, these four principles transform detection from a vague feeling of "something missing" into a repeatable, auditable practice that can be embedded in any materials AI workflow.

Mitigation Principles

Preserving productive scientific friction does not require rejecting AI tools; it requires redesigning them so that efficiency and epistemic depth become mutually reinforcing. Five interlocking mitigation principles offer concrete guidance for achieving this balance.

Friction Mapping begins every project by cataloging which forms of resistance are epistemically productive and which are merely wasteful. Researchers explicitly label each anticipated difficulty—conceptual mismatches in phase stability, methodological convergence failures, material recalcitrance during synthesis, or social critique during peer review—according to its expected epistemic yield. Prasittisopin and colleagues illustrate how such mapping can be integrated into AI-driven materials design pipelines [17], ensuring that only unproductive friction is targeted for removal.

Deliberate Friction deliberately reintroduces calibrated resistance at key junctures. Rather than allowing an automated inverse-design loop to converge in a single pass, the system can be engineered to pause and present the user with two competing conceptual framings—e.g., a purely data-driven stability ranking versus a physics-informed free-energy argument—requiring explicit reconciliation before proceeding [20-22]. This principle draws directly on the manifesto of Radicchi *et al.* [22], who advocate balancing automation with retained epistemic friction in AI4Materials workflows.

Anomaly Preservation mandates that edge cases flagged by the model remain visible and interactively explorable rather than being silently filtered. Hartung and co-authors demonstrate in their review of agentic models that preserving such anomalies is essential for materials innovation [20], because the most transformative discoveries often originate precisely where models exhibit the greatest resistance. An implementation might involve a dedicated "friction dashboard" that surfaces the highest-uncertainty predictions and invites the user to manipulate underlying assumptions before the model re-trains.

Transparency by Design embeds interpretability mechanisms into the architecture from the outset. Channing *et al.* argue that friction and human engagement must remain central even in highly automated materials research [21], which requires that every black-box component be accompanied by an auditable trace of decision steps. For graph neural networks applied to battery materials, this could mean exposing layer-wise attention maps that reveal which atomic interactions the model considers most salient, thereby restoring the opportunity for conceptual engagement.

Learning-First Design prioritizes user understanding over raw speed of output. Instead of optimizing solely for the lowest mean-absolute error, the objective function includes a term that rewards models whose predictions are accompanied by explanations that demonstrably advance researcher insight. This principle aligns with the educational findings of Krotenko [18], who shows that conceptual understanding improves only when difficulty is retained rather than erased. When these five principles are applied in concert, AI systems in materials science shift from friction-erasing engines into friction-curating partners, safeguarding the epistemic resources that have historically driven the field forward.

Relation to Other Failure Modes

The failure mode of scientific friction loss does not exist in isolation; it intersects with and amplifies several other recognized failure modes in AI-assisted research [23-26].

First, friction loss directly contributes to epistemic debt. When AI systems smooth away conceptual difficulties, they accumulate hidden assumptions that later prove costly to repay. The 2025 paper on scientific friction itself identifies this dynamic in materials science pipelines [3], where the apparent efficiency of automated discovery masks an accumulating debt of unexamined physical intuitions. Recht [26] similarly describes how frictionless reproducibility in data-driven materials science creates exactly this form of debt by hiding the messy details that ground claims.

Second, friction loss exacerbates automation bias—the tendency of users to over-trust system outputs precisely because the outputs arrive without visible resistance. When every prediction is delivered smoothly and instantaneously, researchers lose the habitual skepticism that once accompanied slower, friction-rich workflows. Zohar *et al.* highlight this risk in their analysis of the hidden costs of removing productive struggle [8], noting that automation bias is most pronounced precisely when epistemic friction has been engineered away.

Third, friction loss manifests as scaffolding failure. In educational and discovery contexts alike, removing support structures before conceptual mastery has been achieved leaves learners or researchers unable to proceed independently. Borrego and colleagues document this exact pattern in AI-assisted materials curricula [18], where premature removal of methodological friction prevents the development of robust independent reasoning. The friction-loss failure mode, therefore, functions as a meta-failure that intensifies epistemic debt, automation bias, and scaffolding collapse, forming a mutually reinforcing web that threatens the long-term epistemic health of materials AI.

Implications for Materials AI Practice

The identification of scientific friction as a distinct failure mode carries immediate, actionable implications for three stakeholder groups within the materials AI community [27-29].

For authors, three practices become essential. First, every manuscript must explicitly identify what forms of friction the proposed AI system removes and justify whether that removal is epistemically warranted. Second, authors must argue for or against the value of retained friction using the typology and mechanisms articulated here. Third, authors should design and report friction-aware variants of their tools—e.g., an optional “friction mode” that reintroduces deliberate conceptual resistance. These steps align directly with the community-level concerns raised by Channing *et al.* regarding the need for human engagement in AI-driven materials research [21].

For reviewers, two new evaluation criteria emerge. Reviewers must ask whether the submitted work has preserved or eliminated productive friction and must question any claim of “seamless” or “frictionless” performance. When a paper reports only accuracy gains without addressing epistemic depth, reviewers should request a difficulty audit or anomaly-tracking analysis. Such scrutiny echoes the calls for balanced automation found in the AI4Materials manifesto [22].

For the broader community, two strategic initiatives are required. First, dedicated studies of friction in real-world materials workflows should be funded and published, building on the conceptual foundations laid by Geirsdotter Bækkelund *et al.* [1] and Suleimenov *et al.* [2] while grounding them in contemporary AI practice. Second, the community must develop and adopt friction-aware design guidelines—checklists, software templates, and review rubrics—that embed the detection and mitigation principles outlined above. Together, these changes would transform materials AI from a purely efficiency-driven enterprise into one that deliberately curates the productive difficulties essential for sustained scientific progress.

Table 2 translates the paper’s conceptual analysis into a friction-aware evaluation framework that can be used to assess whether AI systems in materials science preserve or erode the epistemic conditions of genuine understanding.

Table 2. Friction-aware evaluation framework for AI systems in materials science

Evaluation dimension	Key diagnostic question	Warning sign of friction loss	Recommended friction-aware response	Primary stakeholder
Difficulty retention	Does the system preserve important conceptual difficulty rather than bypassing it?	Output is accurate, but the user cannot reconstruct what conceptual tension was resolved	Conduct a difficulty audit and require explicit articulation of the hidden conceptual struggle	Authors/reviewers
Anomaly visibility	Are high-residual, exceptional, or contradictory cases preserved for human examination?	Outliers are filtered, regularized away, or omitted from reporting	Implement anomaly-tracking protocols and require reporting of informative residual cases	Authors/tool designers
Process examinability	Are intermediate reasoning steps, descriptors, or uncertainty structures accessible to domain experts?	The system behaves as an oracle and exposes only final scores or rankings	Build transparency by design through interpretable intermediates and auditable traces	Tool designers/reviewers
Learning gain	Does interaction with the system deepen the user’s mechanistic understanding?	User performance improves, but explanation quality does not	Incorporate learning-first design and require pre/post explanation checks in evaluation	Educators/authors
Closure timing	Does the system terminate the inquiry too quickly relative to the complexity of the scientific problem?	A candidate solution is accepted before rival explanations or pathways are explored	Introduce deliberate friction through staged comparison and forced alternative interpretation	Authors/tool designers
Empirical resistance	Does the workflow remain responsive to	Model outputs appear cleaner and more stable	Preserve material mismatch signals and	Tool designers/experimental

sensitivity	material recalcitrance and synthesis mismatch?	than experimental reality warrants	integrate friction dashboards around failed transfer cases	researchers
Critical contestability	Can peers, collaborators, or reviewers productively challenge the system's outputs and assumptions?	Automated recommendations are treated as persuasive endpoints rather than debatable claims	Embed social friction through review prompts, challenge interfaces, and justification requirements	Community/reviewers
Epistemic balance	Is the workflow optimizing only efficiency, or also the depth of understanding?	Seamlessness and speed are presented as unqualified indicators of progress	Balance performance metrics with friction-aware criteria tied to explanation, uncertainty, and anomaly engagement	Authors/community

Conclusion

Scientific friction—the productive resistance, conceptual struggle, and material recalcitrance that have long been the engine of genuine insight in materials science—is being systematically erased by AI systems optimized solely for speed and smoothness. This failure mode analysis has defined scientific friction, distinguished its four epistemically valuable types, identified the four mechanisms of its loss, and articulated a typology of four resulting failure modes. Detection and mitigation principles now exist to preserve what must not be lost. The field stands at a choice point: continue building frictionless pipelines that deliver superficial mastery or redesign AI tools to curate productive difficulty. A friction-aware materials AI that respects and preserves conceptual struggle is not only possible but necessary if the community is to maintain the depth of understanding required for transformative discovery. The future of the discipline depends on recognizing that some resistance is not a bug to be fixed but the very feature that makes scientific progress meaningful.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 02 Mar 2024 Revised: 02 Apr 2024 Accepted: 16 May 2024

Published online: 18 July 2024

References

Geirsdotter Bækkelund N, Isdal Haugland A, Sande J, Rogde Mjånes JO. Ill-structured problems in capstone courses–investigating students' approaches and experience. Eur J Eng Educ. 2025;50(1):1-16. https://doi.org/10.1080/03043797.2025.2585122.

Suleimenov I, Gabrielyan O, Vitulyova Y. Dialectics of scientific revolutions from the point of view of innovations theory. Wisdom. 2022;4(24):25-35.

Madsen AK, Munk AK, Søltoft JI. Friction by machine: How to slow down reasoning with computational methods. In: Ethnographic Praxis in Industry Conference Proceedings. New York, NY: EPIC; 2023;2023(1):65-88.

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):011405.

Kemper J, Jankowski S. Silicon Valley's frictionless future: The design philosophy of frictionlessness. In: De Gruyter handbook of automated futures: Imaginaries, interactions and impact. Berlin: De Gruyter; 2024. pp. 287-302.

Ramsoomair N. Pressing matters: How AI irons out epistemic friction and smooths over diversity. Atlantis. 2025;46(1):42-55.

Yan L, Pammer-Schindler V, Mills C, Nguyen A, Gašević D. Beyond efficiency: Empirical insights on generative AI's impact on cognition, metacognition and epistemic agency in learning. Br J Educ Technol. 2025;56(5):1675-85.

Wright CS. Cognitive castes: Artificial intelligence, epistemic stratification, and the dissolution of democratic discourse. arXiv preprint arXiv:2507.14218. 2025.

Van den Bulck H, Kulig M, Hyzen A, Puppis M, Paulussen S. Epistemic welfare and public service media's algorithmic recommender systems: A theoretical framework, operationalization and relevance for governance. Eur J Commun. 2025;40(6):598-618.

Dessì D, Osborne F, Reforgiato Recupero D, Buscaldi D, Motta E, Sack H. AI-KG: An automatically generated knowledge graph of artificial intelligence. In: International Semantic Web Conference. Cham, Switzerland: Springer; 2020. pp. 127-43.

Cong-Lem N. Rethinking evidence-informed policy and practice in the age of generative artificial intelligence. Lond Rev Educ. 2025;23(1):1-8.

Hernández AD, Mosquera D, Gallegos F. Lessons from the margins: Contextualizing, reimagining, and hacking generative AI in the Global South. Harv Data Sci Rev. Cambridge, MA: MIT Press; 2025;7(4). https://doi.org/10.1162/99608f92.30a2934f.

Obiso T, Lai K, Nath A, Krishnaswamy N, Pustejovsky J. Dynamic epistemic friction in dialogue. In: Proceedings of the 29th Conference on Computational Natural Language Learning. Stroudsburg, PA: Association for Computational Linguistics; 2025. pp.323-33.

Prasittisopin L. Artificial intelligence-driven materiality: Asystematic review and perspective. Eng Sci. 2025;35:1587. https://doi.org/10.30919/es1587.

Krotenko TY. Engineering economics and technological education: Atransdisciplinary approach to the training of modern engineers. PNRPU Bull Soc Econ Sci. Perm: Perm National Research Polytechnic University; 2022;(4):63-75.

Malica C, Novoselov KS, Barnard AS, Kalinin SV, Spurgeon SR, Reuter K, et al. Artificial intelligence for advanced functional materials: Exploring current and future directions. J Phys Mater. 2025;8(2):021001.

Hartung T. AI, agentic models and lab automation for scientific discovery—the beginning of scAInce. Front Artif Intell. 2025;8:1649155.

Wang H, Fu T, Du Y, Gao W, Huang K, Liu Z, et al. Scientific discovery in the age of artificial intelligence. Nature. 2023;620(7972):47-60.

Radicchi E, Syed U, Cunico F, Cassetta M, Khan U, Marone P, et al. AI4Materials: A manifesto for AI-driven scientific discovery in materials science. In: Proceedings of the 4th Italian Conference on Artificial Intelligence (Ital-IA 2024). Aachen: CEUR-WS; 2025:1-12.

Lu C, Lu C, Lange RT, Foerster J, Clune J, Ha D. The AI scientist: Towards fully automated open-ended scientific discovery. arXiv preprint arXiv:2408.06292. 2024.

Craps S. Remembering earth: Countering planetary amnesia through the creative arts. In: Dynamics, mediation, mobilization: Doing memory studies with Ann Rigney. Berlin: De Gruyter; 2024. pp. 313-18.

Groß R, Jordan R, editors. KI-Realitäten: Modelle, praktiken und topologien maschinellen lernens. Bielefeld: transcript Verlag; 2023.

Recht B. The mechanics of frictionless reproducibility. Harv Data Sci Rev. 2024;6(1). https://doi.org/10.1162/99608f92.f5b28d6d.

Ghosh R, Spano FC. Excitons and polarons in organic materials. Acc Chem Res. 2020;53(10):2201-11.

Boelt AM, Kolmos A, Holgaard JE. Literature review of students' perceptions of generic competence development in problem-based learning in engineering education. Eur J Eng Educ. 2022;47(6):1399-420.

Yaayin B. Uniqueness of card game intervention in improving students' conceptual understanding of chemical names and their formulae. Eur J Educ Pedagog. 2025;6(4):62-70.