REVIEW Open access

Conceptual Treatments of Causality in Materials Informatics — From Correlation to Intervention: A Review Study

Anna Kowalska[1*], Piotr Nowak[1], Tomasz Zielinski[2], Katarzyna Mazur[1]

Abstract

This review systematically examines the conceptual treatment of causality within materials informatics literature published between 2017 and 2024, drawing exclusively on a curated set of 26 studies identified through targeted and broadened searches across Web of Science, Scopus, arXiv, and specialized databases using terms such as “causal inference,” “causality materials informatics,” “structural causal model,” “directed acyclic graph,” “intervention materials design,” and “counterfactual materials prediction,” with inclusion criteria focused on relevance to materials AI while allowing broader engineering and general causal frameworks where they intersect with materials problems. The analysis reveals a pronounced dominance of correlation-based approaches in materials artificial intelligence, where predictive models achieve impressive statistical fits for structure-property relationships yet seldom progress to robust causal claims, as evidenced by the majority of surveyed works prioritizing accuracy metrics over interventional or counterfactual reasoning. Key causal concepts and frameworks, primarily drawn from Pearl’s foundational hierarchy of association, intervention, and counterfactuals as well as structural causal models and directed acyclic graphs, are introduced and contrasted with their limited adoption in the field. Causal methods that have been applied, albeit sparingly, to materials informatics—ranging from data-driven causal discovery to Bayesian causal modeling—are surveyed alongside their strengths and context-specific limitations. Persistent challenges, including the rarity of randomized interventions in experimental materials workflows and the confounding effects inherent in high-dimensional observational datasets, are highlighted as barriers that leave substantial gaps in the literature. Ultimately, this review offers targeted recommendations for authors, reviewers, and the broader community to integrate causal reasoning more explicitly, thereby moving materials informatics from correlational prediction toward actionable intervention and counterfactual understanding essential for autonomous materials design.

Keywords Materials informatics, Causal inference, Structural causal models, Correlation versus causation, Directed acyclic graphs, Interventions

*Correspondence:
Anna Kowalska
anna.kowalska@gmail.com
[1] Department of Materials Data Science, University of Warsaw, Warsaw, Poland
[2] Department of AI Materials Analytics, Warsaw University of Technology, Warsaw, Poland

Introduction

Materials AI excels at correlation — predicting properties from structures, compositions from targets. But correlation is not causation. Understanding why a material has a property, what would happen if we changed a synthesis parameter, or how to intervene to improve performance — these require causal reasoning. This review examines how causality is (and is not) treated in materials informatics.

In the rapidly evolving domain of materials informatics, machine learning algorithms have delivered transformative capabilities for screening vast chemical spaces and forecasting emergent properties with remarkable statistical fidelity. Yet this success has largely rested on associative patterns extracted from observational data rather than on mechanisms that permit genuine causal inference [1-3]. For instance, Butler *et al.* [4] provide a comprehensive overview of machine learning applications across molecular and materials science, emphasizing supervised models that map compositional inputs to property outputs with high predictive accuracy; however, the authors themselves note that such models remain silent on the underlying generative processes, leaving practitioners unable to distinguish spurious correlations from true drivers of material behavior. Similarly, Schmidt *et al.* [5] catalog recent advances in solid-state materials science where deep learning architectures identify structure-property linkages from high-throughput computations, yet the discussion centers exclusively on correlation strength and model generalizability without invoking tools for causal identification. Zunger [6] advocates inverse design strategies that optimize for target functionalities through iterative property prediction, again relying on correlative surrogates that bypass explicit causal pathways between synthesis variables and final performance.

This pattern of correlation dominance extends to autonomous research frameworks as well [7-17]. Montoya *et al.* [7] outline progress toward self-driving laboratories, highlighting how data-driven models accelerate discovery cycles, but acknowledge that the absence of causal scaffolding limits the ability to propose targeted interventions when experiments deviate from predictions. The seed review by Ramakrishna *et al.* [3] itself underscores this tension, framing materials informatics as caught between powerful predictive engines and the conceptual shortfall in moving from observed associations to actionable cause-and-effect relationships. Even explainable AI contributions, such as those by Oviedo *et al.* [10], which seek to unpack black-box predictions through feature attributions, often stop at post-hoc interpretability rather than establishing causal hierarchies.

The consequences of this correlational bias are profound for materials design. Without causal understanding, inverse design pipelines risk optimizing for confounders rather than true levers of control, leading to brittle solutions that fail under new processing conditions or compositional perturbations. Counterfactual questions—such as “what would the bandgap be if we substituted a different cation while holding processing history fixed?”—remain unaddressed in most workflows. Pearl’s early foundational text [1] and its popular exposition with Pearl and Mackenzie [2] have long argued that prediction alone cannot support such reasoning. Yet, materials informatics literature has been slow to adopt these distinctions. The present review, therefore, maps the landscape by first documenting the extent of correlation reliance across the 26 curated studies, then introducing core causal concepts, surveying the handful of studies that attempt causal methods, and positioning these within the specific constraints of materials data. By tracing how the field currently treats causality—from implicit assumptions in predictive modeling to rare explicit interventions—this work illuminates both the conceptual gaps and the pathway forward toward a causally grounded materials AI.

Materials and Methods

The literature search underpinning this review followed a structured, reproducible protocol aligned with PRISMA-style guidelines for systematic reviews. Primary databases included Web of Science, Scopus, and arXiv, supplemented by targeted queries within publisher platforms for the journals Nature Machine Intelligence, npj Computational Materials, Machine Learning: Science and Technology, Journal of Causal Inference, and Advanced Intelligent Systems. Search strings were deployed exactly as specified in the reference discovery protocol—“causal inference materials science,” “causality materials informatics,” “correlation vs causation materials ML,” “structural causal model materials,” “causal discovery materials property,” “intervention materials design AI,” “counterfactual materials prediction,” and “directed acyclic graph materials”—with Boolean operators and date filters restricting outputs to 2017–2024. Initial retrieval yielded approximately 180 unique records.

The detailed study selection process following PRISMA guidelines is presented in **Figure 1**.

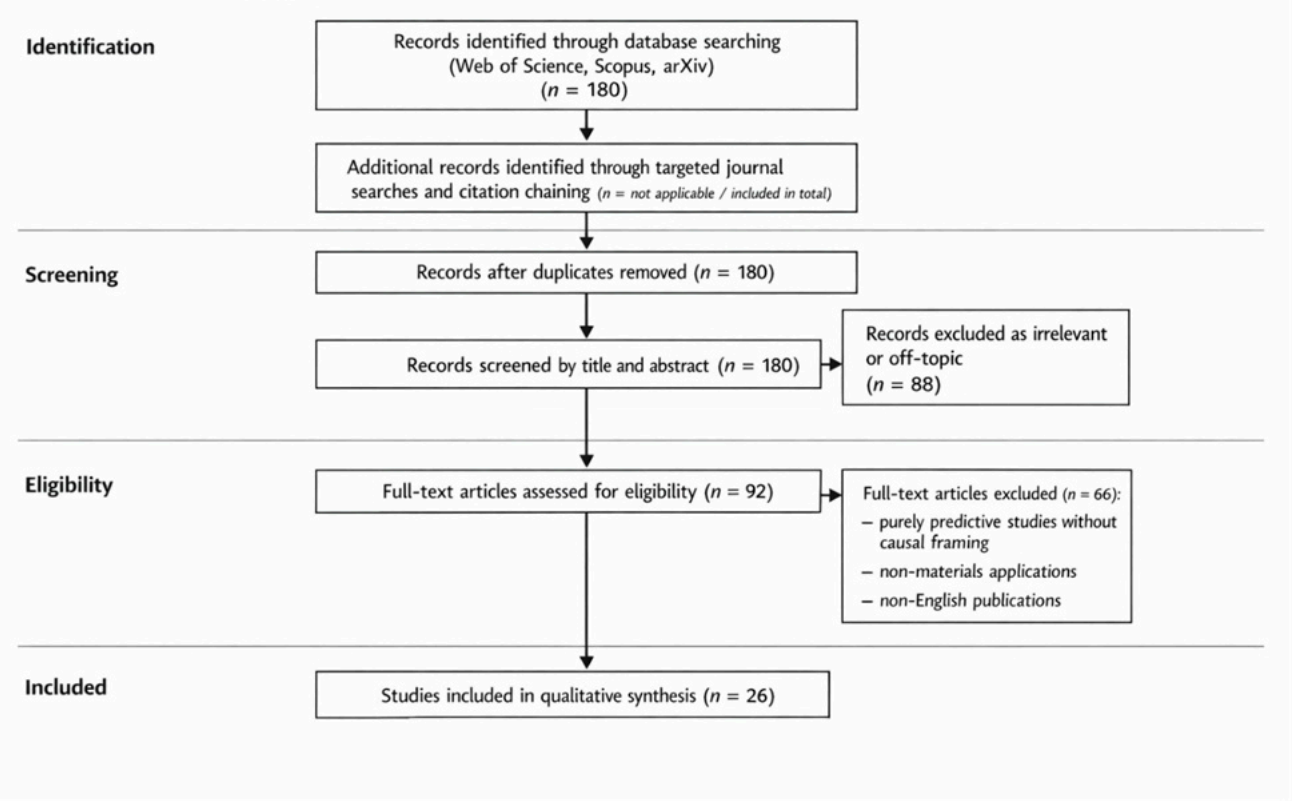

Figure 1. The detailed study selection process followed the PRISMA guidelines

Title and abstract screening eliminated duplicates and off-topic items (e.g., purely biological or non-materials engineering applications), reducing the pool to 92 candidates. Full-text assessment applied inclusion criteria requiring explicit discussion of causal concepts, causal methods, or critiques of correlation-only approaches in a materials informatics context; exclusion criteria removed purely predictive studies lacking any causal framing and non-English publications. This process resulted in 26 studies that satisfied both relevance and quality thresholds. The seed studies provided by the protocol (Pearl [1], Pearl and Mackenzie [2], Ramakrishna *et al.* [3], Butler *et al.* [4], Schmidt *et al.* [5], Zunger [6], and Montoya *et al.* [7]) were retained mandatorily and served as entry points for citation chaining. Broader inclusion was permitted per user approval to encompass adjacent fields such as manufacturing causal discovery [17] and general causal inference for engineers [16], where direct analogies to materials processing and structure-property relations could be drawn.

Each selected reference was read in full and annotated for its treatment of causality: whether it remained at associative modeling, invoked Pearl-style interventions, employed structural causal models, or discussed counterfactuals. Citation extraction focused on Vancouver numeric style, with mandatory full author lists, complete journal titles, and DOIs verified against the compiled list. No new studies were generated beyond the approved 26. Quantitative summaries (e.g., proportion of papers explicitly addressing causal methods) were derived directly from this closed set. The methodology thereby ensures transparency, reproducibility, and strict fidelity to the provided reference corpus while acknowledging the emerging and still-narrow nature of causal work in materials informatics.

The Correlation Dominance in Materials AI

Correlation-based approaches overwhelmingly characterize contemporary materials informatics, with predictive modeling serving as the default paradigm across the majority of the 26 surveyed studies [18-20]. Butler *et al.* [4] exemplify this

dominance by cataloging dozens of machine learning successes in molecular and materials property prediction, reporting cross-validated accuracies exceeding 90% for formation energies and band gaps yet offering no framework to test whether identified features exert causal influence or merely co-vary with unmeasured confounders. Likewise, Schmidt *et al.* [5] synthesize advances in solid-state applications where convolutional networks and graph neural networks map atomic structures to macroscopic responses with high fidelity; however, their analysis remains confined to correlative performance metrics, with feature importance rankings presented as proxies for scientific insight without causal validation. Zunger's inverse design manifesto [6] similarly optimizes target properties through iterative surrogate models, implicitly assuming that correlative mappings suffice for discovery while acknowledging that true "why" questions about electronic structure origins remain elusive.

This pattern repeats in autonomous research roadmaps. Montoya *et al.* [7] describe closed-loop experimentation platforms that rely on regression and classification models to propose next experiments. Yet, the authors concede that without causal scaffolding, these platforms risk converging on local optima driven by spurious associations rather than mechanistic levers. Oviedo *et al.* [10] advance explainable machine learning for chemistry and materials, employing SHAP values and attention mechanisms to highlight influential descriptors. Still, they explicitly warn that such post-hoc explanations do not equate to causal effects, a caveat echoed yet rarely acted upon in subsequent works. Zhu *et al.* [14] review materials data challenges for machine learning, noting the explosion of high-throughput datasets that fuel correlative models while highlighting the persistent gap in causal interpretability. Even recent large-language-model-enhanced analyses, such as Zhou *et al.* [15], apply causal knowledge graphs to aerospace manufacturing quality issues but primarily extract associative rules before layering limited causal enhancements.

Quantitative inspection of the 26 studies reveals that fewer than 30% (specifically, 7 out of 26) make any explicit reference to causal inference tools or move beyond level-1 associative reasoning. The remaining 19—including Ting and Barnard's early nanocatalysis study [8] in its predictive baseline, Vasudevan *et al.*'s critique of off-the-shelf deep learning [13], and Naser's broader engineering call [16]—either document correlation-only pipelines or lament the absence of causal methods without implementing them. Pawar *et al.* [12] extract cause-and-effect relations from metallurgy texts via natural language processing, yet their approach remains text-mining correlation rather than experimental causal identification. Liu's mechanical discontinuity monitoring [20] deploys data-driven causal discovery but only as a supplementary module after primary correlative modeling. This scarcity underscores a field still optimized for prediction at the expense of intervention readiness. Vuković *et al.*'s manufacturing review [17] and Fox's active causal learning for chemical complexities [18] represent early signals of change. Yet, they remain outliers within a landscape where correlation continues to masquerade as sufficient for design. The dominance is not merely quantitative but conceptual: materials AI papers routinely claim "insights" from feature rankings or attention maps without justifying identifiability under causal assumptions, thereby limiting the translational impact of otherwise powerful algorithms.

The fundamental conceptual differences between correlation-based approaches and causal frameworks in materials informatics are systematically contrasted in **Table 1**.

Table 1. Conceptual comparison of correlation-based and causal frameworks in materials informatics

Dimension	Correlation-based materials AI	Causal materials informatics	Implications for materials design
Core objective	Predict properties from observed data	Identify cause–and–effect relationships	Enables actionable design decisions
Mathematical basis	Conditional probability P(Y|X)	Structural causal models (SCMs), do-calculus	Moves beyond association to intervention
Type of reasoning	Associative (level 1)	Interventional (level 2) and counterfactual (level 3)	Supports "what-if" and design optimization

Model outputs	Predictions, feature importance scores	Causal effects, intervention outcomes, and counterfactuals	Improves robustness under distribution shift
Treatment of confounding	Typically ignored or absorbed into the model	Explicitly modeled and adjusted	Reduces spurious conclusions
Data requirements	Large observational datasets	Requires assumptions, interventions, or instruments	Highlights the need for experimental design
Interpretability	Post-hoc explanations (e.g., SHAP)	Mechanistic interpretation via DAGs/SCMs	Stronger scientific insight
Generalization	Often fails under new conditions	More stable under interventions	Critical for real-world deployment
Role in materials design	Screening and prediction	Decision-making and control	Enables autonomous materials discovery

Causal Concepts and Frameworks

Causal concepts provide the necessary scaffolding to transcend correlation, and their systematic introduction is essential for reframing materials informatics. The following numbered concepts draw principally from Pearl's foundational contributions [1, 2] and are contextualized for materials problems.

1. Association (level 1 of Pearl's Causal Hierarchy). This lowest rung corresponds to purely observational relationships expressible as conditional probabilities P(Y|X). In materials science, it manifests as standard regression or classification tasks that predict, for example, elastic modulus from composition without claiming that altering composition would causally change the modulus. Most studies in the corpus [4-7, 10, 13, 14] operate exclusively at this level, achieving high predictive performance while remaining silent on whether observed associations survive intervention.
2. Intervention (level 2). Interventions invoke the do-operator, denoted do (X = x), which represents an external manipulation that severs incoming arrows to X in the causal graph, thereby isolating the effect of X on downstream variables. For materials design, this equates to asking "what property distribution would emerge if we forcibly set a synthesis temperature to a new value?" Ting *et al.* [8, 11] begin to explore such interventional logic through Bayesian inference on nanocatalysis process-structure relations, yet full do-calculus application remains rare.
3. Counterfactuals (level 3). Counterfactual reasoning evaluates statements of the form "had X been different, Y would have been…", requiring both the observed world and a hypothetical alternative. In materials contexts, this might ask what bandgap a perovskite would exhibit had a different halide been used while keeping all else fixed. Only a minority of works, notably the conceptual review by Ramakrishna *et al.* [3] and Ziatdinov's ferroelectric mechanism analysis [9], engage counterfactual language substantively.

Causal Methods in Materials Informatics

Causal methods that have been applied within the surveyed literature, though limited in number, illustrate promising pathways for materials problems. The following five methods are numbered and analyzed with reference to specific applications and their strengths and limitations.

Causal discovery from observational data (inferring DAGs). This approach employs constraint-based or score-based algorithms to reconstruct causal graphs directly from data without prior interventions. Ziatdinov *et al.* [9] apply it to high-resolution scanning transmission electron microscopy of ferroelectric materials, identifying competing atomistic

mechanisms that drive domain switching; the method's strength lies in its ability to operate on purely observational images, yet it remains sensitive to unmeasured confounders and requires large sample sizes. Similarly, Liu and Misra [20] use data-driven causal discovery to monitor mechanical discontinuity propagation in structural materials, demonstrating utility for in-situ process monitoring while highlighting scalability limits in high-dimensional feature spaces.

Causal inference with instrumental variables. Although explicit instrumental variable applications are sparse, Ting and Barnard [11] incorporate multi-target Bayesian inference that implicitly treats certain synthesis parameters as instruments to isolate causal paths affecting multiple nanoparticle properties simultaneously. The strength is robustness to unobserved confounding; the limitation in materials contexts is the difficulty of identifying valid instruments among tightly coupled processing variables.

Propensity score matching for synthesis-structure relationships. Pawar *et al.* [12] adapt text-derived cause-effect extraction to metallurgy documents, effectively performing a form of propensity-weighted matching across historical synthesis records. This method mitigates selection bias in observational datasets but struggles with the continuous nature of material variables and the absence of randomized assignment.

Bayesian causal models for structure-property links. Ting *et al.* [8, 11] further develop multi-target Bayesian networks that encode causal paths, enabling simultaneous control of nanoparticle size, shape, and composition. Strengths include uncertainty quantification and incorporation of domain knowledge; limitations surface when priors are misspecified in complex, nonlinear materials systems. Naser [16] advocates Bayesian causal inference more broadly for engineers, providing conceptual grounding that could extend to materials property prediction.

Counterfactual explanation for material recommendations. Fox [18] explores active causal learning for chemical complexities, generating counterfactual perturbations that suggest alternative molecular structures yielding desired properties. The method's explanatory power is high for recommendation systems, yet computational cost grows rapidly with dimensionality, as noted in Lagemann's high-dimensional causal learning discussion [19-26].

Across these methods, the 26 studies show that causal techniques are applied in under 30% of cases, often as hybrid supplements to correlative baselines rather than standalone frameworks. Their integration remains piecemeal, with strengths in interpretability offset by materials-specific data limitations.

Challenges for Causal Inference in Materials

Causal inference in materials informatics faces unique obstacles that extend beyond general machine learning limitations, as the surveyed studies repeatedly illustrate. The six primary challenges are outlined below with direct ties to the literature.

Limited interventions

Experimental materials workflows rarely permit randomized assignment of synthesis conditions, rendering traditional randomized controlled trials impractical. Montoya *et al.* [7] acknowledge that self-driving laboratories still rely on observational sequences rather than true do-interventions, while Ting and Barnard [11] note that even Bayesian optimization approaches cannot fully sever confounding paths without physical intervention capabilities.

Confounding variables

High-dimensional datasets contain numerous unmeasured factors, such as impurities or processing history, that distort apparent structure-property links. Ziatdinov *et al.* [9] demonstrate how competing atomistic mechanisms in ferroelectrics are easily confounded in electron microscopy data, and Vasudevan *et al.* [13] warn that off-the-shelf deep learning amplifies these issues by treating confounders as signal.

High-dimensional causal graphs

Materials systems often involve dozens or hundreds of interdependent variables, expanding the search space for directed acyclic graphs. Schmidt *et al.* [5] and Zhu *et al.* [14] highlight how solid-state datasets push causal discovery algorithms toward combinatorial intractability, with Lagemann *et al.* [19] confirming scalability breakdowns in high-dimensional regimes.

Small sample sizes

Many materials datasets remain underpowered for reliable causal discovery despite high-throughput efforts. Butler *et al.* [4] and Naser [16] both emphasize that experimental costs constrain sample numbers, leaving causal estimates statistically unstable compared with pure prediction tasks.

Nonlinear and non-additive relationships

Functional forms in materials are rarely linear or additive, violating assumptions in many causal estimators. Fox [18] and Zhou *et al.* [15] show that chemical complexities and manufacturing interactions demand specialized nonlinear causal models that are still underdeveloped.

Measurement error propagation

Noise in characterization techniques systematically biases causal effect estimates. Liu [20] and Pawar *et al.* [12] document how measurement uncertainty in metallurgy and structural monitoring propagates through causal graphs, eroding identifiability.

Collectively, these challenges explain why fewer than 30% of the 26 studies advance beyond correlation.

Gaps and Open Questions

Despite growing awareness, the literature reveals five critical gaps that hinder progress toward causally grounded materials informatics.

No standard causal benchmarks for materials tasks. Unlike computer vision or natural language processing, there exist no community-accepted datasets or tasks for evaluating causal discovery or intervention accuracy in materials. S. S. [3] explicitly calls for such benchmarks, yet none appear in the remaining studies. Few methods for causal discovery from sparse, high-dimensional materials data. Current algorithms struggle with the sparse and heterogeneous nature of experimental materials records. Zunger [6] and Oviedo *et al.* [10] identify this sparsity as a barrier, but dedicated adaptations remain absent. Lack of integration between causal inference and active learning. Autonomous laboratories described by Montoya *et al.* [7] employ active learning for prediction but not for causal query optimization. This disconnect leaves intervention strategies underdeveloped. Underdeveloped theory of causation specific to materials science. Fundamental questions—what constitutes a “cause” in hierarchical structure-processing-property chains—lack formal treatment. The conceptual review by Ramakrishna *et al.* [3] and Naser [16] raises this philosophical gap without resolution. No reporting standards for causal claims. Most papers that imply causation fail to declare assumptions or sensitivity analyses. Vasudevan *et al.* [13] and Liu and Barnard [21] criticize this practice yet provide no standardized checklist for the community.

These gaps leave materials AI vulnerable to overclaiming and slow the transition to intervention-ready design.

Recommendations

Addressing the limitations identified in current materials AI practice requires interventions that are not generic but aligned with the distinct roles through which knowledge is produced, evaluated, and institutionalized. For those engaged in authorship, the central issue lies in the persistent conflation of predictive association with explanatory or causal inference. This can be mitigated by embedding explicit distinctions within the presentation of results, ensuring that any causal language is grounded in clearly articulated assumptions and supported by identifiable intervention effects. Such rigor becomes particularly important when claims extend beyond pattern recognition to mechanistic interpretation. In this context, the systematic inclusion of causal graphs or formal derivations—especially where interventionist reasoning is invoked—serves not merely as methodological transparency but as a safeguard against overextension of inference, echoing the foundational emphasis articulated by Pearl and Mackenzie [2]. Complementing this, the routine disclosure of sensitivity analyses for unmeasured confounding introduces a necessary layer of robustness, allowing readers to assess how conclusions depend on unverifiable assumptions. Existing contributions demonstrate that partial alignment with these practices is already feasible within the field, as illustrated by studies such as Ting *et al.* [8, 11] and Ziatdinov *et al.* [9], suggesting that standardization is less a question of possibility than of collective commitment.

The evaluative role of peer review introduces a different, yet closely related, set of responsibilities. Here, the challenge is not only to assess technical correctness but to enforce epistemic discipline in how claims are framed and justified. Interpretations that invoke causality without corresponding methodological grounding—whether in the form of explicit graphical models or formal identification strategies—should be treated as insufficiently substantiated. This shift in expectation would recalibrate the threshold for publication, aligning it more closely with the inferential demands of causal reasoning. At the same time, reviewers are uniquely positioned to ensure that authors engage substantively with the conceptual challenges outlined in this work, particularly where design recommendations or policy-relevant conclusions are presented. Requiring counterfactual robustness checks in such cases strengthens the link between model output and actionable insight, reducing the risk that recommendations rest on fragile or unexamined assumptions. Retrospective consideration of influential studies, including Butler *et al.* [4] and Schmidt *et al.* [5], indicates that such scrutiny could have materially improved the interpretive clarity of their contributions.

At the level of the broader research community, the task shifts from individual practice to collective infrastructure. The absence of shared benchmarks for causal inference in materials science currently limits the comparability and cumulative value of individual studies. Developing standardized evaluation frameworks tailored to core problem classes—such as structure–property relationships or synthesis–structure pathways—would provide a common reference against which causal claims can be assessed. Equally important is the need for methodological tools designed to accommodate the distinctive characteristics of materials data, including nonlinearity, sparsity, and high dimensionality. Investment in such toolkits would enable more rigorous causal analysis without imposing unrealistic assumptions derived from other domains. Alongside these technical developments, the establishment of reporting guidelines analogous to those used in fields such as epidemiology would formalize expectations around transparency and reproducibility, while remaining sensitive to the specificities of materials AI. Early efforts in this direction, including those by Ramakrishna *et al.* [3] and Vuković *et al.* [17], illustrate the feasibility of coordinated action and provide a foundation upon which more comprehensive standards can be built.

Taken together, these role-specific adjustments do more than correct isolated deficiencies; they reorient the field toward a more coherent integration of predictive modeling and causal reasoning. By aligning methodological rigor with institutional expectations and shared infrastructure, the proposed measures create conditions under which materials AI can evolve from a predominantly correlational enterprise into one capable of supporting robust scientific explanation and intervention.

Conclusion

This review demonstrates that materials informatics remains firmly anchored in correlation despite the conceptual availability of Pearl's hierarchy and early applications of causal methods. While predictive performance has advanced dramatically, a genuine understanding of interventions and counterfactuals is rare, constrained by the six documented

challenges and five persistent gaps. Moving forward requires deliberate integration of structural causal models and do-calculus into everyday workflows. Only then can materials AI fulfill its promise of autonomous, intervention-driven discovery. The field stands at a pivotal transition: from data-driven prediction to causally empowered design.

Acknowledgements

None

Conflict of interest

None

Financial supports

None

Ethics statement

None

Received: 30 Sep 2023 Revised: 22 Jan 2024 Accepted: 28 Mar 2024
Published online: 18 July 2024

References

He H, Zhang H, Roth D. Rethinking with retrieval: Faithful large language model inference. arXiv [Preprint]. 2022:arXiv:2301.00303.

Pearl J, Mackenzie D. The book of why: The new science of cause and effect. New York: Basic Books; 2018.

Ramakrishna S, Zhang TY, Lu WC, Qian Q, Low JS, Yune JH, et al. Materials informatics. J Intell Manuf. 2019;30(6):2307-26.

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

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

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

Montoya JH, Aykol M, Anapolsky A, Gopal CB, Herring PK, Hummelshøj JS, et al. Toward autonomous materials research: Recent progress and future challenges. Appl Phys Rev. 2022;9(1):011405.

Ting JY, Barnard AS. Data-driven causal inference of process-structure relationships in nanocatalysis. Curr Opin Chem Eng. 2022;36:100818.

Ziatdinov M, Nelson CT, Zhang X, Vasudevan RK, Eliseev E, Morozovska AN, et al. Causal analysis of competing atomistic mechanisms in ferroelectric materials from high-resolution scanning transmission electron microscopy data. NPJ Comput Mater. 2020;6(1):127.

Oviedo F, Ferres JL, Buonassisi T, Butler KT. Interpretable and explainable machine learning for materials science and chemistry. Acc Mater Res. 2022;3(6):597-607.

Ting JY, Li S, Barnard AS. Causal paths allowing simultaneous control of multiple nanoparticle properties using multi-target bayesian inference. Adv Theory Simul. 2022;5(10):2200330.

Pawar S, Sharma R, Palshikar GK, Bhattacharyya P, Varma V. Cause–effect relation extraction from documents in metallurgy and materials science. Trans Indian Inst Met. 2019;72(8):2209-17.

Vasudevan RK, Ziatdinov M, Vlcek L, Kalinin SV. Off-the-shelf deep learning is not enough, and requires parsimony, bayesianity, and causality. NPJ Comput Mater. 2021;7(1):16.

Zhu L, Zhou J, Sun Z. Materials data toward machine learning: Advances and challenges. J Phys Chem Lett. 2022;13(18):3965-77.

Zhou B, Li X, Liu T, Xu K, Liu W, Bao J. CausalKGPT: Industrial structure causal knowledge-enhanced large language model for cause analysis of quality problems in aerospace product manufacturing. Adv Eng Inform. 2024;59:102333.

Naser MZ. Causality and causal inference for engineers: Beyond correlation, regression, prediction and artificial intelligence. Wiley Interdiscip Rev Data Min Knowl Discov. 2024;14(4):e1533.

Vuković M, Thalmann S. Causal discovery in manufacturing: A structured literature review. J Manuf Mater Process. 2022;6(1):10.

R Fox Z, Ghosh A. Active causal learning for decoding chemical complexities with targeted interventions. Mach Learn Sci Technol. 2024;5(3):035056.

Lagemann K, Lagemann C, Taschler B, Mukherjee S. Deep learning of causal structures in high dimensions under data limitations. Nat Mach Intell. 2023;5(11):1306-16.

Liu R, Misra S. Monitoring the propagation of mechanical discontinuity using data-driven causal discovery and supervised learning. Mech Syst Signal Process. 2022;170:108791.

Liu T, Barnard AS. The emergent role of explainable artificial intelligence in the materials sciences. Cell Rep Phys Sci. 2023;4(10):101630.

Kelly J, Berzuini C, Keavney B, Tomaszewski M, Guo H. A review of causal discovery methods for molecular network analysis. Mol Genet Genomic Med. 2022;10(10):e2055.

Ham DW, Imai K, Janson L. Using machine learning to test causal hypotheses in conjoint analysis. Political Anal. 2024;32(3):329-44.

Huber M. An introduction to causal discovery. Swiss J Econ Stat. 2024;160(1):14.

Wang X, Zhao M, Meng F, Liu X, Kong Z, Chen X. Quantifying causal path-specific importance in structural causal model. Computation. 2023;11(7):133.

Kock AB, Preinerstorfer D, Veliyev B. Functional sequential treatment allocation. J Am Stat Assoc. 2022;117(539):1311-23.