

ORIGINAL RESEARCH

Open access

Topology without Physics: Structural Abstraction Limits in Graph-Based Materials Models

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

Abstract

The advent of computational and data-driven approaches in materials engineering has transformed discovery pipelines, leveraging machine learning and graph-based representations to navigate vast chemical spaces. However, these models often prioritize topological abstractions over intrinsic physical mechanisms, leading to epistemic constraints in predictive accuracy and interpretability. This manuscript introduces a conceptual framework that dissects the structural abstraction limits inherent in graph-based materials models, emphasizing the trade-offs between computational efficiency and physical fidelity. By synthesizing insights from materials informatics and representation learning, we explore how graph neural networks decouple topological features from underlying physics, potentially hindering autonomous discovery systems and inverse design workflows. The framework delineates layers of abstraction, from data ingestion to inference, highlighting feedback loops that amplify abstraction-induced uncertainties. Implications extend to high-throughput computation, multimodal datasets, and uncertainty quantification, advocating for integrated infrastructures that balance abstraction with mechanistic reintegration. This analysis fosters a deeper understanding of computational steering in materials AI, guiding future developments toward more robust, physics-aware discovery paradigms without empirical validation. Ultimately, addressing these limits could enhance the reliability of data-driven materials engineering ecosystems.

Keywords Materials informatics, Uncertainty quantification, Graph neural networks, Representation learning, Computational discovery, Data infrastructures

*Correspondence:

Meenal Joshi

meenal.joshi@gmail.com

¹ Department of Materials Data Science, Faculty of Engineering, Savitribai Phule Pune University, Pune, India

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

Introduction

The field of computational materials engineering has undergone a profound shift over the past decade, driven by the integration of data-centric methodologies and advanced algorithmic frameworks. This evolution is rooted in the recognition that traditional experimental approaches, while foundational, are inherently limited by time, cost, and scalability constraints when exploring the immense combinatorial space of potential materials. High-throughput computation has emerged as a cornerstone, enabling the systematic screening of thousands of candidate structures through density functional theory and related simulations

[1]. Concurrently, the rise of machine learning in materials science has accelerated this process, offering surrogate models that approximate complex physical properties with remarkable efficiency [2-4].

At the heart of these advancements lies materials informatics, a discipline that harnesses large-scale datasets to inform design and optimization. Multimodal materials datasets, encompassing structural, electronic, and thermodynamic information, serve as the bedrock for training sophisticated models [5, 6]. These datasets are often curated from high-throughput repositories, facilitating the application of deep learning architectures such as

graph neural networks (GNNs), which represent materials as interconnected graphs of atoms and bonds [7-9]. Such representations capture topological relationships effectively, allowing for predictions of properties like bandgaps, stability, and mechanical strength without exhaustive simulations [10-12].

Yet, the data-driven paradigm introduces epistemic challenges that warrant careful examination. In graph-based models, the emphasis on structural abstraction—distilling materials into nodes, edges, and features—often sidelines the nuanced physical interactions that govern real-world behavior. For instance, while GNNs excel in learning from stoichiometric or crystallographic data [11, 13], they may overlook quantum mechanical subtleties or environmental dependencies that are not explicitly encoded. This abstraction facilitates scalability but imposes limits on the models' ability to generalize beyond trained domains, particularly in inverse materials design where desired properties must map back to viable structures [14, 15].

High-throughput infrastructures further amplify these dynamics. Autonomous discovery systems, which couple simulation with experimentation in closed loops, rely on AI to steer exploration [16, 17]. However, when guided by abstracted graph models, these systems risk propagating uncertainties derived from incomplete physical representations. Uncertainty quantification in materials AI becomes critical here, as it addresses not only statistical variances but also structural biases inherent in the data pipelines [18, 19]. Representation learning architectures, including crystal graph convolutional networks and attention mechanisms, attempt to mitigate this by incorporating multi-fidelity data or global features [7, 10, 20], but the core tension between topology and physics persists.

Epistemic constraints manifest in several ways within computational materials engineering. First, the decoupling of topological features from physical laws can lead to models that prioritize pattern recognition over causal understanding, limiting their utility in novel discovery scenarios [3, 21]. Second, in simulation-experiment coupling, abstracted models may fail to align with empirical realities, necessitating iterative refinements that consume resources [22, 23]. Third, foundation models for science, which aim to unify diverse datasets under a single framework, often inherit these abstraction limits, affecting downstream tasks like property prediction in polycrystalline or disordered materials [6, 9, 24].

This manuscript positions itself at the intersection of these challenges, introducing a novel conceptual framework to interrogate the structural abstraction limits in graph-based materials models. By framing the issue through computational workflow dynamics and representation-inference interactions, we seek to illuminate infrastructure trade-offs that influence discovery steering logics. The framework underscores the need for integrative approaches that acknowledge abstraction's role while advocating for mechanisms to reinfuse physical insights, thereby enhancing the epistemic robustness of data-driven materials ecosystems.

Theoretical Background & Literature Synthesis

Materials data infrastructures

The epistemic architecture of contemporary computational materials engineering is fundamentally scaffolded by large-scale data infrastructures that consolidate, standardize, and operationalize heterogeneous materials knowledge. These infrastructures, built upon high-throughput computation, density functional theory workflows, and automated simulation pipelines, have enabled the systematic aggregation of materials properties at unprecedented scale [1, 5]. Through iterative computational screening campaigns, repositories now encode thermodynamic stability, electronic structure, mechanical response, and transport characteristics across vast chemical and structural spaces. The resulting infrastructures do not merely store information; they actively structure the epistemic horizons of discovery by defining which materials domains are computationally legible and which remain underrepresented.

A defining strength of these infrastructures lies in their multimodal integration capacity. Contemporary databases increasingly synthesize crystallographic descriptors, spectroscopic signatures, microstructural imaging, and thermodynamic parameters into unified data ecosystems [6]. Such integration enables cross-modal inference, allowing machine learning systems to correlate structural motifs with emergent physicochemical properties. However, the harmonization required for interoperability introduces abstraction layers that reshape material reality into computationally tractable forms. Complex defect chemistries, metastable phase transitions, and dynamic environmental responses are often simplified into static

graph or tensor representations, generating epistemic compression where physically consequential nuances may be attenuated or omitted [3, 19].

Within this infrastructural context, uncertainty quantification emerges as a critical interpretive stabilizer. Variability in simulation fidelity, exchange–correlation functional selection, convergence thresholds, and experimental calibration propagates through data pipelines, necessitating probabilistic frameworks capable of contextualizing predictive outputs [18]. Autonomous discovery platforms—operating atop these infrastructures—depend on such quantified uncertainties to prioritize candidate screening and allocate experimental validation resources [16, 17]. Yet, their performance remains tightly coupled to the representational assumptions embedded within curated datasets.

Collaborative accessibility further amplifies infrastructural influence. Web-based materials informatics platforms democratize database interaction, enabling distributed modeling, shared benchmarking, and cross-institutional discovery initiatives. However, these platforms frequently privilege topological and compositional descriptors optimized for machine readability over mechanistic or process-dependent variables [5]. This descriptor prioritization subtly reorients discovery logics, privileging structurally describable phenomena while marginalizing context-sensitive physical behaviors. Consequently, data infrastructures function not only as repositories but as epistemic filters shaping the contours of computational exploration.

Representation learning architectures

Representation learning constitutes the algorithmic core through which materials data infrastructures become computationally actionable. Among available paradigms, graph neural networks (GNNs) have emerged as the dominant architecture for encoding material structures, formalizing atoms as nodes and interatomic interactions as edges within relational topologies [2, 4, 7–9, 12, 13, 20, 21, 24, 25]. This graph formalism enables hierarchical feature propagation, allowing models to learn embeddings that capture coordination environments, bonding motifs, and lattice connectivity patterns across scales.

Crystal graph convolutional networks exemplify this paradigm, iteratively propagating node features through message-passing operations to infer structure–property

relationships [6, 10, 11, 13]. Their demonstrated efficacy across ordered crystals, disordered alloys, and molecular frameworks has positioned them as foundational tools in materials informatics. By encoding local atomic neighborhoods alongside global connectivity, these architectures enable scalable prediction across compositional spaces previously inaccessible to conventional physics-based simulations.

Despite their success, the epistemic consequences of topological abstraction remain a subject of growing scrutiny. In decoupling geometry from underlying electronic and quantum mechanical processes, graph architectures risk privileging relational structure over governing physics. Attention mechanisms, adaptive edge updates, and equivariant message passing have been introduced to enhance representational expressivity [10, 20, 23], yet these augmentations remain embedded within discretized structural manifolds that may inadequately capture dynamic or field-dependent phenomena [16, 25].

Explainable artificial intelligence techniques have begun interrogating these representational spaces. Feature attribution analyses frequently reveal that model saliency concentrates on connectivity patterns, coordination counts, and bonding motifs rather than deeper electronic descriptors [18]. This alignment suggests that while embeddings encode structural logic effectively, they may underrepresent emergent physical drivers such as phonon interactions, defect energetics, or charge redistribution. The resulting interpretive asymmetry foregrounds a trade-off: abstraction enhances scalability and computational tractability, yet constrains fidelity when modeling complex systems such as polycrystalline interfaces, amorphous phases, or defect-rich materials [9].

AI-Guided discovery systems

The convergence of data infrastructures and representation learning architectures has catalyzed the emergence of AI-guided discovery systems characterized by closed-loop experimentation. Within these pipelines, predictive models iteratively generate hypotheses, guide candidate selection, and incorporate experimental or simulated feedback to refine subsequent exploration cycles [14, 16, 17]. Such systems operationalize materials discovery as a dynamic optimization process rather than a static screening exercise.

Graph-based inference engines are particularly effective within high-throughput environments, where they accelerate inverse mapping from desired properties to candidate structures [15, 22]. By navigating latent structural manifolds, these systems enable targeted exploration of functional materials for energy storage, catalysis, and electronic applications. However, the abstraction layers underpinning graph representations introduce steering biases. When predictive confidence is derived primarily from topological similarity, exploration trajectories may cluster around structurally familiar regions, inadvertently constraining novelty [19, 21].

The rise of scientific foundation models extends this paradigm further by pretraining representation systems across multimodal scientific corpora [3, 15]. These architectures aspire to unify chemical, structural, and textual knowledge into transferable embeddings capable of cross-domain generalization. Yet, they inherit epistemic constraints embedded within graph abstractions and curated training distributions, affecting interpretability, uncertainty calibration, and extrapolative reliability [6, 18].

Literature across autonomous discovery ecosystems increasingly emphasizes the need to rebalance abstraction with physical reintegration. Hybrid frameworks incorporating physics-informed constraints, simulation-aware embeddings, and experimental feedback assimilation have been proposed to stabilize discovery dynamics [1, 5]. Such integrations aim to prevent epistemic drift, wherein algorithmic optimization diverges from physically realizable design spaces.

Computational design paradigms

Inverse materials design represents a conceptual apex within computational discovery, reframing materials engineering from predictive analysis to generative synthesis. Rather than estimating properties from known structures, inverse paradigms algorithmically construct candidate materials optimized for targeted functionalities [14, 15, 22]. Graph neural networks and generative architectures facilitate this transition by exploring latent chemical spaces and proposing structurally viable configurations.

However, abstraction again mediates generative fidelity. Designs emerging from topological embeddings may satisfy connectivity constraints while neglecting thermodynamic stability, kinetic feasibility, or synthesis accessibility [9, 11,

24]. The resulting candidates occupy computationally valid yet physically tenuous regions of design space, highlighting the epistemic gap between structural plausibility and realizability.

High-throughput computational infrastructures partially mitigate this gap by supplying training corpora that encode stability filters and energetic constraints [1]. Nevertheless, scalability introduces its own trade-offs. Expanding dataset volume enhances model generalization but often necessitates descriptor simplification, reinforcing abstraction layers that distance generative reasoning from physical grounding.

Multitask learning and multi-fidelity modeling attempt to reconcile these tensions by integrating heterogeneous data scales—from coarse simulations to high-accuracy quantum calculations and experimental observations [6, 19]. These paradigms distribute epistemic weight across fidelity tiers, enabling models to learn stability gradients alongside structural embeddings. Yet, the dominance of topological descriptors persists, continuing to shape the generative logic of computational design [4, 7, 8].

Uncertainty & interpretability

Uncertainty quantification and interpretability frameworks function as epistemic counterbalances within materials AI ecosystems. In graph-based learning environments, abstraction amplifies epistemic distance between representation and reality, making confidence estimation essential for responsible inference [18, 19]. Probabilistic deep learning approaches—such as Bayesian neural networks and ensemble graph models—embed predictive distributions within structural reconstructions, enabling systems to express graded confidence across candidate predictions [17].

Interpretability research further interrogates the internal reasoning of materials AI models. Techniques including atoms-in-molecules networks, saliency mapping, and subgraph attribution analyses reveal how learned representations prioritize specific structural features [18, 19]. These investigations consistently indicate a weighting toward topological connectivity and coordination environments rather than deeper electronic or thermodynamic variables.

This interpretive asymmetry reinforces infrastructural trade-offs. Enhancing model transparency often requires

reintegrating physical descriptors, simulation constraints, or mechanistic priors into representation layers [3, 16, 21]. Consequently, uncertainty and interpretability do not operate as peripheral analytical tools but as structural correctives that realign abstraction with physical realism.

Synthesis insight

Across infrastructures, architectures, and discovery systems, a recurring epistemic tension emerges: abstraction enables scalability, interoperability, and algorithmic acceleration, yet simultaneously compresses physical nuance. The literature collectively indicates that future materials AI ecosystems must evolve toward hybrid representational regimes—where topological efficiency is continuously counterbalanced by mechanisms of physical reintegration, uncertainty contextualization, and interpretive transparency.

Proposed conceptual framework

To address the structural abstraction limits in graph-based materials models, we introduce the Topology-Physics Decoupling Framework (TPDF), a layered conceptual architecture that maps the dynamics of abstraction across data ingestion, model construction, and discovery inference. TPDF conceptualizes materials representations as stratified pipelines, where topological elements are progressively abstracted from physical underpinnings, leading to emergent trade-offs in computational steering.

At the base layer, data infrastructures feed multimodal inputs into graph encoders, abstracting atomic topologies while filtering physical descriptors like electronic densities. This abstraction can be conceptualized as a mapping function $A \left(\begin{matrix} T \\ \rightarrow \\ T' \end{matrix} ; P \right)$, where T denotes topological features (nodes, edges), P represents physical parameters (e.g., interaction potentials), and T' is the abstracted topology, capturing the interaction between raw data and model readiness without empirical tuning.

Ascending layers involve representation learning, where GNNs propagate T' through convolutional operations, amplifying efficiency but introducing decoupling risks. Feedback loops within TPDF reintegrate abstracted outputs with upstream data, mitigating losses via iterative refinement logics. For instance, uncertainty propagation may be expressed as $U = f(\Delta T, \Sigma P)$, where ΔT quantifies topological distortion and ΣP aggregates physical

variances, illustrating the workflow dynamics that steer discovery toward balanced outcomes.

The apex layer focuses on inference and design, where abstracted topologies inform predictions, but epistemic risks arise from unrecovered physics. TPDF incorporates steering mechanisms that dynamically adjust abstraction levels, fostering integrative discovery without predictive claims. These layered abstraction dynamics and feedback reintegration pathways are conceptually synthesized in **Figure 1**.




Figure 1. Topology-Physics Decoupling Framework (TPDF): Stratified Abstraction Dynamics in Graph-Based Materials Modeling

Stratified architecture of the Topology-Physics Decoupling Framework illustrating progressive abstraction from physics-rich data infrastructures to topology-dominant inference layers, alongside feedback mechanisms for physical reintegration.

The stratified abstraction architecture underlying TPDF, spanning data ingestion to inference steering, is structurally summarized in **Table 1**.

Table 1. Layers of Structural Abstraction in Graph-Based Materials Modeling Pipelines

Layer	Computational Function	Topological Representation Role	Physical Fidelity Element
Data Ingestion Layer	Aggregates multimodal materials datasets	Converts atomic structures into graph inputs (nodes/edges)	Elemental composition and field representation
Graph Encoding Layer	Encodes structural topology via graph construction	Formalizes bonding relations and coordination motifs	Long-range interactions and quantum correlations
Representation Learning Layer	Learns embeddings via GNN message passing	Propagates relational features across lattice networks	Dynamical thermodynamic behavior
Latent Embedding Layer	Compresses structures into latent manifolds	Encodes topology similarity and structural clustering	Emergent features and gradients
Inference & Design Layer	Predicts properties / generates candidates	Maps topology to functional outputs	Systemic features and kinetics

A third dynamic within TPDF captures data-model-discovery coupling as $C = g(L_{abstr}, F_{phys})$, where L_{abstr} measures abstraction layers and F_{phys} denotes filtered physics, emphasizing the systemic balance required for robust materials engineering ecosystems.

Analytical implications

The Topology-Physics Decoupling Framework (TPDF) offers interpretive lenses for dissecting the implications of structural abstraction in graph-based materials models, revealing how abstraction layers influence computational workflows and epistemic structures. In materials informatics ecosystems, TPDF highlights the propagation of decoupling effects through data pipelines, where topological prioritization can skew inference toward surface-level patterns, potentially overlooking deeper physical dependencies. This dynamic implies that high-throughput computation, while efficient, may inadvertently reinforce

abstraction biases if not counterbalanced by integrative mechanisms [1, 5]. Key epistemic trade-offs emerging from topology-dominant modeling regimes are synthesized in Table 2.

Table 2. Epistemic Trade-Offs Between Topological Efficiency and Physical Fidelity in Materials AI Systems

Modeling Dimension	Topology-Dominant Advantage	Physics-Grounded Advantage	Resulting Trade-Off
Computational Scalability	Rapid screening across vast chemical spaces	Computationally intensive simulations	Speed vs. mechanistic depth
Representation Expressivity	Efficient encoding of bonding networks	Captures electronic and quantum effects	Structural clarity vs. physical completeness
Inverse Design Capacity	Enables latent space generative exploration	Ensures thermodynamic plausibility	Generative breadth vs. realizability
Uncertainty Quantification	Scalable probabilistic prediction	Physically contextualized confidence	Statistical mechanistic uncertainty
Autonomous Steering	Efficient exploration guidance	Physically constrained navigation	Optimization speed vs. physical alignment

Consider the interaction in representation learning architectures: as graph neural networks abstract topologies, the resulting embeddings facilitate scalable property predictions but introduce trade-offs in interpretability [2, 4, 7-9, 13]. TPDF interprets this as a layered decoupling, where each convolutional step widens the gap between T' (abstracted topology) and P (physical parameters), affecting downstream tasks like uncertainty quantification [18, 19]. For instance, in autonomous discovery systems, this implies steering logics that favor topological exploration over physical validation, leading to infrastructures that excel in volume but lag in fidelity [14, 16, 17].

Epistemic risk structures emerge prominently in inverse materials design, where TPDF elucidates how abstracted graphs map properties to structures with incomplete physical reinstatement. This can manifest as workflow dynamics that amplify uncertainties in disordered materials, necessitating feedback loops to recalibrate [6, 11, 15]. The framework's feedback components suggest implications for multimodal datasets, implying that coupling diverse data modalities could mitigate decoupling by enriching T' with residual physical cues [3, 6].

Computationally, TPDF implies trade-offs in discovery steering, where abstraction enables rapid iteration but constrains the exploration of physically novel spaces [21, 22, 24]. In simulation-experiment coupling, this translates to interpretive insights on alignment challenges, where graph-based abstractions may misalign with empirical feedbacks, prompting adaptive infrastructures [22, 23]. Furthermore, for foundation models in science, TPDF implies a need for layered safeguards to prevent abstraction-induced overgeneralization [3, 15].

A key implication involves the interaction between abstraction and uncertainty, which may be expressed as

$$R = h(A_{level}, U_{prop}), \quad (1)$$

where A_{level} denotes the degree of topological abstraction and U_{prop} captures propagated uncertainties, illustrating how higher abstraction correlates with amplified epistemic risks in AI-guided systems [16, 18, 19]. This formalization underscores systems-level insights into balancing computational efficiency with physical reintegration.

Overall, TPDF's analytical implications extend to epistemic risk management, advocating for discovery logics that incorporate abstraction audits to enhance robustness in materials AI ecosystems [18, 19, 21]. By interpreting these dynamics, the framework guides infrastructure evolution toward more cohesive representation-inference paradigms, fostering integrative advancements without empirical assertions.

Results and Discussion

Structural abstraction as an epistemic design variable

Integrating insights from the Topology-Physics Decoupling Framework (TPDF), the discussion foregrounds structural abstraction not merely as a modeling choice but as a systems-level epistemic design variable shaping the trajectory of computational materials engineering. Graph-based architectures have redefined discovery workflows by enabling scalable encoding of crystalline and molecular systems, embedding relational structures into machine-interpretable manifolds [2, 7-9]. Yet, this transformation carries a dual character. Topological abstraction simultaneously expands computational reach while constraining the representational bandwidth through which physical phenomena are expressed.

From a pipeline perspective, abstraction propagates upstream and downstream. At the infrastructural level, curated datasets optimized for graph ingestion privilege structural connectivity, thereby standardizing materials knowledge in forms amenable to message-passing inference. Downstream, predictive systems inherit these encoded priors, reinforcing structural similarity logics in screening and optimization tasks. The result is an epistemic feedback loop wherein abstraction becomes self-stabilizing—structural descriptors guide inference, and inference success reinforces descriptor dominance.

This duality necessitates reconsideration of decoupling thresholds within high-throughput ecosystems. When abstraction operates without compensatory physical reintegration, epistemic blind spots may accumulate, particularly in domains governed by defects, metastability, or environmental coupling [1, 5]. TPDF thus reframes abstraction not as a binary condition but as a tunable systems parameter requiring infrastructural governance.

Representation–physics tensions in learning architectures

Within representation learning, TPDF reveals a persistent interpretive tension between embedding efficiency and physical completeness. Architectures such as crystal graph convolutional networks operationalize materials as relational graphs, enabling high-dimensional embeddings capable of supporting property prediction and generative design [4, 10-13, 25]. Their computational advantages derive from the compression of geometric and electronic complexity into transferable structural descriptors.

However, this compression introduces representational asymmetries. Dynamic physical interactions—including

lattice vibrations, defect migration, electronic polarization, and temperature-dependent phase behavior—often remain external to topological encodings. Even advanced augmentations such as attention weighting or equivariant message passing operate within discretized structural manifolds, limiting their capacity to internalize field-dependent phenomena.

TPDF interprets this condition as representational decoupling rather than representational failure. The issue is not that graph models misrepresent physics, but that they selectively encode it. This selective encoding suggests the need for hybrid representational logics in which topological embeddings remain computational backbones while physical simulation feedbacks modulate latent spaces. Such hybridization could reshape AI-guided discovery systems by embedding corrective signals directly into representation layers [14, 16, 17].

Uncertainty amplification across abstraction layers

A central implication of TPDF lies in its reinterpretation of uncertainty quantification. In conventional modeling discourse, uncertainty is treated as a statistical property arising from data sparsity, measurement error, or model variance. TPDF extends this view by positioning abstraction itself as an uncertainty amplifier.

Each layer of topology–physics decoupling introduces epistemic distance between representation and material reality. As structural embeddings propagate through predictive architectures, minor abstraction-induced distortions may accumulate, manifesting as confidence inflation or miscalibrated prediction intervals [18, 19]. This layered amplification is particularly consequential in closed-loop discovery environments, where uncertainty estimates guide experimental allocation and candidate prioritization.

Interpretability frameworks emerge here as epistemic bridging mechanisms. Attribution mapping, subgraph saliency, and atoms-in-molecules analyses enable interrogation of model reasoning, revealing where topological inference diverges from physical plausibility. TPDF suggests that uncertainty and interpretability should operate as coupled correctives—one quantifying abstraction risk, the other localizing its structural origin.

Implications for inverse design and generative workflows

Computational design paradigms provide a fertile domain for observing topology–physics decoupling in action. Inverse materials design leverages graph-based generative models to propose candidate structures optimized for targeted functionalities [15, 22]. Within this generative context, abstraction accelerates exploration by enabling traversal of latent chemical spaces unconstrained by explicit physical simulation.

Yet, the same abstraction introduces plausibility risks. Generated candidates may satisfy structural heuristics while occupying thermodynamically unstable or synthetically inaccessible regions of design space [24]. TPDF interprets this phenomenon as generative decoupling—where structural feasibility diverges from physical realizability.

Multimodal and multi-fidelity integrations offer partial mitigation. By embedding hierarchical simulation data and experimental priors into generative training regimes, models can learn stability gradients alongside connectivity rules [6]. However, unless physical constraints are recursively reintegrated into generative loops, abstraction-driven exploration may continue to privilege novelty over realizability.

Foundation models and cross-domain abstraction scaling

The emergence of scientific foundation models extends topology–physics decoupling into broader epistemic territories. These architectures unify structural, chemical, and textual corpora into shared embedding spaces, enabling transfer learning across scientific domains [3, 15]. While such scaling enhances generalization capacity, it also propagates abstraction hierarchies across disciplinary boundaries.

TPDF suggests that decoupling risks scale alongside representational universality. When graph-derived structural embeddings are integrated into multimodal foundation systems, their abstraction assumptions may influence downstream reasoning in property prediction, synthesis planning, and experimental design. Steering mechanisms capable of reinfusing domain-specific physics—without

collapsing cross-domain interoperability—thus become critical infrastructural priorities.

Data infrastructures, bias, and closed-loop alignment

Structural abstraction limits are equally visible within materials data infrastructures. Dataset construction processes often standardize materials knowledge into formats optimized for graph ingestion, privileging ordered crystalline systems while underrepresenting disordered, defect-rich, or metastable phases [3, 6]. These infrastructural biases propagate into model training distributions, shaping inference reliability.

Closed-loop experimentation intensifies these effects. Autonomous discovery platforms iteratively retrain on newly generated data, reinforcing representational priors embedded in earlier abstraction layers [17, 22, 23]. Without adaptive corrective feedbacks, discovery trajectories may narrow, converging on structurally familiar regions rather than expanding epistemic coverage.

TPDF frames this condition as infrastructural decoupling drift—where iterative optimization amplifies abstraction biases over successive learning cycles. Embedding experimental feedback, anomaly detection, and physics-aware recalibration into closed loops becomes essential for maintaining alignment between topological inference and empirical reality.

Toward sustainable abstraction governance

Collectively, these dynamics position TPDF as more than an interpretive lens; it becomes a governance heuristic for sustainable computational ecosystems. Balanced abstraction does not imply abandoning graph architectures or high-throughput infrastructures. Rather, it calls for dynamic management of decoupling thresholds across data, representation, and discovery layers.

Sustainable abstraction governance may involve adaptive descriptor enrichment, physics-informed latent modulation, uncertainty-weighted acquisition, and feedback-coupled generative constraints. Such strategies preserve computational scalability while preventing epistemic erosion. In this sense, TPDF reframes materials AI not as a purely algorithmic enterprise but as an infrastructural ecology requiring systemic calibration.

Conclusion

The Topology-Physics Decoupling Framework (TPDF) provides a conceptual scaffold for interrogating structural abstraction limits in graph-based materials modeling, illuminating how representational design choices propagate across computational discovery ecosystems. By formalizing abstraction layers and their feedback dynamics, the framework reveals how topological decoupling shapes epistemic reliability from data curation through AI-guided inference and generative design.

TPDF underscores that abstraction is neither inherently detrimental nor universally beneficial; its impact is contingent on how effectively physical context is reintegrated across infrastructures and learning architectures. Through this lens, uncertainty quantification, interpretability analytics, and multimodal data fusion emerge as corrective instruments capable of stabilizing decoupled representations.

As computational materials ecosystems continue to evolve—expanding toward autonomous laboratories, foundation models, and cross-domain discovery platforms—frameworks such as TPDF offer interpretive guidance for maintaining epistemic integrity. By advocating integrative strategies that balance structural efficiency with physical fidelity, the framework contributes to the development of resilient, transparent, and sustainable materials informatics paradigms without prescribing rigid methodological mandates.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 15 Apr 2022 Revised: 28 Jun 2022 Accepted: 16 Jul 2022

Published online: 18 September 2022

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

- 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(1):83.
- Fung V, Zhang J, Juarez E, Sumpter BG. Benchmarking graph neural networks for materials chemistry. *npj Comput Mater.* 2021;7(1):84.
- Musil F, Grisafi A, Bartók AP, Ortner C, Csányi G, Ceriotti M. Physics-Inspired structural representations for molecules and materials. *Chem Rev.* 2021;121(16):9759-815.
- Chen C, Ye W, Zuo Y, Zheng C, Ong SP. Graph networks as a universal machine learning framework for molecules and crystals. *Chem Mater.* 2019;31(9):3564-72.
- 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(1):65.
- 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.
- Choudhary K, DeCost B. Atomistic line graph neural network for improved materials property predictions. *npj Comput Mater.* 2021;7(1):185.
- Reiser P, Neubert M, Eberhard A, Torresi L, Zhou C, Shao C, et al. Graph neural networks for materials science and chemistry. *Commun Mater.* 2022;3(1):93.
- Dai M, Demirel MF, Liang Y, Hu J-M. Graph neural networks for an accurate and interpretable prediction of the properties of polycrystalline materials. *npj Comput Mater.* 2021;7(1):103.
- Pettersson L, Verdozzi C, Marques MAL. Crystal graph attention networks for the prediction of stable materials. *Sci Adv.* 2021;7(49):eabi7948.
- Goodall REA, Lee AA. Predicting materials properties without crystal structure: deep representation learning from stoichiometry. *Nat Commun.* 2020;11(1):6280.
- Park CW, Wolverton C. Accurate and scalable graph neural network force field and molecular dynamics with direct force architecture. *npj Comput Mater.* 2021;7(1):73.
- Xie T, Grossman JC. crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys Rev Lett.* 2018;120(14):145301.
- Hatakeyama-Sato K, Oyaizu K. Integrating multiple materials science projects in a single neural network. *Commun Mater.* 2020;1(1):49.
- Chen C, Ong SP. A universal graph deep learning interatomic potential for the periodic table. *Nat Comput Sci.* 2020;2(11):718-28.
- Chen Z, Wei T, Ruzsinszky A, Stillinger FH, Carter EA, Debenedetti PG. Graph dynamical networks for unsupervised learning of atomic scale dynamics in materials. *Nat Mach Intell.* 2019;1(3):158-68.
- Rossi K, Cumby J. Graph representation forecasting of patient's medical conditions: toward a digital twin. *Sci Rep.* 2021;11(1):18423.
- Zhong X, Gallagher B, Liu S, Kailkhura B, Hiszpanski A, Han TY-J. Explainable machine learning in materials science. *npj Comput Mater.* 2022;8(1):204.
- Zubatyuk R, Smith JS, Leszczynski J, Isayev O. Accurate and transferable multitask prediction of chemical properties with an atoms-in-molecules neural network. *Sci Adv.* 2019;5(8):eaav6490.
- Louis S-Y, Zhao Y, Nasiri A, Wang X, Song Y, Liu F, et al. Graph convolutional neural networks with global attention for

improved materials property prediction. *Phys Chem Chem Phys.* 2020;22(32):18141-8.

Xie T, Grossman JC. Hierarchical visualization of materials space with graph convolutional neural networks. *J Chem Phys.* 2018;149(17):174107.

Yan K, Liu Q, Yan Z, Wang X, Wu B, Zhu D, et al. Organic reaction mechanism classification using machine learning. *Nature.* 2021;598(7881):451-6.

Jorgensen PB, Jacobsen KW, Schmidt MN. Neural message passing with edge updates for predicting properties of

molecules and materials. arXiv preprint arXiv:1806.03146. 2018.

Karamad M, Magar R, Shi Y, Siahrostami S, Gates ID, Farhad S. Orbital graph convolutional neural network for material property prediction. *Phys Rev Mater.* 2020;4(9):093801.

Schütt KT, Sauceda HE, Kindermans P-J, Tkatchenko A, Müller K-R. SchNet – A deep learning architecture for molecules and materials. *J Chem Theory Comput.* 2018;14(6):3151-61.