

ORIGINAL RESEARCH

Open access

Data Is Not Neutral: A Conceptual Framework for Value-Laden Measurement Choices in Materials Informatics

Sanjay Kulkarni¹, Meenal Joshi^{1*}, Rohan Patil²

Abstract

Materials informatics has become a central paradigm in materials science, leveraging machine learning and large-scale datasets to accelerate property prediction, discovery, and design. However, prevailing approaches often treat data as a neutral substrate for modeling, obscuring the value-laden processes through which data is generated. Measurement choices—what properties to quantify, which materials to prioritize, and which experimental or computational protocols to employ—are inherently shaped by epistemic commitments, practical constraints, and broader societal priorities. These choices embed values into data infrastructures, systematically influencing which material phenomena become visible and which remain obscured in downstream models. This manuscript advances a conceptual framework that interprets measurement choices as value-mediated interfaces linking scientific priorities to data constitution and modeling feedback in materials informatics. The framework elucidates how value horizons, choice architectures, data formation processes, and modeling circuits interact to produce steering logics, trade-offs, and path-dependent dynamics. By reframing data bias as a constitutive outcome of value-conditioned measurement rather than a purely technical artifact, the framework reveals characteristic failure modes—including value lock-in, patterned absences, and self-reinforcing feedback—that constrain epistemic exploration. Integrating insights from materials informatics, data bias studies, and philosophical analyses of scientific practice, the framework provides a diagnostic lens for understanding the non-neutrality of data in iterative AI-driven workflows. Rather than prescribing methodological interventions, it foregrounds the epistemic consequences of measurement decisions, inviting greater reflexivity in shaping data landscapes over time. This perspective repositions materials informatics as an evolving epistemic system whose possibilities and limits are co-produced by values, measurements, and models.

Keywords Materials informatics, Data bias, Value-laden measurement, Epistemic values, Measurement choices, Machine learning in materials science

*Correspondence:

Meenal Joshi
meenal.joshi@gmail.com

¹ Department of Materials Informatics and Machine Learning, Faculty of Engineering, Savitribai Phule Pune University, Pune, India

² Department of AI Materials Simulation, Faculty of Engineering, IIT Bombay, Mumbai, India

Introduction

The integration of machine learning and data-intensive approaches has fundamentally altered the landscape of materials science. Once dominated by trial-and-error experimentation and physics-based simulation, the field now routinely employs statistical models to predict properties, screen candidates, and guide synthesis [1, 2]. This shift, often described as the realization of the “fourth paradigm” of science in materials contexts, relies on the availability of large, structured datasets drawn from experimental repositories, computational databases, and literature mining [3, 4]. Materials informatics promises to compress discovery timelines, optimize resource allocation, and address pressing challenges in energy storage, catalysis, and structural performance [5, 6].

Yet this optimism rests on an implicit premise: that the data underpinning these models is neutral with respect to the questions asked and the conclusions reached. In practice, data in materials science is never a passive record of reality. It emerges from deliberate measurement acts—decisions about which properties to characterize, which materials classes to sample, what resolution or accuracy to pursue, and which phenomena to foreground or background. These decisions are constrained by finite resources, instrumentation limits, and research priorities, but they are also directed by values [7, 8]. Epistemic values, such as preferences for generalizability over precision or robustness over novelty, coexist with non-epistemic considerations, including funding agency mandates, industrial relevance, and societal goals such as decarbonization or the substitution of critical materials [9, 10].

The value-laden character of measurement is not unique to materials science; it reflects broader features of scientific practice, in which underdetermination by evidence allows non-epistemic factors to shape methodological choices [11, 12]. In materials informatics, however, the stakes are amplified by the iterative, feedback-rich nature of the workflow. Models trained on value-shaped data generate predictions that influence subsequent experiments, closing loops that can reinforce or challenge initial priorities [13, 14]. If measurement choices systematically exclude certain material families, property regimes, or failure modes, the resulting datasets may narrow the field's epistemic horizon, embedding biases that propagate through generations of models [15, 16].

Recent literature has begun to grapple with these issues through the lens of data quality, bias mitigation, and best practices [17, 18]. Studies have documented how imbalances in training data—arising from historical focus on well-studied compounds or easily synthesizable systems—can skew predictions toward familiar chemistries [3, 19]. Others have emphasized curation standards to improve interoperability and reliability [4, 6]. Yet these discussions tend to treat bias as a technical artifact to be corrected rather than a symptom of deeper value commitments. Missing is a conceptual apparatus that interprets measurement choices as sites where values actively steer the constitution of data and, by extension, the trajectories of knowledge production.

This manuscript addresses that gap by proposing an original conceptual framework for understanding value-laden measurement choices in materials informatics. The framework foregrounds the systemic interactions among value sources, choice spaces, data-generation processes, and modeling feedback, without reducing the analysis to predictive claims or empirical validation. It draws on interpretive synthesis of recent literature to illuminate trade-offs, steering dynamics, and epistemic consequences, offering informaticians a reflective tool for navigating the non-neutrality of their data foundations. By reframing data as value-infused rather than value-free, the framework invites greater awareness of how measurement decisions shape not only what is known but also what can be known in the field.

Theoretical Background and Literature Synthesis

The emergence and scope of materials informatics

Since approximately 2020, materials informatics has transitioned from a collection of exploratory machine-learning applications into a central methodological paradigm for materials discovery and design. Early efforts in the field relied predominantly on hand-crafted descriptors, regression models, and high-throughput screening pipelines to map structure–property relationships within relatively constrained chemical spaces [1, 4]. While these approaches demonstrated proof-of-concept success, they were limited by representational rigidity and dependence on expert-defined features.

Recent advances have substantially broadened this scope. Graph neural networks, equivariant architectures, generative models, and foundation-style approaches now enable direct learning from atomic structures, microstructures, and multiscale representations without explicit descriptor engineering [7, 8]. These developments have expanded materials informatics from predictive modeling toward inverse design, hypothesis generation, and autonomous experimentation. Community platforms such as JARVIS and MaterialsAtlas exemplify this maturation by integrating heterogeneous

datasets, standardized features, benchmarking protocols, and open evaluation pipelines within shared infrastructures [4, 16].

Contemporary reviews consistently highlight a structural shift from isolated supervised models trained on narrow datasets toward large-scale, multi-fidelity workflows that integrate experimental measurements, first-principles simulations, and literature-derived knowledge [2, 6]. This evolution has enabled demonstrable advances across application domains, including battery materials, thermoelectrics, and high-entropy alloys [5, 6]. At the same time, it has rendered materials informatics increasingly dependent on the composition, coverage, and epistemic assumptions embedded in the underlying data ecosystem.

Data acquisition and the role of measurement choices

The data foundations of materials informatics are intrinsically heterogeneous, drawing from experimental characterization, density functional theory calculations, and curated community databases [3, 4]. Among these sources, experimental data retain a privileged epistemic status as ground truth, yet they are also the most constrained by cost, accessibility, and logistical feasibility. Measurement is therefore not an exhaustive sampling of material reality but a selective process shaped by the availability of instrumentation, constraints on sample preparation, and targeted property priorities.

Choices between synchrotron-based techniques and laboratory-scale measurements, between thin-film and bulk samples, or between electronic, mechanical, and thermal properties directly influence what is included in the collective data record [1]. Recent discussions emphasize that such choices are rarely neutral; they are guided by strategic considerations, including funding priorities, anticipated application relevance, and institutional expertise [9, 17]. As a consequence, materials datasets exhibit persistent structural asymmetries—well-studied oxides and crystalline metals are richly represented, while amorphous systems, metastable phases, and complex composites remain sparsely characterized [19].

These asymmetries are not merely incidental artifacts but cumulative outcomes of repeated measurement decisions over time. As datasets scale, early selection effects become entrenched, shaping both the feasible learning tasks and the apparent regularities models discover. Measurement choices thus function as implicit filters that pre-structure the informational landscape on which materials informatics operates.

Bias and data quality in machine learning workflows

An expanding body of literature has examined how such structural imbalances affect machine-learning performance in materials contexts [3, 15, 16]. Imbalanced sampling has been shown to impair generalization, particularly in extrapolative regimes involving underrepresented chemistries, compositions, or structural motifs. In response, the field has developed a range of technical strategies, including active learning, uncertainty-driven sampling, data augmentation, and redundancy exploitation, to improve efficiency and coverage.

Best-practice guidelines increasingly emphasize rigorous data curation, metadata standardization, and uncertainty quantification as prerequisites for reliable model development [1, 18]. These advances have undoubtedly improved technical robustness. However, they largely conceptualize bias as a statistical or algorithmic deficiency to be corrected post hoc. The upstream conditions that generate imbalance—namely, the selective priorities embedded in measurement and data collection—remain comparatively under-theorized.

As a result, existing approaches often mitigate symptoms without interrogating causes. Bias is treated as noise or sparsity rather than as an expression of structured preferences and constraints that shape the material knowledge base itself.

Philosophical underpinnings: Values in measurement and scientific practice

Philosophical analyses of scientific practice have long recognized that measurement is neither purely objective nor value-free [11]. The underdetermination of theory by data ensures that multiple measurement protocols can be empirically adequate, leaving room for epistemic values—such as simplicity, precision, or robustness—and non-epistemic values—such as economic relevance or societal impact—to influence methodological choices [12].

In data-intensive sciences, these value influences extend beyond individual experiments to the aggregation, classification, and prioritization of data at scale. Although explicit engagement with these issues remains limited in materials informatics, parallel discussions in fields such as climate modeling, biomedical AI, and computational chemistry demonstrate how value commitments systematically shape data landscapes and downstream inferences.

This interpretive tradition reframes measurement not as a neutral conduit between material reality and data, but as a steering mechanism that channels attention, resources, and epistemic authority toward particular regions of the problem space. Such a perspective is especially salient for materials informatics, where automated learning systems amplify the consequences of early selection decisions.

Integration and gaps in current understanding

Taken together, the literature reveals a persistent tension at the heart of materials informatics. The field depends critically on large, high-quality, and representative datasets, yet the processes that generate these datasets are shaped by selective, value-influenced measurement choices [6]. While technical methods for addressing imbalance and uncertainty continue to advance, a conceptual account of how values operate as constitutive elements of data generation remains largely absent.

No existing framework systematically explains how value systems interact with measurement choice spaces to structure data availability, learning trajectories, and model behavior over time. In the absence of such integration, biases risk being continuously re-encoded as apparent regularities, reinforcing path dependencies in materials discovery.

This manuscript addresses this gap by introducing a novel conceptual lens that interprets measurement choices as dynamic steering mechanisms governed by trade-offs among epistemic, practical, and normative values. By foregrounding these dynamics, the framework provides a basis for re-thinking data quality, bias, and model performance as emergent properties of value-laden scientific systems rather than purely technical artifacts.

Proposed conceptual framework

The proposed framework conceptualizes measurement choices in materials informatics as value-mediated interfaces that not only connect scientific priorities to data constitution but also systematically generate epistemic asymmetries, blind spots, and path-dependent dynamics within materials AI systems. Rather than treating measurement as a neutral precursor to modeling, the framework positions it as an active site of epistemic steering through which value commitments become embedded in data infrastructures over time. The framework comprises four interconnected domains: value horizons, measurement choice architecture, data constitution processes, and modeling feedback circuits. These domains are linked through bidirectional interactions that produce steering logics, epistemic trade-offs, and reflexive loops that shape both data landscapes and learning trajectories. The four domains of the framework and their epistemic roles are summarized in Table 1.

Table 1. Conceptual domains and epistemic roles in the proposed framework

Framework domain	Core function	Epistemic role	Typical consequences
Value horizons	Orient research priorities through epistemic and non-epistemic commitments	Define what counts as relevant, credible, or desirable knowledge	Privileging certain properties, materials, or applications

Measurement choice architecture	Structures feasible actions for data acquisition	Filters material reality into observable variables	Systematic inclusion and exclusion of phenomena
Data constitution processes	Transform measurements into usable datasets	Stabilize selective representations of material systems	Historically contingent data landscapes
Modeling feedback circuits	Feed model outputs back into research priorities	Reinforce or challenge prior value commitments	Path-dependence, reinforcement, or reorientation

Value horizons encompass the epistemic commitments (e.g., preferences for predictive accuracy, interpretability, or generalizability) and non-epistemic commitments (e.g., alignment with sustainability goals, industrial scalability, or critical materials security) that orient research programs. These horizons are neither fixed nor uniform; they evolve through disciplinary norms, funding environments, institutional incentives, and broader societal pressures. As such, value horizons function as dynamic orientation fields that condition which material questions are deemed salient and which forms of evidence are prioritized.

The measurement choice architecture represents the structured space of possible actions within these horizons—what properties to target, which material subspaces to sample, which protocols to adopt, and what fidelity levels to pursue. This space is constrained by instrumental availability, temporal limitations, and economic feasibility, yet it is navigated through value-driven selection rules rather than exhaustive optimization. Choices made at this level filter material phenomena into data, rendering some aspects of material reality salient while systematically occluding others. Importantly, these filters do not operate randomly but reflect recurring value alignments that accumulate across iterations.

Data constitution processes transform selected measurements into structured, usable records. Beyond raw acquisition, this includes preprocessing, annotation, standardization, and integration across sources, each step carrying forward value imprints inherited from prior decisions. The resulting datasets, therefore, embody a selective and historically contingent view of material reality, shaped not only by technical constraints but by the cumulative effects of value-conditioned choice architectures.

Modeling feedback circuits completes the system. Trained models generate predictions, uncertainty estimates, and inferred patterns that feed back into value horizons—either reinforcing existing priorities by demonstrating success in favored regimes or, less frequently, prompting reorientation by exposing persistent gaps, errors, or uncertainties in underrepresented regions. These feedback circuits introduce temporal dynamics into the framework: past value-laden choices condition future ones through the stabilization of particular data landscapes and modeling competencies.

Measurement choices can be expressed as a value-conditioned selection operator that maps the material world into the observed dataset used for learning. This makes non-neutrality explicit: values do not merely “bias” modeling after the fact, they parameterize what gets measured, thereby shaping the training distribution and the downstream prediction landscape. The iterative nature of materials informatics can then be captured as a feedback recursion in which model outputs update priorities and thus reshape subsequent measurement.

$$\begin{aligned} & [D_{\{t+1\}} \\ &= \{C\} \vdash (\{M\} \vdash (\Omega; \cdot, V_t, A_t) \vdash), \quad V_{\{t+1\}} \quad ; (1) \\ &= \{U\} \vdash (V_t; \cdot f_{t(D_t)} \vdash) \end{aligned}$$

The framework foregrounds several interaction dynamics that are central to understanding epistemic behavior in materials informatics. Steering logics describe how value horizons systematically channel measurement choices toward particular subspaces, producing path-dependent data accumulation. Trade-offs emerge as tensions between competing epistemic goods—such as breadth versus depth, accuracy versus coverage, or novelty versus reliability—that cannot be resolved through purely technical means but require implicit or explicit value judgments. Feedback structures enable

reflexivity, allowing model outcomes to either entrench existing commitments or destabilize them, thereby creating the possibility of adaptive reorientation or epistemic lock-in.

The framework is not intended as a prescriptive guide for measurement design or model selection. Rather, it functions as a diagnostic lens for interpreting how value commitments become embedded in data infrastructures and propagate through learning workflows over time. By emphasizing interaction dynamics rather than isolated decisions, the framework enables systematic identification of epistemic vulnerabilities and structural blind spots that remain invisible under purely statistical or performance-oriented accounts. As shown in Figure 1, the framework depicts AI-driven research as a value-mediated cyclic system linking value priorities, measurement choices, data formation, and modeling feedback, highlighting how upstream commitments propagate through iterative refinement and reinforcement loops.




Figure 1. Directed cyclic graph illustrating the value-mediated architecture of AI-driven research acceleration.

Revealed failure modes and epistemic reorientation

Epistemic reorientation

The framework proposed here effects an epistemic reorientation in how materials informatics conceptualizes data, bias, and model performance. Rather than treating measurement as a preparatory or logistical precursor to modeling, the framework relocates it to the core of knowledge formation. Measurement choices are recast as value-mediated interfaces that actively structure what becomes observable, learnable, and generalizable within materials AI systems. This shift moves analysis away from downstream artifacts—such as biased datasets or underperforming models—toward the upstream decision architectures that condition the very space of possible inferences. In doing so, the framework alters the

explanatory target: from correcting errors within data to understanding how data landscapes themselves are dynamically constituted. Key failure modes revealed by the framework are summarized in [Table 2](#).

Table 2. Failure modes revealed by value-laden measurement choices

Failure mode	Structural origin	Why is it often misdiagnosed	Epistemic consequence
Value lock-in	Repeated alignment of measurement choices with dominant value horizons	Interpreted as model confidence or robustness	Narrow exploration and overconfidence in favored regimes
Patterned absences	Systematic exclusion at the measurement stage	Treated as incidental data sparsity	Persistent blind spots in materials space
Illusory trade-off resolution	Implicit value judgments framed as technical optimization	Assumed solvable via algorithms or tuning	Silent privileging of certain epistemic goods
Epistemic path-dependence	Reinforcing feedback between data and models	Viewed as a natural research convergence	Reduced adaptability and loss of alternative trajectories

Failure modes revealed by the framework

Viewed through this lens, several systemic failure modes become visible that remain obscured in conventional materials informatics narratives.

First, value lock-in masquerading as model confidence

When value horizons consistently privilege particular properties, material classes, or fidelity regimes, the resulting measurement choices yield dense data coverage across aligned subspaces. Models trained on such landscapes often exhibit high apparent confidence and stability, which can be misinterpreted as epistemic robustness. The framework reveals that this confidence may instead reflect value lock-in: the reinforcement of already favored regimes through cumulative data accumulation. Underrepresented regions remain epistemically opaque, not due to intrinsic complexity, but because steering logics have persistently directed attention elsewhere.

Second, patterned absences are misread as data sparsity

Standard accounts frame gaps in materials datasets as accidental sparsity or technical limitations. The proposed framework reframes these gaps as patterned absences produced by systematic selection. Phenomena outside dominant value horizons—such as metastable phases, non-equilibrium states, or materials with marginal industrial relevance—are filtered out repeatedly at the measurement stage. Over time, these absences harden into structural blind spots that constrain what models can plausibly predict or even recognize as relevant targets.

Third, illusory trade-off resolution through technical optimization

Measurement decisions frequently involve tensions between competing epistemic goods, such as precision versus coverage or interpretability versus representational richness. In practice, these trade-offs are often treated as technical optimization problems. The framework exposes this as a category error: such tensions are resolved implicitly through value judgments rather than algorithmic tuning. When these judgments remain unarticulated, downstream modeling choices inherit unresolved trade-offs, leading to systems that appear optimized while silently privileging particular epistemic ends.

Fourth, self-reinforcing feedback and epistemic path-dependence

Modeling feedback circuits amplifies early value-laden choices by directing future measurement toward regions of prior success. This reinforcement stabilizes research trajectories but also entrenches path-dependence, narrowing exploratory diversity over time. While feedback can occasionally surface anomalies that prompt reorientation, the dominant dynamic is cumulative reinforcement, in which past value commitments increasingly delimit future discovery spaces.

Reframing data neutrality

Taken together, these failure modes support a broader reframing of data neutrality in materials informatics. The framework demonstrates that non-neutrality is not an incidental defect introduced by imperfect sampling or insufficient curation; it is a constitutive feature of data-intensive scientific practice under conditions of underdetermination. Because measurement choices inevitably filter material reality through evolving value horizons, neutrality cannot be recovered solely through bias elimination. Instead, neutrality must be reconceptualized as a reflective achievement: the outcome of sustained awareness of steering logics, trade-offs, and feedback dynamics.

This reframing shifts the normative center of gravity from the impossible task of producing value-free datasets toward the more tractable goal of making value influences explicit, examinable, and contestable. By rendering the dynamics of data constitution visible, the framework opens conceptual space for more pluralistic and resilient forms of materials knowledge—without prescribing specific methodologies or policy interventions. In this sense, it repositions materials informatics not merely as a technical enterprise, but as an evolving epistemic system whose limits and possibilities are co-produced by values, measurements, and models over time.

Results and Discussion

The conceptual framework proposed here invites materials informatics to move beyond technical discussions of data quality and bias mitigation toward a deeper engagement with the value dimensions of scientific practice. While existing literature has productively addressed statistical imbalances through active learning, augmentation, and curation standards [17, 18], these approaches typically operate downstream of measurement decisions and treat value influences as externalities to be corrected. In contrast, the framework positions values as intrinsic to the constitution of data from the outset, shaping which aspects of material reality enter the evidentiary record.

This perspective aligns with broader philosophical reflections on underdetermination and inductive risk in data-intensive sciences [11, 12]. Yet, it offers a field-specific interpretation tailored to the iterative, feedback-driven character of materials workflows. The emphasis on steering logics and reflexive circuits extends prior analyses of path-dependence in computational materials databases [3, 19] by tracing how value horizons interact with choice architectures to produce cumulative asymmetries. It also complements emerging work on uncertainty quantification and epistemic awareness in machine learning [10, 14], interpreting such awareness not merely as a modeling desideratum but as a response to value-mediated selectivity in data generation.

Several challenges arise when operationalizing this framework. First, value horizons are rarely articulated explicitly within research programs; they often remain tacit, embedded in funding calls, publication norms, or community benchmarks. Rendering them visible requires reflective practices—perhaps through structured value elicitation in project design or interdisciplinary dialogue—that may initially slow progress but ultimately enhance robustness. Second, the framework's focus on trade-offs highlights tensions that resist algorithmic resolution. Decisions about breadth versus depth or novelty versus reliability demand normative deliberation, raising questions about whose values should guide such choices in publicly funded or industrially oriented work.

Despite these challenges, the framework provides a constructive lens for navigating the field's maturation. As materials informatics scales toward autonomous discovery and multi-fidelity integration [2, 6, 7], the risks of unexamined value steering intensify: models may optimize efficiently within narrow value-constrained subspaces while missing broader

transformative opportunities. By foregrounding interaction dynamics and feedback structures, the framework equips practitioners to anticipate and mitigate such risks, promoting more intentional governance of measurement choices.

Ultimately, acknowledging data's non-neutrality does not undermine the legitimacy of materials informatics; rather, it strengthens it by aligning methodological choices more closely with explicitly examined priorities. This alignment can foster greater equity in knowledge production—ensuring that underrepresented material classes or societal needs receive due consideration—and enhance the field's capacity to address complex, multi-objective challenges in sustainability, energy, and beyond.

Conclusion

Materials informatics stands at a pivotal juncture where accelerating computational capabilities meet growing societal demands for responsible innovation. The conceptual framework presented here asserts that data in this domain are never neutral: they emerge from value-laden measurement choices that shape sampling, constitute datasets, and feed back into future priorities. By illuminating the interplay among value horizons, choice architectures, data processes, and modeling circuits, the framework offers an integrative interpretive structure for understanding these dynamics.

This perspective encourages a shift from viewing data bias primarily as a statistical artifact toward recognizing it as a symptom of deeper value commitments. It highlights steering logics that produce patterned absences, trade-offs that require normative resolution, and reflexive loops that enable both reinforcement and adaptive change. These insights do not prescribe specific protocols but provide a reflective apparatus for interrogating how measurement decisions shape epistemic possibilities.

Embracing this non-neutrality need not paralyze progress; instead, it invites deliberate stewardship. Materials informaticians can cultivate greater awareness of value influences, periodically reassess priorities in light of model insights, and design workflows that expose rather than obscure trade-offs. Such practices promise more transparent, contestable, and ultimately more robust knowledge generation—better equipped to serve diverse scientific and societal ends in an era of urgent materials challenges.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Zivic F, Kaplarevic Malisic A, Grujovic N, Stojanovic B, Ivanovic M. Materials informatics: A review of ai and machine learning tools, platforms, data repositories, and applications to architected porous materials. *Mater Today Commun.* 2025;48:113525.
<https://doi.org/10.1016/j.mtcomm.2025.113525>.
- Persaud D, Ward L, Hattrick-Simpers J. Reproducibility in materials informatics: Lessons from 'a general-purpose machine learning framework for predicting properties of inorganic materials'. *Digit Discov.* 2024;3:281-6.
<https://doi.org/10.1039/D3DD00199G>.
- Katsura Y, Akiyama M, Morito H, Fujioka M, Sugahara T. Systematic searches for new inorganic materials assisted by materials informatics. *Sci Technol Adv Mater.* 2024;25:2428154.
<https://doi.org/10.1080/14686996.2024.2428154>.
- Hu J, Stefanov S, Song Y, Omeel SS, Louis S-Y, 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;8:65.
<https://doi.org/10.1038/s41524-022-00750-6>.
- Kusne AG, Yu H, Wu C, Zhang H, Hattrick-Simpers J, DeCost B, et al. On-the-fly closed-loop materials discovery via bayesian active learning. *Nat Commun.* 2020;11(1):5966.
- Wang B, Du Y. Accelerating materials discovery through active learning: Methods, challenges and opportunities. *Innov Inform.* 2025;1:100013.
<https://doi.org/10.59717/j.xinn-inform.2025.100013>.
- Wang X, Musielewicz J, Tran R, Ethirajan SK, Fu X, Mera H, et al. Generalization of graph-based active learning relaxation strategies across materials. *Mach Learn Sci Technol.* 2024;5:025018.
<https://doi.org/10.1088/2632-2153/ad37f0>.
- Zhou Q, Chen X, Wang J. Machine learning assisted material discovery: A small data approach. *Acc Mater Res.* 2025;6:685-94.
<https://doi.org/10.1021/accountsmr.1c00236>.
- Kumagai M, Ando Y, Tanaka A, Tsuda K, Katsura Y, Kurosaki K. Effects of data bias on machine-learning-based material discovery using experimental property data. *Sci Technol Adv Mater Methods.* 2022;2:302-9.
<https://doi.org/10.1080/27660400.2022.2109447>.
- Karande P, Gallagher B, Han TY-J. A strategic approach to machine learning for material science: How to tackle real-world challenges and avoid pitfalls. *Chem Mater.* 2022;34:7650-65.
<https://doi.org/10.1021/acs.chemmater.2c01333>.
- Elliott K. Values in science. *Elements in the philosophy of science.* 2022.
<https://doi.org/10.1017/9781009052597>.
- Elliott KC, Korf R. Values in science: What are values, anyway? *Eur J Philos Sci.* 2024;14:1-24.
<https://doi.org/10.1007/s13194-024-00615-3>.

Pyzer-Knapp EO, Manica M, Staar P, Morin L, Ruch P, Laino T, et al. Foundation models for materials discovery – current state and future directions. *npj Comput Mater.* 2025;11:61.
<https://doi.org/10.1038/s41524-025-01538-0>.

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

Wines D, Gurunathan R, Garrity KF, DeCost B, Biacchi AJ, Tavazza F, et al. Recent progress in the jarvis infrastructure for next-generation data-driven materials design. *Appl Phys Rev.* 2023;10(4).

Choudhary K, Wines D, Li K, Garrity KF, Gupta V, Romero AH, et al. JARVIS-leaderboard: a large scale benchmark of materials design methods. *npj Comput Mater.* 2024;10:93.
<https://doi.org/10.1038/s41524-024-01259-w>.

Hestroffer JM, Charpagne MA, Latypov MI, Beyerlein IJ. Graph neural networks for efficient learning of mechanical properties of polycrystals. *Comput Mater Sci.* 2023;217:111894.
<https://doi.org/10.1016/j.commatsci.2022.111894>.

Park CW, Wolverton C. Developing an improved crystal graph convolutional neural network framework for accelerated materials discovery. *Phys Rev Mater.* 2020;4(6):063801.

Shi X, Zhou L, Huang Y, Wu Y, Hong Z. A review on the applications of graph neural networks in materials science at the atomic scale. *Mater Genome Eng Adv.* 2024;2:e50.
<https://doi.org/10.1002/mgea.50>.