ORIGINAL RESEARCH Open access

What Counts as Knowledge in AI-Driven Materials Science? A Philosophical and Practical Reframing

Maria Hernandez[1*], Carlos Vega[1]

Abstract

The integration of artificial intelligence (AI) into materials science represents a profound epistemic shift, challenging longstanding assumptions about the nature and validation of scientific knowledge. Rather than merely accelerating computational tasks, AI reconfigures the epistemological landscape by generating outputs that blur the boundaries between data, inference, and insight. This paper diagnoses a central problem: "knowledge inflation," where AI's predictive prowess is prematurely equated with genuine understanding, leading to overconfidence in materials-related decisions. Tensions arise between the opacity of AI-driven predictions and the demands for explanation and mechanistic clarity inherent to materials science, where structure-property relationships and causal processes have traditionally grounded epistemic warrant. Such discrepancies risk undermining the reliability of knowledge claims in domains like alloy design and sustainable material selection. To address this, we propose a reframed epistemic framework tailored to materials AI, centered on actionability as the criterion for knowledge: what can be responsibly acted upon in practical contexts. This includes a novel typology distinguishing epistemic categories of AI outputs, from mere predictive signals to robust decision warrants, with conditions for elevation between them. By emphasizing responsibility and scope, this reframing aims to safeguard epistemic integrity while harnessing AI's potential.

Keywords Materials science, AI epistemology, Knowledge typology, Epistemic responsibility, Actionability, Predictive inference

*Correspondence:
Maria Hernandez
maria.hernandez@outlook.com
[1] Department of Materials Science and Machine Learning, Faculty of Engineering, Polytechnic University of Valencia, Valencia, Spain

Introduction

The advent of artificial intelligence (AI) in materials science has prompted a reevaluation of what constitutes knowledge within the discipline. Traditionally, materials knowledge has been anchored in empirically derived mechanisms, structure-property relations, and process understandings that enable reliable interventions in physical systems. These elements form the bedrock of scientific practice, allowing researchers to design, select, and deploy materials with confidence in their performance and implications [1, 2]. In contrast, AI-generated outputs—such as neural network predictions, unsupervised learning latent embeddings, or representation model similarity scores—introduce novel forms of inference that do not always align with these classical epistemic standards. This misalignment raises fundamental questions about the status of AI contributions: when does an AI output transcend mere computation to become bona fide knowledge? What epistemic thresholds must be met for a prediction to qualify as a material insight? And who bears responsibility for claims derived from such systems?

This reconfiguration is not merely a technical enhancement but an epistemological disruption. In traditional materials science, knowledge emerges from iterative cycles of observation, hypothesis formulation, and verification, often grounded in physical principles like thermodynamics or quantum mechanics. For instance, understanding phase transformations in alloys relies on mechanistic accounts that link atomic arrangements to macroscopic properties, providing not just descriptive accuracy but explanatory depth [3-5]. AI, however, often operates through correlative patterns extracted from vast datasets, yielding outputs that excel in forecasting but may lack transparency regarding underlying causes. This shift challenges the positivist legacy in science, which holds that knowledge is justified by verifiable evidence and logical coherence [4, 6]. In materials contexts, where decisions impact real-world applications—from structural integrity in aerospace to environmental sustainability in energy storage—the stakes are high. Misclassifying an AI prediction as knowledge could lead to suboptimal designs or unforeseen failures, amplifying epistemic risks.

Consider the applied domains of materials design and screening. In designing new composites for lightweight applications, traditional approaches emphasize invariant relations between composition, microstructure, and mechanical properties, ensuring that knowledge claims are actionable and reproducible [7]. AI enables high-throughput screening by identifying candidate materials through probabilistic matches. Yet, these suggestions may not reveal why a particular composition is favored, potentially conflating statistical association with causal necessity. Similarly, in sustainability decisions, such as selecting low-carbon footprint polymers, AI might prioritize optimization scores without accounting for lifecycle mechanisms, leading to claims that appear robust but rest on fragile foundations [8]. These examples illustrate how AI expands the epistemic horizon while introducing ambiguities: outputs that are highly accurate in narrow contexts may falter when generalized, casting doubt on their status as knowledge.

At the heart of this inquiry are three foundational questions. First, when does an AI output become knowledge? Classical epistemology posits that knowledge requires justified true belief, but in AI-driven materials science, justification often hinges on model performance rather than conceptual grasp [9]. Second, what distinguishes a prediction from a material insight? Predictions offer probabilistic forecasts, whereas insights demand integration with domain-specific principles, such as defect dynamics or interfacial behaviors, to yield understanding [10]. Third, who is epistemically responsible for AI-derived claims? In collaborative settings, responsibility may diffuse across human experts, algorithms, and data curators, complicating accountability in cases of error or misapplication [11].

This paper’s contribution is strictly epistemological, not methodological. It does not advocate for specific AI techniques or protocols but instead offers a philosophical reframing to clarify what counts as knowledge in this hybrid landscape. By diagnosing tensions with classical epistemologies and proposing a materials-specific typology, we aim to foster a more rigorous discourse. This approach draws on recent reflections in AI scholarship, adapting them to the unique demands of materials science, where knowledge must support tangible actions, such as processing optimization or deployment strategies [12]. In doing so, we underscore that epistemic value lies not in computational sophistication but in the responsible integration of AI outputs into scientific practice.

The ensuing sections build this argument progressively. We first outline the epistemic background, contrasting classical views with AI’s incursions. Subsequently, we introduce a reframed typology that categorizes AI outputs by epistemic strength and actionability. This framework emphasizes that knowledge in materials AI is defined by what can be responsibly acted upon, prioritizing practical consequences over abstract metrics. Through this lens, we seek to mitigate the risks of epistemic overreach while preserving the transformative promise of AI in advancing materials innovation.

Epistemic Background: Knowledge, Understanding, and AI in Materials Science

Classical views of scientific knowledge and understanding

Classical epistemologies of science have long emphasized the interplay of empiricism, explanation, mechanism, and actionability as pillars of knowledge. Empiricism, rooted in observational data, posits that knowledge arises from sensory evidence corroborated by repeatable phenomena [13]. Explanation extends this by requiring accounts that unify disparate observations under coherent principles, as in deductive-nomological models where laws predict and elucidate outcomes [14]. Mechanisms add depth by identifying the underlying processes—causal chains or entities—that generate observed effects, transforming mere description into comprehension [15]. Finally, actionability integrates these elements, defining knowledge as that which enables informed intervention and aligning with pragmatic philosophies, where truth is gauged by practical utility [16].

In scientific contexts, these dimensions ensure that claims are not isolated assertions but embedded in a web of justification. For instance, knowledge is not just an accurate prediction but the capacity to manipulate variables with foreseeable results, underscoring reliability and scope [17]. The transition from classical research paradigms to AI-mediated discovery necessitates a comparison of foundational pillars; these shifts in evidence, depth, and justification are summarized in **Table 1.**

Table 1. Comparative epistemic dimensions

Dimension	Classical materials science	AI-driven materials science
Primary evidence	Physical observation and repeatable experiment.	Statistical correlation and pattern recognition in large datasets.
Explanatory goal	Mechanistic depth (the “Why”).	Predictive accuracy (the “What”).
Invariants	Laws of thermodynamics, quantum mechanics, and symmetry.	Latent representations and learned embedding spaces.
Validation	Theoretical coherence and counterfactual robustness.	Performance metrics (RMSE, \$R^2\$) and cross-validation scores.
Actionable basis	Causal understanding of structure-property relations.	Probabilistic forecasts and optimization scores.

What “understanding” traditionally means in materials science

In materials science, understanding has historically centered on mechanisms, invariants, and structure-property relations that bridge atomic-scale phenomena to macroscopic behaviors. Mechanisms reveal how processes like diffusion or dislocation motion dictate material responses, providing a causal narrative essential for reproducibility [18]. Invariants, such as conserved quantities in phase equilibria, offer stable principles that transcend specific instances, enabling generalization across material classes [19]. Structure-property relations encapsulate these, linking crystallographic arrangements or chemical compositions to attributes like strength or conductivity, forming the epistemic core of the field [20].

This conception prioritizes depth over breadth, ensuring that understanding supports practical endeavors like alloy formulation or defect engineering. Unlike superficial correlations, it demands integration with physical laws, fostering confidence in applications where failure has tangible costs [21].

How AI outputs enter materials discourse as putative knowledge

AI outputs infiltrate materials discourse through conceptual pathways that elevate raw computations to knowledge claims. Predictions from supervised models, for example, are often framed as insights when aligned with domain expectations, blurring the line between statistical inference and scientific warrant [22]. Embeddings from representation learning provide latent spaces that suggest relational structures, interpreted as emergent knowledge despite their data-driven origins [23].

Scores from optimization algorithms similarly gain epistemic status by implying hierarchies of material viability, integrating into narratives of discovery [24].

These pathways rely on analogical reasoning, in which AI mimics human intuition, yet they risk conflating pattern recognition with principled understanding, thereby altering how claims are validated within the discipline [25].

The problem of "accuracy-as-knowledge" in materials AI

The equation of accuracy with knowledge poses a core epistemic shortfall in materials AI. Performance metrics, while indicative of predictive fidelity, do not suffice for epistemic warrant, as they overlook explanatory gaps and contextual limits [26]. High accuracy in controlled datasets may mask brittleness in novel scenarios, leading to claims that appear justified but lack robustness [27]. This problem stems from AI's inductive bias, which prioritizes correlation over causation, in contrast to materials science's demand for mechanistic grounding [28].

Epistemically, accuracy alone fosters a reductive view, where knowledge is quantified rather than qualified, potentially eroding the discipline's foundational rigor [29].

Consequences

Misclassifying AI outputs as knowledge incurs significant epistemic risks in materials science. In design contexts, false confidence may lead to errors, such as selecting brittle composites under unmodeled conditions, thereby compromising structural integrity [30]. Sustainability misjudgments arise when predictions ignore lifecycle mechanisms, leading to environmentally suboptimal choices [31]. Broader consequences include eroded trust in scientific claims, where overreliance on opaque outputs hampers collaborative verification [32].

These risks highlight the need for epistemic vigilance, ensuring that AI enhances rather than undermines the action-oriented nature of materials knowledge [33, 34].

From knowledge typology to applicability boundaries and scientific silence

The typology outlined above does not merely classify AI outputs by epistemic strength; it also establishes the conditions under which non-action becomes the most responsible scientific response. In materials science, where decisions often carry irreversible economic, environmental, or safety consequences, the absence of an epistemic warrant must be treated not as a temporary inconvenience but as a decisive boundary. This reframes applicability domains not as technical diagnostics but as normative limits on knowledge claims.

Within this perspective, the applicability domain defines the region in which elevation through the epistemic hierarchy is possible. Predictive signals may exist well beyond this domain, but they are epistemically confined to the lowest tier and cannot be legitimately elevated. Explanatory relations that emerge near the boundary may appear compelling, yet without robustness or domain alignment, they remain conditionally informative rather than actionable. Mechanistic accounts, design rules, and causal claims, by contrast, presuppose interiority within the applicability domain, where physical constraints, representational validity, and contextual stability jointly hold.

Scientific silence emerges precisely at the point where this presupposition fails. When an AI output cannot satisfy the epistemic conditions required for elevation—whether due to extrapolation, proxy misalignment, or unresolved uncertainty—the correct response is not qualification or narrative repair, but abstention. Silence, in this framework, is not epistemic emptiness; it is an active refusal to convert insufficient knowledge into action.

This reconceptualization directly counters a pervasive tendency in AI-driven materials research: the pressure to always produce an answer. The typology makes explicit that not all outputs are candidates for epistemic ascent, and that forcing them to be elevated constitutes overreach. A prediction outside the applicability domain does not merely become less accurate; it becomes epistemically illegitimate as a basis for action. Silence, therefore, functions as a safeguard against both technical error and ethical failure.

Importantly, silence is tier-dependent. At the level of predictive signals, silence may delay exploration. At higher tiers—particularly those involving causal claims or decision warrants—silence prevents unjustified interventions that could have potentially severe downstream consequences. This asymmetry explains why uncertainty alone cannot define abstention thresholds; instead, silence must be triggered by a mismatch between epistemic tier and decision stakes.

By embedding applicability domains within the knowledge typology, the framework clarifies why many contemporary failures in materials AI are not failures of modeling, but failures of epistemic governance. Models are routinely asked to support claims—mechanistic, causal, or decision-oriented—that their outputs are structurally incapable of warranting. The result is not simply an error, but misplaced authority.

Scientific silence thus becomes the final containment mechanism in the epistemic lifecycle of materials AI. It marks the boundary beyond which outputs may exist computationally but cannot exist epistemically. Treating silence as a first-class outcome restores symmetry to the decision space: models may recommend action, recommend caution, or explicitly refuse to recommend at all. Only the first of these has traditionally been recognized as success; this framework insists that the latter two are equally essential for trustworthy materials discovery and deployment. **Figure 1** maps the epistemic hierarchy, providing a framework for identifying the applicability boundaries inherent in the current model.

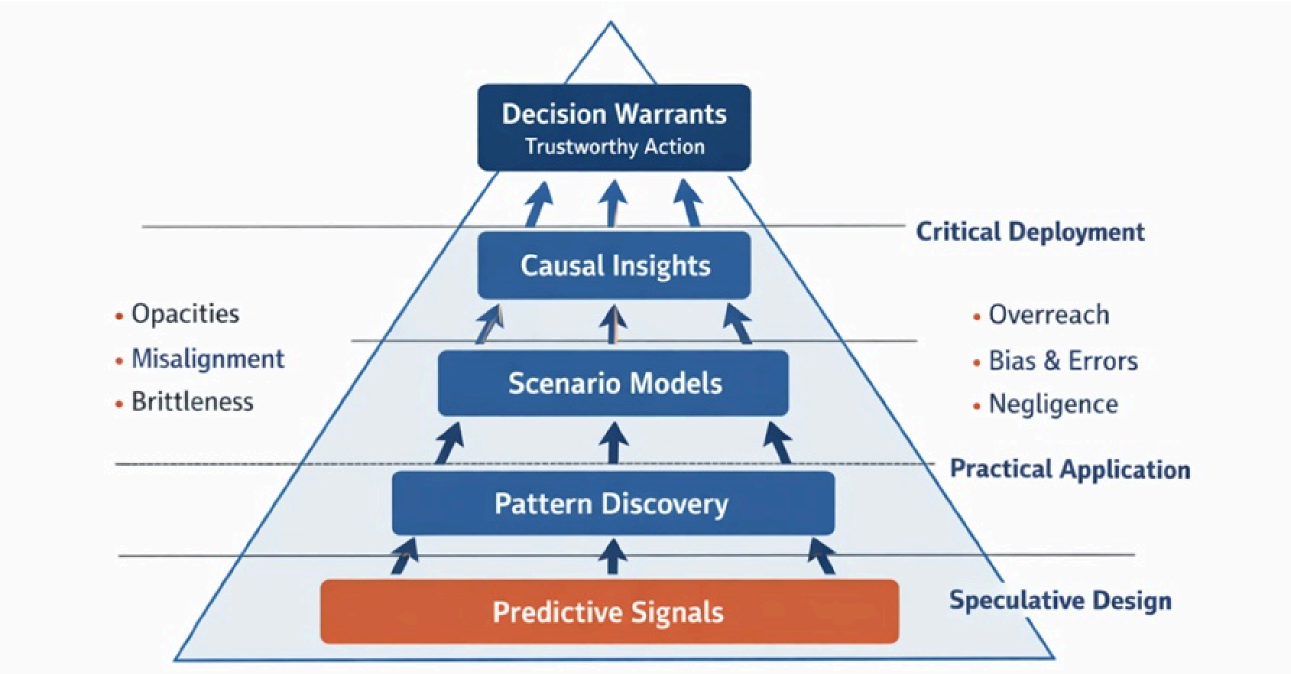

Figure 1. Epistemic hierarchy, applicability boundaries, and zones of scientific silence

Constructive epistemic frameworks

Building upon the typology introduced in the previous section, this part develops constructive epistemic frameworks for validating AI outputs in materials science. The intent is not to prescribe technical pipelines or evaluation metrics, but to articulate epistemic conditions under which AI outputs may be legitimately elevated from computational artifacts to actionable knowledge. These frameworks are explicitly materials-specific, reflecting the field's sensitivity to physical realism, scale coupling, and downstream consequences. They are grounded in the central thesis that knowledge is defined not by predictive success alone, but by responsible actionability.

Rather than treating validation as a retrospective exercise applied after model development, the frameworks conceptualize epistemic validation as an integral component of knowledge production itself. By embedding epistemic criteria upstream—at the level of observation, explanation, mechanism, and responsibility—the framework aims to prevent familiar pathologies in materials AI, including semantic overreach, proxy substitution, and unjustified extrapolation. In this sense, epistemic rigor functions as a containment principle rather than a corrective one. To prevent epistemic overreach, we propose a formal hierarchy for AI outputs. Table 2 outlines the typology of these outputs and the specific epistemic conditions required to elevate a computational signal to a decision warrant.

Table 2. Typology of AI outputs and conditions for elevation

Epistemic tier	Output type	Primary condition for elevation	Practical application
Tier 1: Signal	Predictive scores/latent embeddings	Observational grounding: Tethering to physical invariants/referents.	Initial screening; hypothesis generation.
Tier 2: Relation	Feature attributions/attention maps	Relational coherence: Alignment with structure-property paradigms.	Directing experimental focus; processing optimization.
Tier 3: Insight	Causal claims/counterfactuals	Mechanistic alignment: Integration with physical kinetics (e.g., diffusion).	Design rules for novel material classes.
Tier 4: Warrant	Actionable knowledge	Responsibility and accountability: Expert audit of stakes and consequences.	Deployment in high-stakes structural or safety systems.

The four dimensions that structure these frameworks—observation, explanation, mechanism, and responsibility—correspond to progressively stronger forms of epistemic commitment. Each dimension introduces additional constraints that an AI output must meet to warrant elevation within the knowledge typology. Together, they ensure that the transition from output to knowledge is disciplined, transparent, and commensurate with the stakes of materials decision-making.

Observation: Data fidelity and epistemic foundations

Observation constitutes the epistemic foundation of materials science. Traditionally, it refers to empirical access to physical reality through techniques such as diffraction, microscopy, spectroscopy, or mechanical testing. In AI-mediated contexts, observation must be reinterpreted to address the fidelity of data representations and outputs relative to these physical referents. Predictive signals generated by AI systems only qualify as epistemically meaningful insofar as they remain tethered to observable material phenomena.

Within this framework, observational validity requires more than dataset size or statistical coverage. It demands congruence with known invariants of materials systems—such as stoichiometric constraints, symmetry relations, or conservation laws—and vigilance against artifacts introduced by biased sampling, spurious correlations, or overfitting. Outputs that violate such constraints may appear numerically plausible while remaining physically incoherent.

For AI outputs to be elevated beyond the lowest tier of the typology, they must demonstrate explicit observational grounding. In the case of learned embedding spaces, this grounding requires demonstrable correspondence with measurable properties, such as lattice parameters, coordination environments, or phase labels. Without such empirical tethering, embeddings risk becoming abstract surrogates detached from material reality. This framework, therefore, treats observational alignment as a necessary precondition for subsequent epistemic elevation. Absent this grounding, even sophisticated models remain speculative and cannot responsibly inform actions such as material screening or selection [1].

https://doi.org//a045179042

Explanation: Interpretability and relational coherence

Explanation extends observational grounding by articulating relations among variables, transforming isolated signals into intelligible structures. Within this epistemic framework, explanation is not equated with visualization or post-hoc rationalization, but with the capacity of AI outputs to support coherent accounts that are compatible with materials-science reasoning. Interpretability thus becomes an epistemic requirement rather than a user-interface convenience.

Explanatory validity hinges on relational coherence: explanations must consistently link inputs, representations, and outputs in ways that align with established structure–property–processing paradigms. Feature attributions, attention patterns, or graph-based message passing can serve this role when they reveal stable and domain-relevant dependencies, such as the influence of local coordination on electronic or mechanical behavior. When such relations generalize across contexts, they may support elevation to the level of design rules.

However, the framework also acknowledges the epistemic risks posed by opacity in high-capacity models. Complex architectures may generate internally consistent explanations that are nevertheless physically irrelevant. To mitigate this risk, the framework emphasizes cross-validation of AI-generated explanations against theoretical models or well-established empirical regularities. Hybrid interpretability—where AI explanations are constrained or audited using domain theory—serves as a safeguard against narrative overreach and ensures that explanatory claims remain proportionate to their epistemic support. Only under these conditions can explanations responsibly inform processing decisions or guide experimental exploration [13].

Mechanism: Physical alignment and causal depth

Mechanistic understanding has long occupied a privileged epistemic position in materials science, enabling reasoning across scales and supporting intervention rather than mere description. In the proposed framework, mechanistic accounts represent a decisive elevation in epistemic commitment. AI outputs qualify as mechanistic knowledge only when they can be coherently integrated with underlying physical processes, such as diffusion, phase transformation kinetics, electronic structure evolution, or microstructural development.

This requirement explicitly rejects the conflation of correlation with causation. Predictive success alone is insufficient for mechanistic status; outputs must be reconciled with causal principles that remain stable under intervention. A central condition here is counterfactual robustness: mechanistic claims must support credible answers to “what-if” questions, such as how a system would respond to altered composition, processing conditions, or environmental exposure. Without this robustness, causal interpretations risk collapsing into correlative mirages.

Mechanistic alignment also plays a critical role in sustainability-sensitive applications. When AI outputs are used to forecast long-term performance, degradation, or environmental impact, mechanistic grounding prevents erroneous extrapolation by anchoring predictions in lifecycle-relevant processes. By tying AI forecasts to physically meaningful pathways, the framework ensures that elevated knowledge claims remain resilient across operating regimes and temporal scales, rather than being optimized for narrow datasets [18].

Responsibility: Accountability, consequence, and epistemic restraint

While observation, explanation, and mechanism address the internal validity of knowledge claims, responsibility governs their legitimate use. This dimension recognizes that materials AI operates in contexts where decisions can have irreversible consequences, ranging from structural failure to environmental harm. As such, epistemic validation must explicitly incorporate accountability and awareness of consequences.

Within this framework, responsibility requires clear attribution of who bears epistemic and practical accountability for actions informed by AI outputs. It also requires proportionality between epistemic strength and decision stakes. Outputs that may be acceptable for exploratory design become epistemically insufficient when used to justify deployment or scale-up. In such cases, the absence of sufficient warrant does not invite creative interpretation but demands restraint.

Responsibility, therefore, formalizes the role of scientific silence introduced earlier. When epistemic conditions cannot be met—whether due to domain extrapolation, unresolved uncertainty, or missing mechanistic grounding—the responsible outcome is abstention. Silence is not a failure of the framework but its successful execution, preventing the conversion of weak knowledge into high-stakes action.

Responsibility: Human-AI collaboration and accountability

Responsibility integrates the preceding dimensions, framing knowledge as a collaborative product where humans retain oversight. The framework delineates epistemic roles: AI provides provisional outputs, while experts assess scope and consequences. Decision warrants emerge when responsibility is explicitly assigned, incorporating ethical considerations like equity in resource use.

Guidelines include traceability protocols and failure audits to counter mechanisms such as authority transfer. This dimension emphasizes that epistemic responsibility safeguards against risks, ensuring AI-derived knowledge supports sustainable practices without diffusing accountability [30].

These frameworks collectively offer a scaffold for epistemic practice, promoting a balanced integration of AI in materials science. By prioritizing actionability, they align with the discipline’s applied ethos, enabling knowledge that is both innovative and trustworthy.

Implications for Materials Practice

The reframed epistemology and typology have profound implications for materials practice, reshaping how AI informs design, screening, and sustainability decisions. By defining knowledge in terms of actionability, this approach encourages a cautious yet productive use of AI, mitigating epistemic pitfalls while amplifying its utility.

In materials design, the typology distinguishes mere predictions from design rules, urging practitioners to seek explanatory and mechanistic depth. For alloy development, AI signals might suggest compositions, but only those aligned with mechanisms—such as solid-solution strengthening—qualify as knowledge, reducing the risk of brittle failures. This reframing promotes hybrid workflows where AI augments human insight, fostering designs that are robust across operational envelopes [2].

Screening processes benefit from emphasizing decision warrants, where outputs must account for scope limitations. In high-throughput virtual screening for catalysts, predictive signals are insufficient; elevation requires causal claims validated against physical constraints, ensuring selections prioritize long-term viability over short-term metrics. This mitigates narrative overreach, which can lead to inefficient resource allocation [22].

Sustainability decisions highlight the responsibility dimension, as AI outputs must integrate environmental mechanisms. In selecting eco-friendly polymers, explanatory relationships must link molecular features to degradation pathways, yielding knowledge that supports circular-economy goals. Failure modes such as proxy substitution—equating energy scores with holistic impact—are averted by requiring comprehensive warrants, thereby enhancing ethical practice [31].

Conceptually, consider a scenario in battery materials: AI predicts high-capacity electrodes, but without mechanistic accounts of ion transport, these remain signals. By applying the frameworks, practitioners probe for causal depth, elevating outputs into actionable insights that inform scalable production. This not only accelerates innovation but also builds epistemic resilience against uncertainties, such as supply chain disruptions.

Future directions involve embedding these frameworks in educational curricula, cultivating a generation attuned to epistemic nuances. Collaborative initiatives could standardize validation protocols, bridging academia and industry. Ultimately, this reframing positions AI as a partner in knowledge creation, ensuring materials science advances responsibly toward societal challenges.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 01 Aug 2023 Revised: 22 Oct 2023 Accepted: 15 Nov 2023
Published online: 18 January 2024

References

Abolhasani M, Brown KA. Role of AI in experimental materials science. MRS Bull. 2023;48:1-2.

Sha W, Guo Y, Yuan Q, Tang S, Zhang X. Artificial intelligence to power the future of materials science and engineering. Adv Intell Syst. 2020;2(4):1900143.

DeCost BL, Hattrick-Simpers JR, Trautt Z. Scientific AI in materials science: a path to a sustainable and scalable paradigm. Mach Learn Sci Technol. 2020;1(3):033001.

Liu Y, Yang Z, Yu Z, Liu Z, Liu D, Lin H, et al. Generative artificial intelligence and its applications in materials science: Current situation and future perspectives. J Mater Sci Technol. 2023;158:150-70.

Cheng Y, Wang T, Zhang G. Artificial intelligence for materials science. Cham: Springer; 2021.

Green ML, Maruyama B, Schrier J. Autonomous (AI-driven) materials science. Appl Phys Rev. 2022;9(3):030401.

Goswami L, Deka MK, Roy M. Artificial intelligence in material engineering: A review on applications of artificial intelligence in material engineering. Adv Eng Mater. 2023;25(20):2300104.

Pilania G. Machine learning in materials science: From explainable predictions to autonomous design. Comput Mater Sci. 2021;193:110360.

Kalidindi SR. Feature engineering of material structure for AI-based materials knowledge systems. J Appl Phys. 2020;128(4):041103.

Li J, Lim K, Yang H, Ren Z, Raghavan S, Chen PY. AI applications through the whole life cycle of material discovery. Matter. 2020;3(2):393-432.

Choudhary K, DeCost B, Chen C, Jain A. Recent advances and applications of deep learning methods in materials science. NPJ Comput Mater. 2022;8(1):59.

Epps RW, Volk AA, Reyes KG, Abolhasani M. Accelerated AI development for autonomous materials synthesis in flow. Chem Sci. 2021;12(17):6025-36.

Badini S, Regondi S, Pugliese R. Unleashing the power of artificial intelligence in materials design. Materials (Basel). 2023;16(17):5927.

Qayyum F, Kim DH, Bong SJ, Chi SY, Choi YH. A survey of datasets, preprocessing, modeling mechanisms, and simulation tools based on AI for material analysis and discovery. Materials (Basel). 2022;15(4):1428.

Zhong X, Gallagher B, Liu S, Kailkhura B. Explainable machine learning in materials science. NPJ Comput Mater. 2022;8(1):204.

Tamura R, Tsuda K, Matsuda S. NIMS-OS: an automation software to implement a closed loop between artificial intelligence and robotic experiments in materials science. Sci Technol Adv Mater Methods. 2023;3(1):2232297.

López C. Artificial intelligence and advanced materials. Adv Mater. 2023;35(16):2208683.

Vasoya NH. Revolutionizing nano materials processing through IoT-AI integration: opportunities and challenges. J Mater Sci Res Rev. 2023;11:1-11.

Jablonka KM, Ai Q, Al-Feghali A, Badhwar S. 14 examples of how LLMs can transform materials science and chemistry: A reflection on a large language model hackathon. Digit Discov. 2023;2(5):1233-50.

Gao C, Min X, Fang M, Tao T, Zheng X. Innovative materials science via machine learning. Adv Funct Mater. 2022;32(19):2108044.

Jha D, Gupta V, Liao W, Choudhary A, Agrawal A. Moving closer to experimental level materials property prediction using AI. Sci Rep. 2022;12(1):11720.

Guo K, Yang Z, Yu CH, Buehler MJ. Artificial intelligence and machine learning in design of mechanical materials. Mater Horiz. 2021;8(4):1153-72.

Oviedo F, Ferres JL, Buonassisi T. Interpretable and explainable machine learning for materials science and chemistry. Acc Mater Res. 2022;3(11):1071-81.

Takeda S, Kishimoto A, Hamada L, Nakano D. Foundation model for material science. Proc AAAI Conf Artif Intell. 2023;37(13):15681-8.

Geng X, Wang F, Wu HH, Wang S, Wu G. Data-driven and artificial intelligence accelerated steel material research and intelligent manufacturing technology. Mater Genome Eng Adv. 2023;1(1):e10.

Xu Y, Liu X, Cao X, Huang C, Liu E, Qian S, et al. Artificial intelligence: A powerful paradigm for scientific research. Innovation (Camb). 2021;2(4):100179.

Batra R, Song L, Ramprasad R. Emerging materials intelligence ecosystems propelled by machine learning. Nat Rev Mater. 2021;6(8):655-78.

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

Zhao XG, Zhou K, Xing B, Zhao R, Luo S, Li T, et al. JAMIP: an artificial-intelligence aided data-driven infrastructure for computational materials informatics. Sci Bull. 2021;66(6):624-30.

Pyzer-Knapp EO, Pitera JW, Staar PWJ. Accelerating materials discovery using artificial intelligence, high performance computing and robotics. NPJ Comput Mater. 2022;8(1):84.

Gregoire JM, Zhou L, Haber JA. Combinatorial synthesis for AI-driven materials discovery. Nat Synth. 2023;2(6):493-504.

Pan G, Wang F, Shang C, Wu H, Wu G, Gao J. Advances in machine learning-and artificial intelligence-assisted material design of steels. Int J Miner Metall Mater. 2023;30(3):405-23.

Karthikeyan A, Priyakumar UD. Artificial intelligence: machine learning for chemical sciences. J Chem Sci. 2022;134(1):2.

Huang B, von Rudorff GF, von Lilienfeld OA. The central role of density functional theory in the AI age. Science. 2023;380(6651):1359-64.