

ORIGINAL RESEARCH

Open access

Optimization without Causality: Limits of Correlation-Driven Materials Design

Ravi Kumar^{1*}, Neha Sharma¹, Arjun Nair²

Abstract

In the evolving landscape of computational and data-driven materials engineering, machine learning techniques have revolutionized the discovery and optimization of materials by leveraging vast datasets to identify patterns and correlations. However, this reliance on correlation-driven approaches often overlooks the underlying causal mechanisms that govern material properties and behaviors, leading to inherent limitations in the generalizability and robustness of designed materials. This manuscript explores the conceptual boundaries of optimization strategies that prioritize statistical associations over causal understanding within materials informatics ecosystems. We introduce a novel conceptual framework, termed the Correlation Boundary Architecture (CBA), which delineates the epistemic constraints imposed by correlation-centric pipelines in materials design. The CBA integrates representation learning, inference dynamics, and feedback structures to highlight how data-driven optimizations can falter in extrapolative scenarios, such as novel chemical spaces or extreme conditions. By synthesizing recent advancements in graph neural networks, high-throughput computations, and uncertainty quantification, we articulate the trade-offs between computational efficiency and causal fidelity. Implications extend to autonomous discovery systems and inverse design paradigms, suggesting pathways for hybrid frameworks that mitigate correlation biases through enhanced interpretive layers. This work underscores the need for computational steering logics that balance correlative power with causal awareness, fostering more resilient materials engineering practices.

Keywords Materials informatics, Machine learning, Representation learning, Correlation limits, Causal deficits, Data-driven optimization

*Correspondence:

Ravi Kumar

ravi.kumar@gmail.com

¹ Department of Materials Data Engineering, Faculty of Engineering, IIT Delhi, New Delhi, India

² Department of Computational Materials Systems, Faculty of Engineering, IIT Bombay, Mumbai, India

Introduction

The emergence of computational materials paradigms

The field of materials engineering has undergone a profound transformation with the integration of computational methods and data-driven strategies. Traditionally reliant on experimental trial-and-error or physics-based simulations, materials discovery now increasingly employs informatics approaches to accelerate innovation [1, 2]. This shift is driven by the exponential growth in computational power and the availability of multimodal datasets, enabling high-throughput screening and predictive modeling at scales previously unattainable. For instance, advancements in density functional theory coupled with machine learning have facilitated the rapid evaluation of material properties across vast composition spaces [3, 4].

At the core of this paradigm is the utilization of statistical correlations derived from large-scale data repositories. These correlations allow for the identification of trends in material behaviors, such as electronic properties or mechanical

strength, without necessitating a complete mechanistic understanding [5, 6]. Such methods have proven instrumental in applications ranging from polymer design to alloy optimization, where genetic algorithms and generative models sample chemical spaces efficiently [5, 7]. However, this correlative focus introduces subtle yet significant constraints, particularly when optimizations are pursued in the absence of explicit causal linkages.

Challenges in data-driven optimization

Data-driven optimization in materials science often operates under the assumption that strong correlations suffice for effective design. Machine learning models, including graph neural networks and representation learning frameworks, excel at capturing these associations from stoichiometric or structural data [7, 8]. Yet, correlations do not imply causation, and this distinction becomes critical in complex systems where emergent properties arise from intricate interactions [9, 10]. For example, in porous materials or nanoclusters, predictive models may overfit to observed data patterns, failing to account for underlying physical drivers [3, 11].

This challenge is exacerbated in inverse design scenarios, where the goal is to engineer materials with targeted functionalities. Here, correlation-driven approaches might generate candidates that perform well within trained domains but exhibit brittleness in unexplored regimes [6, 12]. Uncertainty quantification techniques attempt to mitigate these issues by estimating model confidence, but they often remain tethered to correlative metrics rather than causal validations [13, 14]. Moreover, the integration of simulation-experiment coupling in closed-loop systems amplifies these limitations, as feedback loops propagate correlation biases through iterative refinements [15, 16].

The role of representation and inference

Effective data-driven materials design hinges on robust representations that encode chemical and structural information. Deep learning architectures, such as those based on graph networks, have advanced this by learning hierarchical features from atomic configurations [8, 17]. These representations facilitate inference in high-dimensional spaces, enabling predictions for properties like adsorption energies or band gaps [4, 18]. However, the inferential process is inherently correlative, relying on pattern recognition rather than causal deduction, which can lead to epistemic gaps in understanding material stability or reactivity [11, 19].

Foundation models for science further amplify this by pre-training on diverse datasets, yet they too prioritize correlative generalization over causal insight [20, 21]. In autonomous discovery systems, this manifests as optimized workflows that favor speed and scalability but overlook causal bottlenecks, potentially resulting in suboptimal or unreliable material candidates [22, 23].

Positioning the conceptual inquiry

Despite these advancements, a systematic examination of the limits imposed by correlation-driven optimization remains underexplored. This manuscript addresses this gap by conceptualizing the boundaries where correlative methods diverge from causal necessities in materials design. Drawing on computational ecosystems, we propose the Correlation Boundary Architecture (CBA) as a framework to interpret these limits, offering insights into how data pipelines, model architectures, and discovery logics interact under causal deficits. Through this lens, we aim to guide future computational strategies toward more integrated approaches.

Theoretical Background & Literature Synthesis

Foundations of materials informatics and machine learning integration

Materials informatics has emerged as a foundational paradigm within computational materials science, representing the systematic convergence of data science, statistical learning, and physics-based modeling infrastructures. Rather than treating materials discovery as an exclusively simulation-driven or experimentally bounded enterprise, materials

informatics reframes the field as a data-intensive knowledge system in which machine learning functions as a central epistemic engine for pattern extraction, inference generation, and discovery acceleration [1, 2]. Within this paradigm, heterogeneous data ecosystems—spanning density functional theory repositories, high-throughput screening outputs, spectroscopy archives, and microstructural imaging datasets—are transformed into structured learning environments where algorithmic models infer structure–property relationships at scale.

Early methodological developments in materials informatics were anchored in descriptor engineering, a process through which high-dimensional material characteristics were compressed into tractable numerical representations suitable for statistical modeling [9, 24]. These descriptors encoded crystallographic symmetries, electronic distributions, bonding environments, and thermodynamic attributes, enabling the translation of complex physicochemical systems into vectorized formats compatible with regression and classification architectures. Descriptor construction often leveraged compressed sensing principles and domain-informed feature engineering, ensuring that physically meaningful invariances—such as rotational symmetry or compositional conservation—were preserved within the encoded space [9, 25].

This descriptor-centric paradigm facilitated one of the earliest successes of machine learning in materials science: correlation-based property prediction. By mapping compositional or structural descriptors to experimentally measured or simulation-derived properties, models could identify statistical regularities that guided materials optimization strategies. Such mappings underpinned accelerated screening workflows, allowing researchers to rank candidate compounds for targeted synthesis or further computation without exhaustively simulating the full design space.

As the field matured, supervised and unsupervised learning frameworks became increasingly embedded within solid-state materials research infrastructures. Supervised architectures—trained on labeled datasets of structures and associated properties—enabled predictive modeling across electronic, mechanical, and thermodynamic domains [2, 26]. In parallel, unsupervised techniques facilitated clustering, anomaly detection, and latent structure identification, revealing hidden taxonomies within compositional and microstructural datasets.

Model selection within these early ecosystems was often shaped by dataset scale and sparsity. Ensemble methods such as random forests and gradient boosting machines proved particularly effective in small-data regimes, where their capacity to manage nonlinear interactions and mitigate overfitting addressed persistent challenges in specialized materials domains [12, 27]. These algorithms supported predictive inference in contexts where experimental data acquisition remained costly or slow, thereby extending the reach of computational exploration.

The integration of machine learning into materials workflows also reconfigured accessibility. By abstracting complex quantum-mechanical simulations into data-driven surrogate models, materials informatics democratized exploratory analysis. Researchers could interrogate vast compositional spaces, prioritize synthesis targets, and explore hypothetical materials systems without requiring prohibitive computational resources [3, 28]. In this sense, machine learning did not merely accelerate discovery pipelines; it redistributed epistemic agency across the materials research ecosystem, enabling broader participation in computational design.

Representation learning and deep architectures

A transformative shift within materials informatics has been the transition from hand-crafted descriptors to autonomous representation learning. In contrast to feature engineering—where domain experts explicitly define relevant attributes—representation learning architectures derive hierarchical features directly from raw or minimally processed inputs. This shift reflects a broader epistemological movement: from encoding human-interpretable descriptors toward constructing machine-interpretable latent spaces capable of capturing multiscale material phenomena.

Graph neural networks (GNNs) have emerged as a particularly consequential class of representation learning models within materials science. By formalizing atomic systems as graphs—where nodes represent atoms and edges encode interatomic interactions—GNNs provide a natural mathematical substrate for modeling crystalline and molecular

structures [8, 29]. Message-passing mechanisms within these networks propagate information across atomic neighborhoods, enabling the emergent encoding of bonding environments, coordination geometries, and long-range structural dependencies.

Such architectures enable end-to-end learning pipelines that map stoichiometry or structural graphs directly to property predictions, bypassing intermediate descriptor construction [7, 30]. This architectural integration enhances correlation detection efficiency while preserving structural fidelity, allowing models to internalize physicochemical regularities embedded within training datasets. As a result, GNNs have demonstrated strong performance across diverse predictive tasks, including bandgap estimation, formation energy prediction, and elastic property inference.

Beyond predictive modeling, deep learning architectures have extended materials informatics into generative and inverse design domains. Generative adversarial networks (GANs), variational autoencoders (VAEs), and diffusion-based models enable the sampling of novel material compositions and structures conditioned on desired property targets [6, 31]. Within polymer science, alloy engineering, and functional ceramics, such generative systems explore latent compositional manifolds to identify candidates exhibiting optimized dielectric, catalytic, or mechanical characteristics [5, 21].

These generative capabilities mark a conceptual transition from passive prediction to active design. Instead of merely forecasting properties of known compounds, deep architectures participate in hypothesis generation, proposing previously unobserved materials configurations aligned with specified performance criteria.

However, the epistemic foundations of these systems remain intrinsically correlative. Representation learning models derive inferential capacity from statistical regularities embedded in training data distributions. Consequently, their extrapolative reach is bounded by the diversity, quality, and representational completeness of underlying datasets [19, 22]. When confronted with underrepresented chemistries, extreme thermodynamic regimes, or novel synthesis pathways, such architectures may revert to interpolation within familiar latent regions rather than genuine discovery.

This data-boundedness introduces critical limitations. Latent representations, while mathematically expressive, may conflate correlation with causation, embedding spurious associations that lack mechanistic grounding. In inverse design contexts, this can produce candidate materials that satisfy algorithmic criteria yet remain physically unrealizable or synthetically inaccessible.

Moreover, deep representations often obscure interpretability. As hierarchical embeddings grow in dimensional complexity, tracing property predictions back to specific structural or electronic drivers becomes increasingly challenging. This opacity complicates scientific validation, particularly in autonomous discovery pipelines where interpretability, uncertainty quantification, and experimental coupling are essential for epistemic trust.

High-throughput computation and autonomous systems

High-throughput computational frameworks automate the screening of material candidates, integrating density functional theory with machine learning surrogates [4, 16]. This synergy accelerates discovery pipelines, as seen in scintillator or battery material explorations, where physics-informed models bridge simulation and prediction [10, 30]. Autonomous discovery systems extend this by incorporating closed-loop experimentation, where models iteratively refine hypotheses based on correlative feedback [15, 23].

Such systems emphasize workflow automation, with tools for data mining and model deployment facilitating seamless integration [13, 29]. Yet, the reliance on correlative surrogates introduces vulnerabilities, particularly in multimodal datasets where disparate data sources may embed inconsistent associations [3, 20].

Inverse design and generative paradigms

Inverse materials design inverts the traditional property-prediction workflow, using optimization algorithms to identify structures meeting predefined criteria [6, 5]. Genetic algorithms and variational autoencoders exemplify this, generating candidates through correlative exploration of chemical spaces [5, 7]. In contexts like full-Heusler compounds or porous frameworks, these paradigms leverage big-data science to uncover hidden patterns [3, 31].

Generative models further enhance this by producing distributions over molecular or crystalline spaces, enabling efficient sampling [6, 18]. However, the absence of causal anchors can lead to designs that are correlation-optimal but physically implausible [11, 26].

Uncertainty quantification and epistemic considerations

Uncertainty quantification addresses the reliability of correlative inferences, employing techniques like Bayesian frameworks or ensemble methods to estimate prediction variances [13, 14]. In materials AI, this is crucial for high-stakes applications, such as energy storage or catalysis, where extrapolation risks are high [4, 23]. Tools for coordination environment identification or force field construction incorporate uncertainty to refine model outputs [11, 27].

Despite these efforts, epistemic uncertainties—stemming from incomplete causal knowledge—persist, as models often conflate aleatoric noise with systemic biases [9, 19]. This synthesis reveals a computational ecosystem rich in correlative power but constrained by causal oversights [1, 2].

Simulation-experiment coupling and feedback dynamics

Coupling simulations with experimental validation forms the backbone of modern materials pipelines, where data-driven models steer iterative refinements [15, 16]. Foundation models pretrained on scientific corpora enhance this coupling by providing transferable representations [20, 21]. In nanocluster or polymer systems, such dynamics enable adaptive optimization, adjusting parameters based on correlative signals [11, 24].

Feedback loops in these systems amplify correlative strengths but also propagate limitations, as unaddressed causal gaps can cascade through discovery cycles [22, 25]. Literature highlights the need for interpretive layers to discern correlation from causation, yet current architectures prioritize efficiency over such depth [8, 29].

Synthesis of computational trade-offs

Integrating these elements, the literature underscores a tension between correlative scalability and causal robustness in materials design [12, 28]. While advancements in deep learning and high-throughput methods have propelled the field, they reveal systemic boundaries where optimization falters without causal grounding [17, 18]. This background sets the stage for conceptualizing these limits through unified frameworks that interpret pipeline dynamics and inference interactions [26, 30].

Table 1. Correlation–Causality Trade-Offs in Computational Materials Design Ecosystems

Dimension	Correlation-Driven Strength	Causality-Driven Requirement	Trade-Off Tension	Pipeline Impact
Computational Speed	High via surrogate models	Slower mechanistic modeling	Efficiency vs realism	Accelerated screening
Data Scalability	Massive dataset compatibility	Limited by mechanistic data	Volume vs depth	Broad exploration
Interpretability	Latent embeddings opaque	Mechanistic traceability needed	Performance vs insight	Validation difficulty

Generative Design	Rapid candidate synthesis	Physical feasibility constraints	Creativity vs plausibility	Unrealizable outputs
Autonomous Discovery	Closed-loop optimization	Experimental causal validation	Automation vs grounding	Risk propagation

Proposed conceptual framework

Overview of the correlation boundary architecture

To address the epistemic constraints in correlation-driven materials design, we propose the Correlation Boundary Architecture (CBA), an original conceptual framework that maps the interfaces between data representations, inferential processes, and discovery outcomes. The CBA conceptualizes materials optimization as a layered system where correlative mechanisms operate within bounded domains, delineating zones of reliable performance from those prone to causal deficits. This architecture comprises three primary structural layers: the Data Ingestion Layer, the Inference Mediation Layer, and the Optimization Steering Layer, interconnected through feedback loops that modulate information flow.

In the Data Ingestion Layer, raw multimodal inputs—such as stoichiometric descriptors or structural graphs—are transformed into correlative embeddings, emphasizing pattern extraction over mechanistic interpretation. The Inference Mediation Layer then processes these embeddings via model architectures, generating predictions based on statistical associations. Finally, the Optimization Steering Layer directs design iterations, incorporating uncertainty signals to navigate correlative landscapes. The structural composition and functional roles of these layers are synthesized in [Table 2](#).

Table 2. Structural Layers and Functional Roles within the Correlation Boundary Architecture

CBA Layer	Core Function	Computational Components	Correlation Role	Causal Limitation Exposure
Data Ingestion Layer	Multimodal data encoding	Descriptors, crystal graphs, spectroscopy data	Extracts statistical structure	Filters mechanistic signals
Inference Mediation Layer	Pattern inference & embedding construction	GNNs, foundation models, surrogate predictors	Amplifies correlations	Embeds spurious associations
Optimization Steering Layer	Design navigation & candidate generation	GANs, VAEs, genetic algorithms	Explores correlative optima	Produces causally implausible designs
Feedback Systems	Iterative refinement	Closed-loop experimentation	Reinforces correlations	Propagates boundary biases

Feedback loops within the CBA enable dynamic adjustments, such as recalibrating representations in response to inference discrepancies, yet they highlight boundaries where correlative signals weaken without causal reinforcement. This structure captures the workflow dynamics inherent to computational materials pipelines, offering interpretive insights into how correlation limits manifest in inverse design or autonomous systems.

Pipeline dynamics and feedback structures

The CBA formalizes data-to-discovery pipelines as sequential yet recursive flows, where correlations propagate through stages of representation, inference, and optimization. A key dynamic is the attenuation of signal fidelity across layers, conceptualized as a trade-off between computational tractability and epistemic completeness. For instance, in high-throughput scenarios, rapid correlative screening accelerates candidate generation but risks overlooking causal interactions that emerge at finer scales.

Feedback loops serve as corrective mechanisms, looping inference outputs back to data refinement, yet they are constrained by the architecture's correlative foundation. This introduces steering logics that prioritize correlation maximization, such as gradient-based optimizations in neural networks, while exposing vulnerabilities in extrapolative contexts.

One conceptual formula that captures the interaction between correlative strength and boundary constraints may be expressed as:

$$\begin{aligned} C(D, M) &= \int_{\Omega} \rho(x) \cdot \kappa(M | x) dx - \lambda \cdot \Delta(C, P) \end{aligned} \quad (1)$$

Here, $C(D, M)$ represents the correlative capacity of the system given data D and model M integrated over the design space Ω with density $\rho(x)$ and conditional kernel κ . The penalty term $\lambda \cdot \Delta(C, P)$ accounts for the divergence Δ between correlative outputs C and physical priors P , symbolizing the causal deficit. This formula interprets the balance where excessive reliance on correlations amplifies boundary effects.

Interpretive layers and epistemic risks

To mitigate these risks, the CBA incorporates interpretive layers that analyze representation-inference interactions, such as assessing embedding robustness against perturbations. These layers reveal epistemic risk structures, where correlation-driven optimizations encounter limits in novel regimes, like transitioning from trained to unseen chemical compositions.

Another formula conceptualizes this risk as:

$$\begin{aligned} R(I) &= \frac{1}{N} \sum_i \left(1 - \sigma(I_i) \tau(B_i) \right) \exp(-d(I_i, B) \gamma) \end{aligned} \quad (2)$$

Where $R(I)$ denotes the epistemic risk for inferences I , averaged over N instances, with σ as correlative confidence, τ as boundary threshold, and d as distance to the correlative boundary B , modulated by decay γ . This expression captures how risk escalates as inferences approach or exceed correlative boundaries.

A third formula addresses feedback loop stability:

$$\begin{aligned} S(F) &= \prod_k \left(\alpha_k \cdot \phi(F_k) + (1 - \alpha_k) \cdot \psi(C_k) \right) \end{aligned} \quad (3)$$

Here, $S(F)$ is the stability of feedback F across K loops, blending adaptation factors ϕ with correlative contributions ψ weighted by α_k . This interprets loop dynamics as multiplicative interactions that sustain optimization within correlative confines but destabilize under causal stress. Key epistemic risk structures arising across the architecture are categorized in **Table 3** and illustrating trade-offs in computational workflows and boundary-conditioned optimization dynamics (**Figure 1**).

Table 3. Epistemic Risk Structures Across Correlation-Driven Materials Pipelines

Risk Category	Origin Layer	Mechanism	Manifestation in	Discovery
---------------	--------------	-----------	------------------	-----------

			Design	Consequence
Spurious Correlation Risk	Inference	Latent pattern overfitting	False property predictions	Misguided optimization
Extrapolation Risk	Optimization	Boundary overreach	Failure in novel chemistries	Low generalizability
Representation Bias	Data / Inference	Skewed dataset encoding	Underrepresented materials	Blind spots in discovery
Generative Instability	Optimization	Latent space distortion	Unrealizable structures	Experimental infeasibility
Feedback Amplification Risk	Systemic	Recursive correlation reinforcement	Escalating prediction errors	Pipeline fragility




Figure 1. Correlation Boundary Architecture (CBA) for data-driven materials optimization. The framework conceptualizes correlation-centric discovery pipelines as a three-layer system comprising data ingestion, inference mediation, and optimization steering. Correlative signals propagate across layers through representation learning and generative design engines, enabling accelerated materials optimization. However, boundary envelopes emerge where statistical associations diverge from causal physical mechanisms. Embedded epistemic risk nodes highlight zones of extrapolation deficit and spurious inference. Inner and outer feedback loops reinforce correlative refinement while simultaneously propagating boundary constraints. The architecture delineates reliable optimization domains from causal deficit frontiers, illustrating systemic limits of correlation-driven materials design.

Analytical implications

Interpretive dynamics in materials pipelines

The Correlation Boundary Architecture (CBA) provides a lens for interpreting the dynamics of data-driven pipelines in materials engineering, revealing how correlation-centric optimizations influence discovery trajectories. In representation learning contexts, the architecture highlights interactions where embeddings derived from stoichiometric or graph-based inputs may capture surface-level associations but fail to encode deeper physical constraints [7, 8]. This interpretive insight suggests that as pipelines scale to multimodal datasets, the risk of correlation inflation—where spurious patterns dominate—escalates, potentially steering optimizations toward non-generalizable solutions [3, 20].

Computational steering logics within the CBA underscore trade-offs in inverse design workflows. For instance, generative models that sample composition spaces rely on correlative gradients to navigate optima, yet the architecture interprets these as bounded by epistemic horizons where causal factors, such as thermodynamic stability, remain unaddressed [6, 5]. This implies a need for enhanced feedback structures that incorporate boundary awareness, modulating inference to prioritize robust correlations over fleeting ones [11, 15].

Epistemic risk structures and uncertainty interactions

Epistemic risks in the CBA manifest through representation-inference mismatches, particularly in high-throughput systems where rapid screenings amplify correlative biases [4, 16]. The framework interprets uncertainty quantification not merely as a variance estimator but as a signal of boundary proximity, where elevated uncertainties indicate transitions from correlative reliability to causal ambiguity [13, 14]. In autonomous discovery setups, this risk structure affects closed-loop iterations, as feedback loops may reinforce correlated patterns without challenging underlying assumptions [22, 23].

A conceptual formula that captures this uncertainty-boundary interaction can be expressed as:

$$U(R, B) = \beta \cdot \left(\int \partial B \nabla R \cdot ds \right) + \eta \cdot \left\| \frac{R}{P} \right\|_2 \quad (4)$$

Here, $U(R, B)$ denotes the uncertainty at representation R relative to boundary B , with surface integral capturing gradient flows across boundaries and a norm term $\frac{\eta}{\| \frac{R}{P} \|_2}$ reflecting deviation from physical priors P , weighted by β and η .

This formula interprets how uncertainties accumulate at correlative edges, informing risk-aware steering in materials AI.

Infrastructure trade-offs in discovery ecosystems

The CBA elucidates infrastructure-level trade-offs in coupling simulation and experimentation, where correlative efficiencies enable scalability but impose limits on extrapolative power [15, 21]. For example, in foundation models pretrained on scientific data, the architecture reveals how transfer learning propagates correlative priors, potentially entrenching biases in downstream optimizations [20, 29]. This implies that discovery ecosystems benefit from layered interpretations that balance computational cost with boundary mitigation, such as hybrid logics integrating correlative speed with causal checkpoints [1, 2].

In machine learning for small datasets or specialized materials, the framework highlights how feature compression techniques constrain the correlative domain, trading dimensionality reduction for epistemic completeness [9, 12]. These implications extend to graph neural networks, where node-embedding dynamics are interpreted as correlative amplifiers susceptible to boundary distortions in complex topologies [8, 27].

Systems-level insights for optimization resilience

At a systems level, the CBA offers insights into enhancing optimization resilience by interpreting feedback as correlative stabilizers. In polymer or alloy design, where genetic algorithms drive iterations, the architecture suggests that loop dynamics can be steered to probe boundary regions, fostering adaptive pipelines that evolve beyond pure correlations [5, 25]. This interpretive approach advocates for computational workflows that embed risk structures, ensuring that discovery logics account for correlation limits in real-time [19, 30].

Another formula conceptualizing resilience trade-offs is:

$$T(O, C) = \max_{\theta} (\theta(O \top C \theta - \mu \cdot Tr(\Sigma C)) \quad (5)$$

Where $T(O, C)$ represents the trade-off for optimization O under correlations C , maximized over parameters θ , penalized by trace of covariance ΣC scaled by μ . This captures the interaction between optimization objectives and correlative variability, interpreting resilience as bounded maximization.

These analytical implications collectively frame the CBA as a tool for navigating correlation-driven constraints, promoting interpretive strategies that enhance the robustness of materials design ecosystems [18, 31].

Results and Discussion

The conceptual exploration of optimization without causality through the Correlation Boundary Architecture (CBA) illuminates fundamental boundaries in computational materials engineering. By synthesizing advancements in materials informatics and machine learning, the framework interprets how correlation-driven paradigms, while computationally advantageous, introduce epistemic vulnerabilities that affect discovery reliability [1-3]. This discussion integrates these insights, examining broader implications for data-driven workflows and potential avenues for mitigation.

Central to the CBA is the recognition that representation learning and deep architectures, such as graph neural networks, excel in correlative pattern extraction but often operate within implicit boundaries that causal deficits exacerbate [7, 8, 27]. In high-throughput and autonomous systems, this manifests as pipeline dynamics where feedback loops perpetuate correlative biases, as evidenced in inverse design applications for polymers or alloys [5, 6, 15]. The interpretive layers of the CBA suggest that incorporating boundary-aware steering could refine these dynamics, allowing systems to flag when correlations diverge from physical plausibility [11, 13].

Uncertainty quantification emerges as a critical interaction point, where the framework's risk structures interpret uncertainties as indicators of correlative limits rather than mere statistical artifacts [14, 19]. This perspective is particularly relevant in multimodal datasets or simulation-experiment couplings, where disparate data sources may embed conflicting associations, leading to infrastructure trade-offs [3, 16, 20]. For instance, foundation models for science, while enabling scalable inferences, risk entrenching domain-specific correlations that falter in cross-disciplinary applications [21, 29].

The CBA's systems-level insights advocate for hybrid computational logics that blend correlative efficiency with causal interpretive elements, without mandating full mechanistic models [4, 9, 12]. In contexts like scintillator discovery or coordination environment prediction, this could involve adaptive feedback that probes epistemic risks, enhancing optimization resilience [10, 28]. Moreover, the framework's formulas provide symbolic tools for conceptualizing these interactions, offering a basis for analyzing workflow trade-offs in generative paradigms [6, 18, 25].

Challenges persist in implementing such interpretations, as current ecosystems prioritize speed over depth, potentially overlooking the long-term costs of correlation reliance [22, 23, 30]. Future computational strategies might leverage the CBA to design more balanced pipelines, fostering innovations that extend beyond correlative horizons [11, 17, 31]. This

discussion underscores the value of conceptual frameworks in guiding materials engineering toward sustainable, robust practices.

Conclusion

In summary, this manuscript has conceptualized the limits of correlation-driven materials design within computational and data-driven ecosystems, introducing the Correlation Boundary Architecture (CBA) as a novel framework for interpreting these constraints. By delineating structural layers, pipeline dynamics, and epistemic risk structures, the CBA highlights the trade-offs inherent in optimizations that sidestep causality, drawing on advancements in machine learning, representation learning, and uncertainty quantification.

The analytical implications and discussions reveal pathways for enhancing discovery resilience, emphasizing interpretive strategies that balance correlative power with boundary awareness. Ultimately, the CBA serves as a conceptual tool to steer materials engineering toward more integrated paradigms, mitigating the epistemic pitfalls of pure correlation reliance and paving the way for causally informed innovations.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 22 Dec 2022

Revised: 15 Mar 2023

Accepted: 05 Apr 2023

Published online: 18 September 2023

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

Ramprasad R, Behler J, Zunger A, Willand A, Ceriotti M. Machine learning in materials informatics: Recent applications and prospects. *npj Comput Mater.* 2017;3(54):1-13.
<https://doi.org/10.1038/s41524-017-0056-5>.

Schmidt J, Marques MRG, Botti S, Marques MAL. Recent advances and applications of machine learning in solid-state materials science. *npj Comput Mater.* 2019;5(83):1-36.
<https://doi.org/10.1038/s41524-019-0221-0>.

Jablonka KM, Ongari D, Moosavi SM, Smit B. Big-data science in porous materials: materials genomics and machine learning. *Chem Rev.* 2020;120(16):8066-129.
<https://doi.org/10.1021/acs.chemrev.0c00004>.

Fung V, Hu J, Ganesh P, Sumpter BG. Machine learned features from density of states for accurate adsorption energy prediction. *Nat Commun.* 2021;12(88):1-11.

Kim C, Batra R, Chen L, Tran H, Ramprasad R. Polymer design using genetic algorithm and machine learning. *Comput Mater Sci.* 2021;186:110067.
<https://doi.org/10.1016/j.commatsci.2020.110067>.

Dan Y, Zhao Y, Li X, Li S, Hu M, Hu J. Generative adversarial networks (GAN) based efficient sampling of chemical composition space for inverse materials design. *Nat Commun.* 2020;11(2312):1-9.

Goodall REA, Lee AA. Predicting materials properties without crystal structure: deep representation learning from stoichiometry. *Nat Commun.* 2020;11(6280):1-9.
<https://doi.org/10.1038/s41467-020-19964-7>.

Zhang Y, Ling C. A strategy to apply machine learning to small datasets in materials science. *npj Comput Mater.* 2018;4(25):1-8.

Ghiringhelli LM, Vybiral J, Ahmetcik E, Ouyang R, Levchenko SV, Draxl C, et al. Learning physical descriptors for materials science by compressed sensing. *New J Phys.* 2017;19(023017):1-19.

Pilania G, McClellan KJ, Stanek CR, Uberuaga BP. Physics-informed machine learning for inorganic scintillator discovery. *J Phys Chem C.* 2018;122(52):29340-50.

Glielmo A, Zeni C, De Vita A. Efficient nonparametric n-body force fields from machine learning. *Phys Rev B.* 2018;97(18):184307.
<https://doi.org/10.1103/PhysRevB.97.184307>.

Glielmo A, Sollich P, De Vita A. Accurate interatomic force fields via machine learning with covariant kernels. *Phys Rev B.* 2017;95(21):214302.
<https://doi.org/10.1103/PhysRevB.95.214302>.

Ward L, Wolverton C. Atomistic calculations and materials informatics: A review. *Curr Opin Solid State Mater Sci.* 2017;21(3):167-76.
<https://doi.org/10.1016/j.cossms.2016.07.002>.

Huo H, Rong Z, Kononova O, Sun W, Botari T, He T, et al. Semi-supervised machine-learning classification of materials synthesis procedures. *Npj Comput Mater.* 2019;5(1):62.

Wei J, Chu X, Sun XY, Xu K, Deng HX, Chen J, et al. Machine learning in materials science. *InfoMat.* 2019;1(3):338-58.

Hutter F, Kotthoff L, Vanschoren J. Automated machine learning: methods, systems, challenges. Springer; 2019.

Behler J. First principles neural network potentials for reactive simulations of large molecular and condensed systems. *Angew Chem Int Ed.* 2017;56(42):12828-40.

Bartók AP, De S, Poelking C, Bernstein N, Kermode JR, Csányi G, et al. Machine learning unifies the modeling of materials and molecules. *Sci Adv.* 2017;3(12):e1701816.
<https://doi.org/10.1126/sciadv.1701816>.

Chen C, Zuo Y, Ye W, Li X, Ong SP. Learning properties of ordered and disordered materials from multi-fidelity data. *Nat Comput Sci.* 2021;1(1):46-53.
<https://doi.org/10.1038/s43588-020-00002-x>.

Court CJ, Cole JM. Auto-generated database of semiconductor band gaps using chemdataextractor. *Sci Data.* 2020;7(1):1.

Tran H, Kim S, Batra R, Tran TM, Kim C, Loeffler L, et al. Polymerization kinetics under confinement. *npj Comput Mater.* 2019;5(1):96.

Mannodi-Kanakthodi A, Chandrasekaran A, Kim C, Huan TD, Pilania G, Botu V, et al. Scoping the polymer genome: A roadmap for rational polymer dielectrics design and beyond. *Mater Today.* 2018;21(8):785-96.
<https://doi.org/10.1016/j.mattod.2017.11.021>.

Doan HA, Agarwal G, Qian H, Counihan MJ, Rodriguez-Lopez J, Moore JS, et al. Quantum chemistry-informed active learning to accelerate the design and discovery of sustainable energy storage materials. *Chem Mater.* 2020;32(15):6338-46.
<https://doi.org/10.1021/acs.chemmater.0c00768>.

Amini M, Rahmani A. Machine learning process evaluating damage classification of composites. *Int J Sci AdvTechnol.* 2023;9(2023):240-50.

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

Ouyang R, Curtarolo S, Ahmetcik E, Scheffler M, Ghiringhelli LM. SISSO: A compressed-sensing method for identifying the best low-dimensional descriptor in an immensity of offered candidates. *Phys Rev Mater.* 2018;2(8):083802.
<https://doi.org/10.1103/PhysRevMaterials.2.083802>.

Cubuk ED, Sendek AD, Reed EJ. Screening billions of candidates for solid lithium-ion conductors: A transfer learning approach for small data. *J Chem Phys.* 2019;150(21):214701.
<https://doi.org/10.1063/1.5093220>.

Zheng C, Chen C, Chen Y, Ong SP. Random forest models for accurate identification of coordination environments from X-ray absorption near-edge structure. *npj Comput Mater.* 2020;6(1):134.

Ward L, Dunn A, Faghaninia A, Zimmermann NER, Bajaj S, Wang Q, et al. Matminer: An open source toolkit for materials data mining. *Comput Mater Sci.* 2018;152:60-9.
<https://doi.org/10.1016/j.commatsci.2018.05.018>.

Bartel CJ, Chen C, Tang X, McKeown Wessner A, Goldsmith BR, et al. Accelerating high-throughput calculations of lithium diffusion in battery cathodes with machine learning. *Chem Mater.* 2020;32(18):8058-66.

Kaufmann K, Maryanovsky D, Mellor WM, Zhu C, Rosengarten AS, Harrington TJ, et al. Discovery of high-entropy ceramics via machine learning. *npj Comput Mater.* 2020;6(1):42.
<https://doi.org/10.1038/s41524-020-0317-6>.