ORIGINAL RESEARCH Open access

The Limits of Scale in Materials AI: A Conceptual Argument Against 'Bigger Is Always Better'

Wei Liu[1*], Zhang Min[1]

Abstract

In the rapidly evolving field of materials artificial intelligence (AI), the prevailing emphasis on scaling data volumes and computational resources has driven significant advancements in predictive modeling and discovery processes. However, this conceptual manuscript interrogates the implicit assumption that larger scales invariably yield superior outcomes, positing instead that unchecked expansion introduces intricate interaction dynamics that undermine the integrity of materials informatics. Through an integrative analysis, we explore how escalating data scales interact with inherent biases, leading to amplified distortions in representational fidelity and epistemic reliability. The framework delineates trade-offs wherein quantitative abundance may erode qualitative depth, fostering feedback structures that perpetuate homogeneity in material explorations at the expense of diversity. Ethical reasoning underscores the epistemic implications, revealing how scale-driven approaches can inadvertently prioritize dominant paradigms, marginalizing underrepresented material classes and contexts. Systems-level insights highlight steering logics that balance scale with interpretive nuance, advocating for calibrated integrations that preserve domain-specific insights. This argument reframes scale not as an unequivocal virtue but as a contingent factor within broader conceptual interpretations, urging a reevaluation of priorities in applied AI for materials science to foster sustainable and equitable progress.

Keywords Materials informatics, Artificial intelligence, Epistemic values, Data scale, Bias dynamics, Interaction trade-offs

*Correspondence:
Wei Liu
wei.liu@gmail.com
[1] Department of Materials Informatics and Data Modeling, School of Materials Science, Shanghai Jiao Tong University, Shanghai, China

Introduction

The integration of artificial intelligence (AI) into materials science has transformed the landscape of discovery and design, enabling unprecedented efficiencies in navigating complex material spaces. Materials AI, often encompassed under the umbrella of materials informatics, leverages machine learning algorithms to process vast datasets, infer patterns, and predict properties that would otherwise require laborious experimental or computational efforts [1]. This synergy has accelerated innovations in areas such as energy storage, catalysis, and structural materials, where predictive models inform the synthesis of novel compounds with tailored functionalities [2]. The allure of AI lies in its capacity to handle multidimensional data, uncovering correlations that elude traditional analytical methods and facilitating a shift from trial-and-error paradigms to data-guided strategies [3].

Central to this evolution is the role of scale—both in terms of data volume and model complexity. Contemporary trends in AI emphasize the benefits of "big data," in which expansive datasets purportedly improve model generalization and accuracy [4]. In materials science, this manifests as repositories such as the Materials Project or Automatic FLOW for Materials Discovery (AFLOW), which aggregate millions of entries on crystal structures, electronic properties, and thermodynamic stabilities [5]. Scaling initiatives draw inspiration from successes in other domains, such as natural language processing, where larger models, such as transformers, have demonstrated emergent capabilities [6]. Proponents argue that increased scale mitigates uncertainties, refines predictions, and democratizes access to advanced simulations, thereby expediting the materials innovation cycle [7].

Yet, this enthusiasm for scale warrants a nuanced conceptual examination, particularly in the context of materials AI, where data generation is constrained by physical realities and resource limitations [8]. Unlike domains with abundant, low-cost data, materials datasets often derive from high-fidelity simulations or experiments that are time-intensive and costly, leading to inherent selectivities in what is captured [9]. This selectivity introduces interaction dynamics that scale may exacerbate rather than resolve, prompting a reevaluation of the “bigger is always better” ethos. Analytically, scale interacts with data provenance, where the origins and curation of datasets influence the interpretive validity of AI outputs [10]. For instance, reliance on computationally derived data may privilege certain material classes, such as inorganic crystals, over organics or amorphous systems, creating feedback structures that reinforce existing knowledge silos [11].

Epistemic reasoning further illuminates these dynamics, highlighting how scale can dilute the contextual richness essential to materials understanding. Material properties are not merely numerical abstractions but emerge from intricate atomic interactions influenced by environmental conditions, processing histories, and measurement variabilities [12]. Large-scale AI approaches risk oversimplifying these nuances, prioritizing statistical aggregates over domain-specific interpretations that account for anomalies or edge cases [13]. Ethical dimensions compound this, as scale-driven paradigms may perpetuate biases embedded in historical data, affecting the equity of material discoveries and their societal applications [14]. For example, the underrepresentation of sustainable or low-abundance materials in large datasets could steer innovation toward resource-intensive options, raising questions about their long-term viability [15].

This manuscript advances a conceptual argument against unbridled scaling by synthesizing literature on materials informatics and interrogating its underlying assumptions through interpretive lenses. We delineate systems-level insights into how scale interfaces with bias, diversity, and epistemic integrity, revealing trade-offs that demand calibrated steering logics [10]. Rather than dismissing scale, the analysis interprets it as a double-edged construct: a facilitator of breadth but a potential inhibitor of depth. Conceptual interpretations emphasize interaction dynamics, such as amplification loops where larger datasets magnify subtle biases, leading to homogenized predictive landscapes [6]. In materials AI, this could manifest as overconfidence in models that excel on benchmark tasks but falter in real-world extrapolations [4]. The multidimensional nature of scale and its non-linear epistemic consequences are summarized in **Table 1**.

Table 1. Scaling dimensions and their epistemic consequences in materials AI

Scaling dimension	Operational meaning in materials AI	Primary epistemic benefit	Emergent limitation at high scale	Conceptual risk identified
Data volume	Expansion of training datasets via simulations, repositories, or experiments	Improved coverage and statistical learning	Masking of contextual gaps and rare phenomena	Epistemic dilution through aggregation
Model capacity	Increase in parameters and architectural complexity	Enhanced pattern extraction	Memorization over generalization	Brittle extrapolation beyond dominant regimes
Computational throughput	High-performance computing and parallelization	Faster discovery cycles	Optimization bias toward efficiency	Speed prioritized over interpretive depth

Dataset reuse	Recursive use of benchmark datasets	Standardization and comparability	Reinforcement of historical skews	Amplified representational bias
Automation level	Reduced human mediation in inference pipelines	Scalability and reproducibility	Loss of domain-guided correction	Reduced epistemic reflexivity

Theoretically, this invites a rethinking of value alignments in scientific modeling, where epistemic goals—such as explanatory power and robustness—may conflict with scale-oriented efficiencies [5]. Drawing on values in science, we consider how AI frameworks embody implicit priorities, such as speed over thoroughness, and how these shape material outcomes [16]. The proposed framework integrates these elements, offering analytical implications for balancing scale with interpretive fidelity. By focusing on feedback structures, it elucidates how excessive scale might contract exploratory spaces, limiting serendipitous discoveries that arise from targeted, smaller-scale inquiries [7].

In synthesizing the literature, we highlight subthemes like data bias and ethical reasoning, underscoring their interplay with scale. Ultimately, this conceptual exploration advocates a paradigm in which scale serves as a tool rather than a dictate, fostering integrative approaches that honor the multifaceted nature of materials science. This sets the stage for a detailed examination of theoretical backgrounds and the articulation of a novel framework that navigates these complexities.

Theoretical Background and Literature Synthesis

Evolution of data-driven approaches in materials science

The emergence of data-driven methodologies represents a fundamental epistemic shift in materials science, marking a transition from predominantly reductionist, physics-first paradigms toward integrative, informatics-centered modes of inquiry. Traditionally, materials discovery and characterization relied on a combination of empirical heuristics, experimental trial-and-error, and theory-driven computational approaches, most notably density functional theory (DFT), which offered physically grounded yet computationally intensive predictions of material behavior [1]. While these methods yielded deep mechanistic insights, their scalability was inherently constrained by computational cost, simplifying assumptions, and the narrow scope of tractable material systems.

Advances in high-performance computing and automated simulation pipelines catalyzed the generation of large, structured materials datasets, enabling a gradual reorientation toward data-centric reasoning. This shift gave rise to materials informatics as a distinct disciplinary domain, wherein machine learning and artificial intelligence are leveraged to extract patterns, correlations, and predictive relationships across vast chemical and structural spaces [8, 9, 17]. Rather than replacing physics-based understanding, data-driven approaches increasingly function as connective tissue between theory and experiment, facilitating rapid hypothesis generation, surrogate modeling, and high-throughput screening workflows [3].

A critical infrastructural development in this evolution has been the establishment of standardized, open materials databases that aggregate computational and experimental property data across diverse material classes [5]. These repositories have enabled unprecedented cross-referencing, reuse, and pattern discovery, lowering barriers for researchers to predict phase stability, electronic structure, and thermomechanical performance without engaging in costly simulations [2]. At the same time, literature synthesizes how the democratization of predictive capability introduces new interaction dynamics between algorithmic inference and domain expertise. In many workflows, model outputs increasingly mediate scientific judgment, subtly reconfiguring how intuition, theory, and evidence are weighted in decision-making processes [1].

Systems-level analyses highlight a central tension in this transition: while data-driven approaches substantially enhance efficiency and exploratory breadth, they also introduce structural dependencies on data quality, representativeness, and provenance that do not diminish with scale alone [9]. As a result, the epistemic authority of AI-enabled predictions becomes inseparable from the assumptions embedded in data generation and curation practices, setting the stage for deeper scrutiny of scaling, bias, and value alignment in materials AI.

Scaling paradigms in AI and their application to materials

Within artificial intelligence research, scaling has emerged as a dominant paradigm, referring to the systematic expansion of model capacity, training data volume, and computational resources to achieve improved performance across benchmark tasks [4]. Materials science has rapidly adopted this paradigm in response to the intrinsic high dimensionality of materials spaces, where compositional, structural, and processing variables interact nonlinearly across multiple length and time scales [6]. Large-scale learning architectures, inspired by successes in computer vision and natural language processing, have been adapted to materials property prediction, often outperforming traditional handcrafted descriptors and linear models [15].

Graph-based representations have proven particularly amenable to scaling in materials contexts, as graph neural networks can encode atomic connectivities and local environments across chemically diverse systems [18]. Empirical studies demonstrate that performance gains often track increases in dataset size and model complexity, reinforcing the perception that scale is a primary driver of predictive accuracy. Consequently, scaling has become closely associated with progress narratives in materials AI, positioning ever-larger models as solutions to longstanding challenges in generalization and transferability.

However, conceptual analyses complicate this narrative by foregrounding trade-offs that accompany aggressive scaling. As models grow in capacity, they may increasingly rely on memorization of dominant patterns rather than learning robust, transferable representations, resulting in brittle performance when confronted with sparsely sampled or novel regions of materials space [7]. Feedback structures further emerge wherein scaled models reinforce prevailing data distributions, amplifying attention toward well-studied material classes while marginalizing unconventional or emerging systems [10]. Such dynamics risk narrowing the effective exploration space, even as nominal coverage expands.

Ethical and normative considerations intersect with these technical concerns. The literature increasingly questions whether scaling-driven optimization aligns with the values of inclusivity, openness, and pluralism in scientific inquiry, particularly when resource-intensive models privilege well-funded institutions or established research agendas [16]. In applied contexts such as alloy design, hybrid strategies that incorporate active learning and adaptive sampling have been proposed to counterbalance the limitations of scaling by iteratively targeting underexplored regimes [6]. These approaches illustrate how scaling, while powerful, requires complementary governance mechanisms to ensure epistemic resilience and exploratory diversity.

Challenges of data bias and representation in large datasets

Data bias constitutes one of the most persistent and structurally embedded challenges in materials AI, arising from asymmetries in how materials data are generated, selected, and curated [10]. Syntheses of recent literature reveal that biases frequently originate upstream, for example, through the overrepresentation of thermodynamically stable crystalline materials in computational repositories, reflecting historical research priorities and methodological convenience [11]. When such skewed datasets are used for training, models internalize and propagate these imbalances, shaping downstream predictions in subtle yet consequential ways.

Analytical studies demonstrate that scaling can exacerbate these effects rather than dilute them. As datasets grow, small initial biases are often amplified through repeated reuse and model retraining, yielding increasingly homogenized predictive landscapes that obscure rare, metastable, or unconventional material behaviors [6]. Cross-domain investigations have proposed quantitative metrics to detect representational disparities, highlighting systematic undercoverage of amorphous, organic, or hybrid materials in widely used benchmarks [4].

The interaction between bias and scale is therefore nontrivial. Larger datasets may create an illusion of comprehensiveness, masking qualitative gaps in representational fidelity [13]. Case studies in synthesizability prediction and materials screening illustrate how such biases directly undermine reliability, leading to overconfident recommendations that fail under experimental validation [5]. In response, recent work advocates for debiasing strategies that explicitly account for scaling dynamics, emphasizing diversity-aware sampling, uncertainty-guided exploration, and taxonomy-balanced dataset construction [19]. Specific amplification pathways through which scale intensifies bias in materials AI workflows are analytically mapped in **Table 2**.

Table 2. Bias amplification pathways under increasing scale

Bias source	Mechanism of introduction	Effect of scaling	Feedback structure	Downstream impact on discovery
Dataset curation bias	Preference for stable, well-characterized materials	Overrepresentation grows with volume	Recursive retraining on skewed data	Homogenized prediction landscapes
Domain availability bias	Abundance of inorganic vs. organic data	Scale amplifies asymmetry	Model-guided data acquisition	Marginalization of underexplored classes
Simulation bias	Reliance on DFT-accessible systems	Expanded but narrow chemical space	Benchmark reinforcement loops	Overconfidence in computational feasibility
Benchmark bias	Optimization toward standard datasets	Performance inflated with size	Metric-driven model selection	Illusory generalization
Measurement bias	Inconsistent experimental conditions	Noise absorbed into aggregates	Aggregation masking variability	Reduced interpretive reliability

Collectively, this body of literature underscores the need for steering logics that deliberately balance scale with representational diversity. Rather than treating data expansion as an inherently corrective force, contemporary frameworks emphasize the active governance of dataset composition to ensure equitable epistemic coverage across materials classes and application domains [20].

Epistemic and ethical dimensions in materials AI

The integration of AI into materials science raises foundational questions about how scientific knowledge is generated, justified, and valued. Epistemic values such as coherence, explanatory depth, robustness, and fruitfulness have long guided theory construction and experimental interpretation in the physical sciences [5]. Literature at the intersection of AI and philosophy of science interrogates how these values are operationalized—or sidelined—within data-driven workflows, where predictive performance often supersedes interpretability or causal understanding [21].

Systems-level analyses reveal feedback loops wherein scale-driven efficiencies incentivize opaque modeling practices, progressively eroding epistemic transparency as models evolve into black-box systems [22]. This opacity complicates scientific accountability, particularly when AI outputs inform high-stakes decisions related to materials selection, sustainability, or safety. Ethical dimensions extend these concerns by foregrounding issues of data provenance, labor distribution, and societal impact associated with large-scale computational infrastructures [14].

Conceptual interpretations highlight persistent trade-offs between exploratory breadth and normative responsibility. While scaled AI systems enable rapid traversal of vast material spaces, they may implicitly marginalize ethical priorities such as environmental sustainability or long-term societal benefit if these values are not explicitly encoded into modeling objectives [15]. In response, reviews on open and responsible science advocate for transparent, value-aware AI frameworks that integrate epistemic rigor with ethical reflexivity [8].

Within materials informatics, this translates into growing calls for value-aligned AI systems that incorporate diverse stakeholder perspectives, including experimentalists, domain theorists, and societal end-users [23]. The emerging consensus emphasizes integrative approaches that treat epistemic and ethical considerations not as external constraints, but as constitutive elements of AI-enabled materials research. Such synthesis positions materials AI as a socio-technical enterprise, whose legitimacy depends on its capacity to respect the interpretive complexity and normative commitments of scientific practice [24].

Proposed conceptual framework

The proposed conceptual framework reframes scale in materials artificial intelligence as a relational and dynamic construct, rather than a unidirectional driver of progress. Instead of treating increases in data volume, model capacity, or computational throughput as linear enhancers of discovery, the framework conceptualizes scale as a system of interacting forces that collectively shape epistemic trajectories, interpretive practices, and research priorities within materials science. In this view, scale functions as a steering logic—a structuring influence that governs how attention, resources, and inferential authority are distributed across material spaces.

At the core of the framework lies the recognition that data volume, model complexity, and computational capacity operate as interdependent nodes within a feedback-rich system. Their interactions generate trade-offs that are not readily apparent when scale is assessed solely through performance metrics. Systems-level insights reveal that scaling decisions implicitly prioritize certain forms of knowledge—typically those that are statistically dominant, computationally convenient, or historically well-represented—while marginalizing others. As a result, scale does not merely accelerate discovery; it actively shapes what kinds of materials knowledge become visible, legible, and actionable.

Analytical implications emerge most clearly in the interaction between scale and bias dynamics. As datasets expand, representational biases arising from selective data curation—such as the overrepresentation of stable crystalline phases or well-characterized inorganic compounds—are amplified through recursive training and reuse [10]. These amplification loops transform scale into an active architect of the knowledge landscape, reinforcing dominant material narratives while attenuating signals associated with rare, metastable, or emergent systems [13]. From an epistemic standpoint, this process produces dilution effects, where quantitative abundance masks qualitative insufficiency, eroding the granularity required for nuanced interpretation.

Interaction dynamics further illuminate how scale interfaces with data quality. Large datasets can obscure inconsistencies, measurement heterogeneity, and contextual gaps, fostering a false sense of robustness grounded in aggregation rather than justification [6]. In such settings, predictive confidence may increase even as epistemic reliability stagnates or declines. The framework, therefore, treats abundance not as a proxy for validity but as a condition that intensifies the need for interpretive safeguards and reflexive oversight.

Ethical reasoning is embedded within the framework as a constitutive dimension rather than an external evaluative layer. Scale-driven optimization regimes often encode implicit value hierarchies, privileging high-throughput discovery and performance maximization over sustainability, inclusivity, or long-term societal relevance [14]. Systems-level analysis reveals persistent tensions between breadth, enabled by large-scale exploration, and depth, which depends on targeted curation, domain expertise, and contextual interpretation. These tensions necessitate steering mechanisms that calibrate expansion with normative and epistemic constraints, rather than assuming alignment by default [16].

Conceptual interpretations also emphasize feedback from model outputs into data-generating practices. As scaled AI systems increasingly guide experimental design and data acquisition, they risk reproducing their own representational biases, creating self-reinforcing cycles of homogeneity that progressively contract exploratory diversity [10]. Within this

https://doi.org//q661501684

framework, such feedback is not treated as an implementation flaw but as a structural consequence of unmoderated scaling.

By foregrounding these interactions, the framework advances integrative reasoning that situates scale as a contingent, relational property whose value emerges through its coupling with bias management, epistemic governance, and ethical alignment. Analytical implications suggest that beyond a certain threshold, unchecked scaling may yield diminishing returns, where marginal improvements in predictive accuracy are offset by losses in interpretability, reliability, and exploratory openness [22].

To formalize the interaction dynamics articulated conceptually, scale can be expressed as a non-linear epistemic function rather than a purely performance-enhancing parameter. In this framing, epistemic reliability is not monotonically proportional to data volume or model capacity; rather, it is mediated by bias amplification, representational dilution, and interpretive attenuation, which intensify under recursive scaling. The following conceptual expression captures the threshold-dependent trade-off between quantitative expansion and qualitative epistemic integrity.

$$\left[E_{\{r\}(S)} \right. = \left. \frac{\{\alpha \;,\; \log(1+S)\}}{\{1 \;+\; \beta\, B(S) \;+\; \gamma\, D(S)\}} \right] \quad (1)$$

This formulation expresses scale as an asymptotically bounded epistemic driver. While early scaling produces logarithmic gains in reliability through expanded coverage and pattern learning, the denominator captures countervailing forces that intensify with scale. Bias amplification and contextual dilution progressively erode epistemic integrity, generating diminishing—or even negative—returns beyond critical thresholds. The model, therefore, encodes scale as a contingent rather than intrinsically beneficial construct.

The framework thus provides an interpretive lens for diagnosing contemporary AI practices in materials. It motivates hybrid strategies that deliberately combine scale with domain-guided refinement, diversity-aware sampling, and reflexive control mechanisms [15]. As illustrated in the conceptual framework **(Figure 1)**, scale in materials AI operates not as a neutral technical parameter but as a steering logic, where the pursuit of data volume, model capacity, and computational throughput engages in complex feedback loops with epistemic and ethical outcomes.

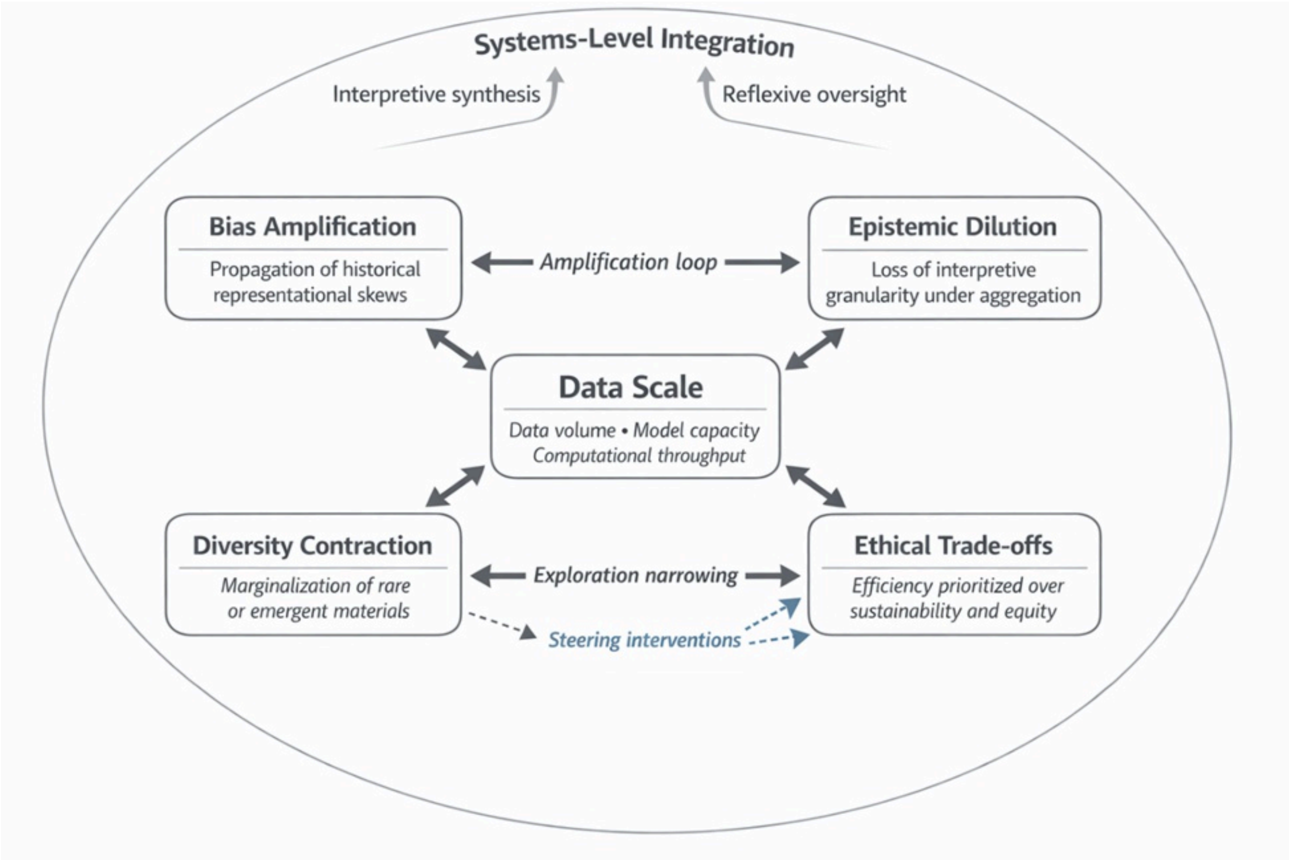

Figure 1. Conceptual framework illustrating scale as a steering logic in materials AI

This framework advances integrative reasoning by reframing scale as a relational entity, whose value emerges through its interactions rather than isolation. Analytical implications suggest that unchecked scaling may engender diminishing returns, where incremental gains in predictive power are offset by losses in epistemic reliability [22]. By focusing on these dynamics, the framework offers a lens for interpreting AI practices in materials and advocates hybrid approaches that blend scale with domain-guided refinements [15].

Analytical implications

The conceptual framework advanced in this manuscript illuminates analytical implications for understanding scale in materials AI, revealing how interaction dynamics between data abundance and bias propagation shape the epistemic landscape of materials informatics [6]. As scale expands, the framework suggests an amplification of inherent distortions, where large datasets, while offering broader coverage, often intensify subtle skews derived from data curation practices [2]. This dynamic implies a trade-off in which quantitative growth in data volume may compromise representational accuracy, particularly for material classes underrepresented in computational repositories, such as amorphous or organic systems [3]. Systems-level insights indicate that this amplification creates feedback structures that favor dominant material paradigms, potentially limiting the exploration of diverse chemical spaces and hindering the identification of novel compounds with unique properties [5]. **Table 3** synthesizes the steering logics proposed in this study, linking scaling risks to actionable epistemic and ethical countermeasures.

Table 3. Steering logics for calibrating scale in materials AI

Scaling risk	Observed failure mode	Steering logic	Epistemic objective	Ethical alignment

Bias amplification	Reinforced dominant material narratives	Diversity-aware dataset construction	Representational balance	Inclusivity across material classes
Epistemic dilution	Loss of contextual interpretation	Hybrid AI–domain expert workflows	Interpretive depth	Accountability in inference
Overconfidence	Narrow confidence intervals with weak justification	Uncertainty-aware modeling	Epistemic humility	Risk-sensitive deployment
Exploration contraction	Reduced novelty discovery	Active learning toward sparse regimes	Exploratory openness	Equitable innovation pathways
Efficiency dominance	Speed prioritized over understanding	Value-aligned optimization objectives	Knowledge robustness	Sustainability and long-term impact

Interpretive analysis extends this to epistemic reliability, highlighting how scale-driven models can foster overconfidence in predictions that align with aggregated patterns but overlook contextual nuances, such as processing histories or environmental influences [9]. For instance, in predictive tasks for material stability, large-scale AI might prioritize statistical correlations over physical interpretability, leading to outputs that are efficient yet epistemically shallow [10]. Ethical reasoning is integral here, underscoring the implications for equity in material innovations; amplified biases could steer developments toward resource-intensive materials, marginalizing sustainable alternatives, and exacerbating societal disparities in applications such as energy storage or biomedical devices [14]. The framework's steering logics propose calibrated integrations, in which scale is modulated by domain-specific constraints to preserve interpretive depth, suggesting analytical pathways for hybrid approaches that combine big data with targeted validation [5].

Further, interaction dynamics between scale and model complexity reveal analytical implications for generalization in materials AI. As models scale in parameter count, they enhance pattern recognition but introduce vulnerabilities to data noise, leading to feedback loops that erode robustness in extrapolative scenarios [15]. Conceptual interpretations frame this as a tension between breadth and precision, where excessive scale may homogenize output landscapes, contracting the space for serendipitous discoveries [10]. Systems-level insights advocate for trade-offs that incorporate bias-aware mechanisms, such as diversity metrics in dataset construction, to counteract these effects and foster more inclusive material explorations [20]. In terms of epistemic values, this implies a reevaluation of priorities, balancing the pursuit of scale with commitments to explanatory power and reproducibility, ensuring that AI advancements contribute to sustainable progress in materials science [21].

The framework also offers analytical implications for handling uncertainty in scaled AI systems. Large datasets enable probabilistic modeling, but interaction with bias can distort uncertainty estimates, leading to unreliable confidence intervals in property predictions [22]. Ethical dimensions highlight the risks posed by such uncertainties in critical applications, such as structural materials for infrastructure, where overly optimistic outputs could have real-world consequences [14]. Steering logics suggest integrative strategies, such as ensemble methods that aggregate scaled models with physics-based constraints, to enhance reliability [15]. Overall, these implications underscore the contingent nature of scale, positioning it as a tool that, when navigated through interpretive nuance, can advance materials informatics without sacrificing epistemic integrity [24].

Results and Discussion

The conceptual framework presented herein prompts a discussion of broader conceptual interpretations in materials AI, particularly how scale interfaces with epistemic and ethical reasoning to influence the field's trajectory [2]. Interaction dynamics reveal that while scale facilitates rapid pattern identification, it can perpetuate homogeneity, as models trained on expansive but biased datasets tend to reinforce existing knowledge structures, potentially stifling innovation in underrepresented material domains [3]. This dynamic invites consideration of feedback structures, where scaled AI outputs feed back into data generation, creating cycles that prioritize commonality over novelty [5]. Systems-level insights suggest that this could limit the diversity of material explorations, emphasizing the need for steering logics that incorporate deliberate diversity interventions to break such cycles [9].

Ethical reasoning further enriches this discussion by highlighting trade-offs between scale-driven efficiencies and value alignment in science [10]. For example, prioritizing high-throughput predictions may align with values of productivity but conflict with equity, as biases amplified by scale could disadvantage certain material classes or applications, such as those in sustainable technologies [14]. Conceptual interpretations frame scale as a relational construct whose benefits emerge from its interactions with other factors, such as data quality and model interpretability [5]. This perspective encourages integrative approaches that balance scale with qualitative assessments, fostering frameworks where AI serves as a complement to human insight rather than a replacement [15].

Moreover, the framework's analytical implications extend to the epistemic foundations of materials informatics, questioning how scale affects the construction of knowledge [10]. Large-scale models offer unprecedented breadth, but their black-box nature can dilute explanatory depth, raising concerns about the verifiability of discoveries [20]. Feedback structures imply that, without safeguards, scale can lead to epistemic dilution, in which quantitative abundance overshadows the need for contextual understanding [21]. Ethical dimensions compound this, as unchecked scale might implicitly embed societal values, influencing which materials are pursued and for what purposes [22]. Steering logics propose calibrated scaling, such as adaptive learning techniques that adjust for bias in real-time, to maintain epistemic robustness [24].

In discussing these elements, the framework underscores the importance of interdisciplinary integration, where materials science converges with informatics and ethics to navigate the complexities of scale [2]. Interaction dynamics with external factors, such as computational resources, further imply that the scale's value is contingent on accessibility, advocating open repositories to democratize AI benefits [3]. Ultimately, this discussion reframes scale not as an end in itself but as a component within larger systems of interpretation, urging a shift toward balanced, value-aware practices in materials AI [5].

Conclusion

This manuscript's conceptual framework reinterprets scale in materials AI as a multifaceted element embedded in interaction dynamics and trade-offs that shape epistemic and ethical landscapes. Through analytical implications and systems-level insights, it highlights how unchecked scaling can amplify biases and homogenize explorations, while calibrated steering logics offer pathways for integrative progress. Ethical reasoning emphasizes the need to align scale with values of equity and sustainability, ensuring that AI advancements in materials science contribute to diverse and responsible innovations. By framing scale as contingent rather than absolute, the framework advocates for nuanced approaches that preserve interpretive depth amid quantitative expansion, fostering a more equitable and robust field.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 16 Apr 2025 Revised: 11 Jun 2025 Accepted: 09 Jul 2025
Published online: 18 January 2026

References

Zhang C, Cheng J, Chen Y, Chan MKY, Cai Q, Carvalho RP, et al. 2023 roadmap on molecular modelling of electrochemical energy materials. J Phys Energy. 2023;5(4):041501.

Herowati W, Akrom M, Hidayat NN, Sutojo T. Investigation of corrosion inhibition efficiency of pyridine-quinoline compounds through machine learning. J Multiscale Mater Inform. 2024;1(1):1-9. https://doi.org/10.62411/jimat.v1i1.10448.

Bayan FM, Fayyad M. The ethics of ai: Navigating the moral dilemmas of artificial intelligence. Arab J Sci Publ. 2024;2663:5798.

Ferrara E. Fairness and bias in artificial intelligence: A brief survey of sources, impacts, and mitigation strategies. Sci. 2024;6(1):3.

Chen F, Wang L, Hong J, Jiang J, Zhou L. Unmasking bias in artificial intelligence: A systematic review of bias detection and mitigation strategies in electronic health record-based models. J Am Med Inform Assoc. 2024;31(5):1172-83.

Panayides AS, Amini A, Filipovic ND, Sharma A, Tsaftaris SA, Young A, et al. AI in medical imaging informatics: Current challenges and future directions. IEEE J Biomed Health Inform. 2020;24(7):1837-57. https://doi.org/10.1109/JBHI.2020.2991043.

Gong S. The significance of materials informatics on material science. Appl Comput Eng. 2024;58:208-14.

Wang Z, Chen A, Tao K, Han Y, Li J. Matgpt: A vane of materials informatics from past, present, to future. Adv Mater. 2024;36(18):2306733.

Mannodi-Kanakkithodi A, McDannald A, Sun S, Desai S, Brown KA, Kusne AG. A framework for materials informatics education through workshops. MRS Bull. 2023;48(6):631-42.

Hu J, Stefanov S, Song Y, Omee SS, Louis SY, Siriwardane EMD, et al. Materialsatlas.org: a materials informatics web app platform for materials discovery and survey of state-of-the-art. npj Comput Mater. 2022;8(1):72.

Gabriel I, Ghazavi V. The challenge of value alignment. In: The oxford handbook of digital ethics; 2022. pp. 336-355.

Tırpan EC. The ethical issues in generative artificial intelligence: A systematic review. Bus Manag Stud Int J. 2024;12(4):729-47.

Esteves RM. Applications, challenges, and ethical implications of generative AI: a systematic review [master's thesis]. Lisboa: NOVA Information Management School, Universidade Nova de Lisboa; 2024.

Lohyna D, Pico Oristrell A. Towards bias-free ai-supported decision-making: Exploring organisational and technology characteristics for algorithmic bias mitigation. Lund: Department of informatics, Lund school of economics and management, Lund university; 2023.

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

Hofmann B. Biases in AI: Acknowledging and addressing the inevitable ethical issues. Front Digit Health. 2025;7:1477585.

DP, LW, JHS. Reproducibility in materials informatics: Lessons from 'a general-purpose machine learning framework for predicting properties of inorganic materials'. Digit Discov. 2024;3(2):281-6.

Liu Y, Yang Z, Zou X, Ma S, Liu D, Avdeev M, et al. Data quantity governance for machine learning in materials science. Natl Sci Rev. 2023;10(7):nwad125.

Li C, Zheng K. Methods, progresses, and opportunities of materials informatics. InfoMat. 2023;5(8):e12425.

Hwang A, Badreddine S, Gifford F, Besold TR. Recipe 2.0: Information presentation for ai-supported culinary idea generation. In: Proceedings of the 14th international conference on computational creativity; 2023.

Artificial Intelligence Research and Development Interagency Working Group. 2020-2024 progress report: Advancing trustworthy artificial intelligence research and development. Natl Sci Technol Counc. 2024.

Maqsood A, Chen C, Jacobsson TJ. The future of material scientists in an age of artificial intelligence. Adv Sci. 2024;11(19):2401401.

Mosqueira-Rey E, Pereira EH, Alonso-Ríos D, Bobes-Bascarán J. A classification and review of tools for developing and interacting with machine learning systems. In: Proceedings of the 37th acm/sigapp symposium on applied computing; 2022:1092-101.

Mammas CS, Mamma AS, Papaxoinis G, Georgiou I. Remote AI supported e-multidisciplinary oncology conference in breast cancer as a technology and method to optimize outcomes in the peripheries. In: Informatics and technology in clinical care and public health; 2022. pp. 309-312.