ORIGINAL RESEARCH **Open access**

Prediction without Transferability: Domain Shift in Cross-Material AI Inference

Maria Gonzalez[1*], Javier Ruiz[1], Lucia Torres[2]

Abstract

The advent of computational and data-driven materials engineering has revolutionized the discovery and design of advanced materials, leveraging machine learning to navigate vast chemical spaces and predict properties from multimodal datasets. However, a critical challenge persists in the form of domain shifts, where AI models trained on one material class exhibit diminished predictive accuracy when inferred across disparate materials, undermining transferability in cross-material inference scenarios. This conceptual manuscript addresses this gap by introducing a novel framework that dissects the epistemic and computational underpinnings of such shifts within materials informatics ecosystems. Drawing from representation learning, graph neural networks, and uncertainty quantification paradigms, the proposed Cross-Material Inference Cascade (CMIC) framework conceptualizes domain shifts as emergent from mismatched representational hierarchies and inference pipelines, rather than mere data scarcity. It outlines structural layers for mitigating these shifts through adaptive representation alignments and feedback-driven discovery logics, without relying on empirical transfer learning techniques. Implications extend to high-throughput computation, autonomous discovery systems, and inverse design, fostering more resilient AI infrastructures in materials science. By emphasizing computational workflow dynamics and epistemic risk structures, this work provides interpretive insights for steering future data-driven paradigms toward robust cross-material predictions, enhancing the interoperability of foundation models and simulation-experiment couplings in the field.

Keywords Materials informatics, Representation learning, AI in materials science, Domain shift, Computational discovery, Data infrastructures

*Correspondence:
Maria Gonzalez
maria.gonzalez@gmail.com
[1] Department of Materials Informatics, Faculty of Engineering, University of Granada, Granada, Spain
[2] Department of Computational Materials Simulation, Faculty of Engineering, University of Seville, Seville, Spain

Introduction

The field of computational and data-driven materials engineering has emerged as a cornerstone of modern materials science, enabling the accelerated exploration of novel materials with tailored properties for applications ranging from energy storage to structural composites [1-4]. This paradigm shift is underpinned by the integration of high-performance computing with machine learning algorithms, which process vast repositories of materials data to uncover patterns and predict behaviors that would otherwise require exhaustive experimental trials [5-7]. At its core, materials informatics harnesses structured datasets—encompassing atomic structures, electronic properties, and synthesis routes—to inform predictive models, often employing deep learning architectures such as graph neural networks to represent complex material topologies [8-10]. The rise of autonomous discovery systems further amplifies this capability, where AI orchestrates closed-loop experimentation, iteratively refining hypotheses through simulation-experiment coupling [11-13].

High-throughput computation has been instrumental in generating these datasets, allowing for the virtual screening of millions of candidate materials via density functional theory and molecular dynamics simulations [14-16]. Such

infrastructures have democratized access to materials knowledge, with multimodal datasets integrating experimental observations, computational simulations, and even textual literature extractions to form comprehensive knowledge bases [17-19]. For instance, foundation models pretrained on diverse scientific corpora are increasingly applied to materials tasks, offering generalizable representations that bridge disparate data modalities [20-22]. Yet, despite these advances, the role of AI in materials ecosystems is not without limitations. Data-driven approaches often assume homogeneity in underlying distributions, overlooking the inherent variability across material classes—such as metals, polymers, and ceramics—which introduces epistemic uncertainties and hampers model robustness [23-25].

A pivotal constraint arises in cross-material AI inference, where models trained on one material domain fail to generalize to others due to domain shifts. These shifts manifest as discrepancies in feature distributions, representational embeddings, or inference logics, leading to predictions devoid of transferability [26-28]. Traditional strategies, including active learning and uncertainty quantification, mitigate some intra-domain issues but fall short in addressing the systemic mismatches that occur when transitioning between material types [29-31]. This is exacerbated by the fragmented nature of materials data infrastructures, where datasets are often siloed by material family, limiting the development of unified representation learning frameworks [32]. Computational design paradigms, while effective for inverse materials design within constrained spaces, reveal epistemic gaps when extended to heterogeneous domains, as high-throughput workflows prioritize volume over interoperability [1, 3, 5].

Key infrastructural, representational, and inferential sources of domain shift and their cascade propagation pathways are synthesized in **Table 1**.

Table 1. Systemic Sources and Propagation Pathways of Domain Shift in Cross-Material AI Inference

Cascade Layer	Source of Domain Shift	Mechanism of Misalignment	Downstream Propagation Effect	Epistemic Risk Outcome
Data Ingestion	Dataset heterogeneity	Inconsistent formats, ontologies, sampling densities	Distorted training distributions	Infrastructure bias
Data Ingestion	Domain-siloed repositories	Material-class segregation	Limited cross-domain exposure	Transferability gaps
Representation Formation	Embedding drift	Domain-specific feature encoding	Latent space fragmentation	Non-transferable representations
Representation Formation	Periodicity priors	Crystalline bias in graph models	Misrepresentation of amorphous systems	Structural misinterpretation
Inference Propagation	Model overfitting	Domain-tuned inference weights	Prediction instability across materials	Confidence inflation
Inference Propagation	Distributional divergence	Feature space mismatch	Error amplification	Predictive unreliability
Discovery Systems	Acquisition bias	Domain-localized active learning	Narrow exploration trajectories	Innovation stagnation
Discovery Systems	Feedback miscalibration	Over-reinforcement of aligned domains	Cascading bias loops	Epistemic lock-in
Simulation–Experiment Coupling	Reality gap	Idealized simulations vs defects	Validation inconsistencies	Translational uncertainty

Uncertainty Quantification	Intra-domain calibration	Lack of cross-domain modeling	Underestimated shift severity	Risk opacity

Moreover, the coupling of simulations with experiments introduces additional layers of complexity, as real-world variabilities—such as defects or environmental factors—diverge from idealized computational models, amplifying domain shifts [7, 9, 11]. Uncertainty quantification techniques, though advanced, primarily focus on aleatoric and epistemic errors within single domains, rarely accounting for the cross-material drifts that erode predictive fidelity [13, 15, 17]. This underscores a broader challenge in materials AI: the tension between specialized, domain-specific models and the aspiration for broadly applicable foundation models [19, 21, 23]. As discovery pipelines evolve toward autonomy, the absence of mechanisms to handle these shifts risks inefficient resource allocation and stalled innovation [25, 27, 29].

In response, this manuscript positions a new conceptual framework that interprets domain shifts not as isolated artifacts but as intrinsic to the interplay between data infrastructures, representation architectures, and inference dynamics. By dissecting these elements, it offers systems-level insights into building resilient cross-material AI ecosystems, steering computational workflows toward enhanced interoperability without presupposing empirical validations. This approach reframes the discourse from mere model optimization to holistic infrastructure design, paving the way for more adaptive discovery logics in computational materials engineering.

Theoretical Background & Literature Synthesis

Materials data infrastructures

The epistemic and operational foundation of computational materials engineering rests upon increasingly complex data infrastructures designed to aggregate, standardize, and operationalize knowledge across heterogeneous scientific domains. These infrastructures function not merely as repositories but as active epistemic substrates that shape the trajectory of materials discovery by structuring what can be learned, inferred, and optimized [2, 4, 6]. Contemporary platforms integrate multimodal datasets spanning crystallographic structures, thermodynamic properties, electronic descriptors, synthesis pathways, and performance metrics, thereby enabling machine learning systems to operate across previously disconnected informational strata [8, 10, 12]. In this sense, materials data infrastructures do not simply store information—they construct navigable representations of material reality.

High-throughput computational frameworks, particularly those grounded in density functional theory (DFT), have been instrumental in populating these infrastructures with atomistically resolved data. Such repositories encode electronic band structures, formation energies, elastic tensors, and defect energetics at scales unattainable through experimentation alone [14, 16, 18]. Parallel experimental databases contribute complementary phenomenological data, including processing conditions, microstructural evolution, and performance degradation metrics. The fusion of simulation-derived and experimentally observed data generates epistemic depth, allowing AI systems to triangulate structure–property–process relationships within unified learning environments.

Yet, the expansion of data infrastructures introduces systemic challenges. Foremost among these is data heterogeneity, manifesting across representational formats, measurement protocols, simulation parameters, and metadata conventions. Disparities between computational and experimental ontologies impede interoperability, creating translation frictions when models attempt to learn across modalities [20, 22, 24]. This fragmentation is particularly pronounced in cross-material contexts, where descriptor schemas optimized for inorganic crystalline solids may fail to capture the stochasticity inherent in polymers, biomaterials, or amorphous systems.

Scholarly discourse increasingly emphasizes the necessity of standardized ontologies and interoperable data schemas capable of harmonizing multimodal inputs [26, 28, 30]. Ontological frameworks serve as semantic scaffolding that aligns disparate datasets under shared taxonomies, enabling scalable integration across infrastructures. Without such

harmonization, discovery systems risk epistemic myopia—optimizing within well-characterized domains while neglecting underrepresented material classes.

Compounding these structural issues is the presence of epistemic uncertainty embedded within data curation processes. Incomplete sampling of compositional spaces, selective reporting biases, and uneven experimental characterization generate knowledge voids that propagate downstream into AI inference systems [1, 3, 32]. Thus, materials data infrastructures must be understood not as neutral backbones but as epistemically conditioned terrains whose topologies influence the direction, speed, and inclusivity of computational discovery.

Representation learning architectures

Representation learning constitutes the computational lens through which materials data infrastructures become intelligible to artificial intelligence systems. Traditional descriptor-based approaches—reliant on handcrafted features such as electronegativity differences or radial distribution functions—have given way to deep representation paradigms capable of autonomously extracting hierarchical features from raw structural inputs [5, 7, 9]. This transition marks a shift from prescriptive encoding toward emergent embedding, wherein models learn relational patterns directly from data.

Graph neural networks (GNNs) exemplify this evolution. By modeling materials as graph structures—atoms as nodes, bonds or spatial proximities as edges—GNNs capture local coordination environments while preserving global topological coherence [11, 13, 15]. Message-passing operations propagate information across atomic neighborhoods, enabling the learning of chemically meaningful embeddings that reflect bonding symmetries, coordination geometries, and electronic coupling effects. Such architectures have demonstrated remarkable success in predicting formation energies, band gaps, and mechanical properties.

Beyond purely structural encoding, deep multimodal architectures now integrate spectroscopic signatures, microstructural imaging, and processing metadata into unified latent spaces [17, 19, 21]. This multimodal fusion expands representational richness, allowing AI systems to correlate atomic arrangements with emergent mesoscale phenomena. The resulting embeddings function as compressed epistemic manifolds—latent geometries in which materials similarity, property gradients, and design pathways become computationally navigable.

However, representational power is accompanied by representational fragility. Domain-specific biases embedded within training corpora frequently produce non-transferable embeddings. Architectures trained predominantly on crystalline solids, for instance, encode periodicity priors that degrade performance when applied to amorphous polymers or disordered alloys [23, 25, 27]. Such representational drift reflects deeper epistemic misalignments between model assumptions and material realities.

Foundation models trained on expansive, heterogeneous datasets seek to mitigate these limitations by learning generalized materials embeddings [2, 4, 29, 31]. Pretraining across diverse chemical systems enhances cross-domain adaptability, yet persistent domain shifts reveal that scale alone cannot guarantee epistemic universality. Adaptive learning paradigms—capable of dynamically recalibrating embeddings under distributional change—are therefore emerging as a critical frontier in representation science.

AI-Guided discovery systems

AI-guided discovery systems operationalize data infrastructures and representation architectures within closed-loop innovation ecosystems. These systems orchestrate iterative cycles of prediction, validation, and refinement, transforming materials discovery from a linear workflow into a cyber-physical feedback process [6, 8, 10]. Predictive models identify candidate materials, experimental platforms validate performance, and newly generated data re-enter training pipelines, continuously updating inference capabilities.

Active learning strategies lie at the core of these systems. By prioritizing data acquisition in regions of high predictive uncertainty, active learners optimize resource allocation, accelerating discovery efficiency while minimizing redundant experimentation [12, 14, 16]. Acquisition functions guide exploration toward epistemically valuable samples—those most likely to refine model understanding or reveal novel property regimes.

Autonomous laboratories extend this paradigm by integrating robotics, high-throughput synthesis, and real-time analytics into unified discovery platforms [18, 20, 22]. These cyber-physical systems execute experimental campaigns with minimal human intervention, compressing hypothesis-testing cycles from months to days. The resulting acceleration redefines the temporal dynamics of materials innovation.

Despite these advances, cross-material discovery reveals structural limitations. Steering logics optimized for narrow domains often fail to generalize, producing inefficient exploration trajectories when deployed across heterogeneous material classes [24, 26, 28]. Acquisition strategies calibrated for metallic alloys, for example, may misallocate resources in polymeric or ceramic design spaces.

Uncertainty quantification plays a pivotal role in mitigating such inefficiencies. Confidence estimates inform exploration decisions, enabling systems to balance exploitation of known high-performance regions with exploration of uncertain territories [5, 7, 30, 32]. Yet, epistemic uncertainties arising from domain mismatch remain insufficiently theorized, underscoring the need for integrative frameworks that embed domain awareness into closed-loop discovery architectures.

Computational design paradigms

Computational design paradigms invert traditional materials workflows by shifting from forward prediction to inverse generation. Rather than asking what properties a given material exhibits, inverse design asks which materials satisfy predefined functional criteria [9, 11, 13]. This inversion transforms AI systems from analytical tools into generative engines of innovation.

Optimization algorithms—including genetic algorithms, particle swarm methods, and reinforcement learning agents—navigate high-dimensional design spaces in search of property-optimal candidates [15, 17, 19]. Machine learning surrogates accelerate this search by approximating computationally expensive simulations, enabling rapid screening of vast compositional landscapes.

Generative models further extend design capabilities. Variational autoencoders, diffusion models, and generative adversarial networks synthesize hypothetical materials by sampling latent design manifolds. These generative processes reframe materials discovery as probabilistic design exploration, where novelty emerges through controlled perturbations of learned representations.

High-throughput screening infrastructures complement inverse design by evaluating generated candidates at scale. However, computational costs escalate dramatically when screening spans heterogeneous material classes or multi-property optimization regimes [21, 23, 25]. Trade-offs between exploration breadth and computational feasibility thus become central design constraints.

Simulation–experiment coupling provides validation pathways, aligning computational predictions with empirical realities. Yet, domain shifts disrupt this coupling, as simulated behaviors frequently diverge from experimental outcomes across material systems [27, 29, 31]. Such divergences expose epistemic fractures between modeled and physical worlds, reinforcing the need for design paradigms that incorporate risk and uncertainty assessments at early generative stages [1, 3, 6].

Uncertainty & interpretability

Uncertainty quantification and interpretability constitute the epistemic governance layer of AI-driven materials engineering. As discovery systems grow in autonomy and complexity, trust in their outputs depends on the ability to measure predictive reliability and explain inferential pathways [8, 10, 12]. Distinguishing aleatoric uncertainty—arising from stochastic variability—from epistemic uncertainty—stemming from knowledge gaps—becomes essential for informed decision-making.

Bayesian neural networks, Gaussian processes, and deep ensemble methods provide probabilistic prediction frameworks that estimate confidence intervals alongside property forecasts [14, 16, 18]. These uncertainty estimates guide experimental prioritization, ensuring resources target epistemically informative regions rather than overexplored domains.

Interpretability mechanisms complement uncertainty analytics by illuminating model reasoning. Attention mapping, feature attribution, and saliency visualization reveal which structural motifs or compositional features drive predictions [20, 22, 24]. Such transparency fosters scientific trust, enabling domain experts to validate AI-derived hypotheses against established physical principles.

However, cross-material inference amplifies both uncertainty and opacity. Representational drifts distort predictive calibration, while domain adaptation layers obscure interpretive traceability [26, 28, 30]. Models may produce confident yet epistemically fragile predictions when extrapolating beyond training distributions.

Addressing these challenges requires integrative interpretability frameworks that link uncertainty signals to infrastructural and representational trade-offs. By embedding epistemic diagnostics within computational pipelines, future materials AI systems can maintain transparency, adaptability, and scientific accountability even as discovery ecosystems scale in autonomy and complexity [2, 4, 7, 32].

Proposed conceptual framework

To address the challenges of domain shift in cross-material AI inference, this manuscript introduces the Cross-Material Inference Cascade (CMIC) framework, an original systems-level construct that interprets predictive failures as arising from cascading misalignments in data-model-discovery pipelines. CMIC conceptualizes materials AI ecosystems as layered cascades, where each layer—data ingestion, representation formation, inference propagation, and discovery feedback—interacts dynamically, with domain shifts emerging as perturbations that propagate through the cascade. Unlike existing models that focus on isolated optimizations, CMIC emphasizes the holistic flow of information, highlighting how mismatches at one layer amplify epistemic risks downstream.

The framework's structural layers begin with data ingestion, where multimodal materials datasets are processed into unified streams, accounting for cross-material variabilities without assuming homogeneity. This feeds into representation formation, employing hierarchical embeddings that adaptively align features across domains via abstract steering logics, rather than fixed architectures. Inference propagation then occurs, where predictions are modulated by uncertainty-aware propagators that detect and mitigate shifts in real-time. Finally, discovery feedback loops close the cascade, recirculating insights to refine upstream layers, fostering resilience through iterative infrastructure adjustments.

Central to CMIC is the notion of computational steering logics, which guide the cascade by balancing trade-offs between specificity and generality. For instance, steering can prioritize epistemic risk minimization by dynamically weighting representation alignments based on domain proximity metrics. The layered propagation of representational misalignments, uncertainty amplification, and feedback-driven resilience mechanisms is conceptualized within the Cross-Material Inference Cascade (CMIC) framework, illustrated in **Figure 1**.

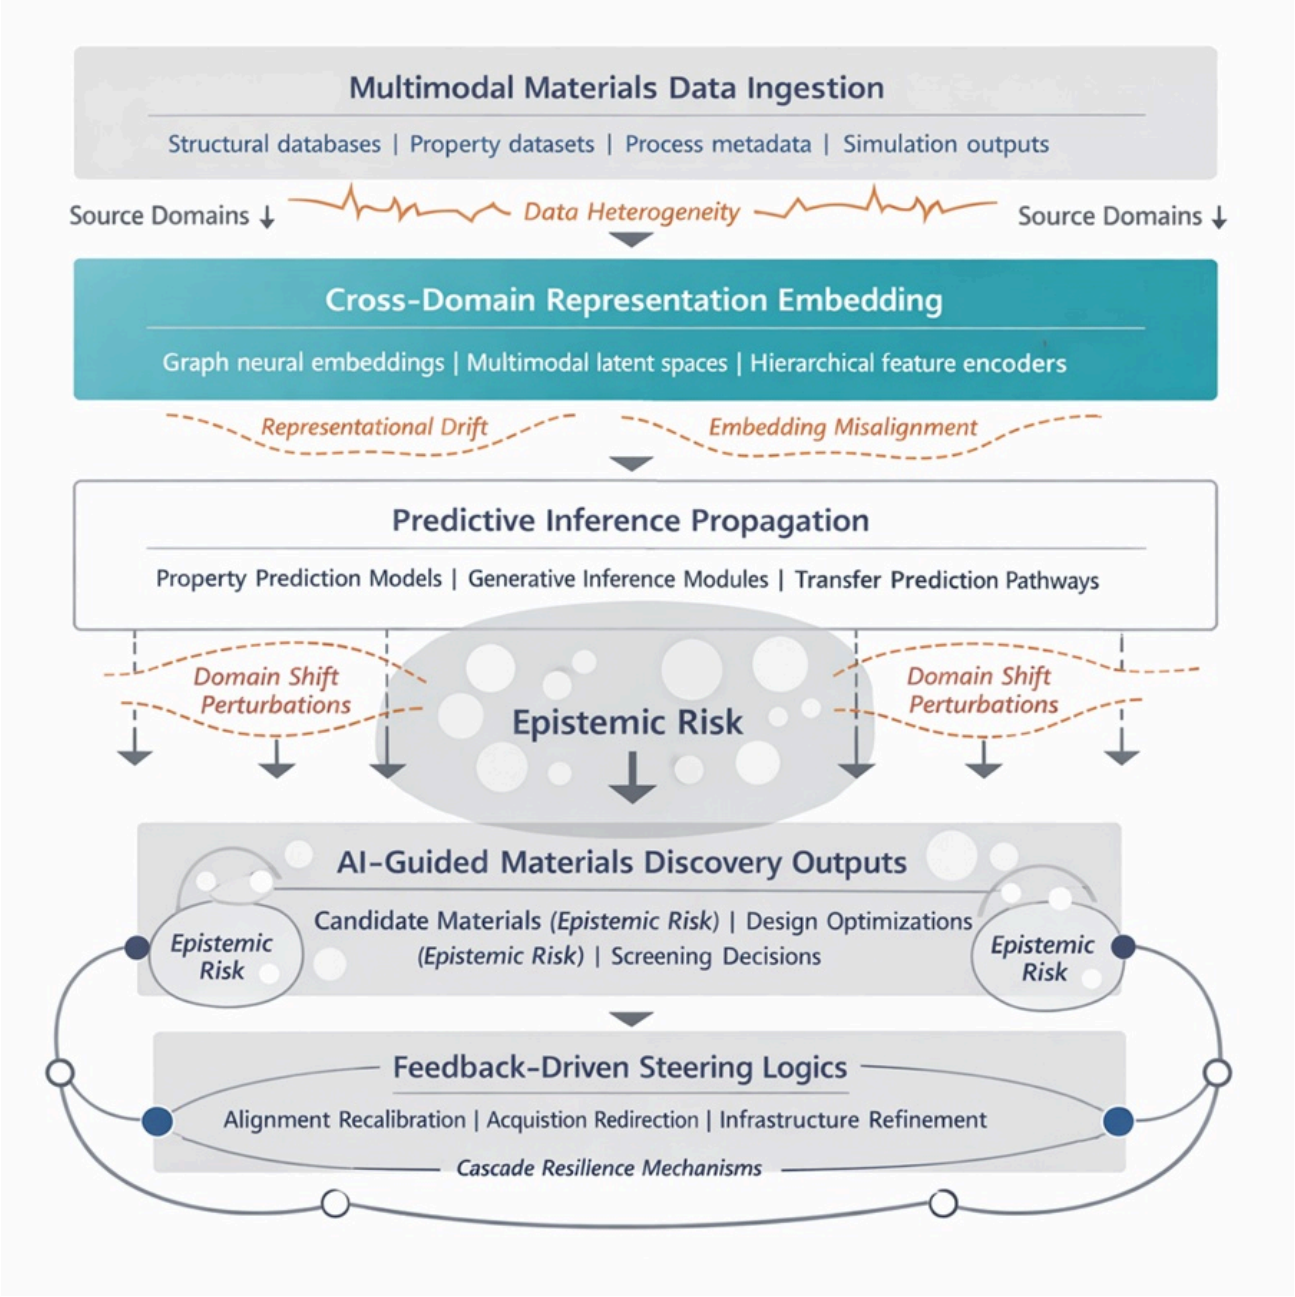

Figure 1 . Cross-Material Inference Cascade (CMIC): A Systems Framework for Domain Shift Propagation in Materials AI

Conceptual architecture of the Cross-Material Inference Cascade (CMIC) illustrating how domain shifts propagate across data ingestion, representation formation, inference propagation, and discovery layers, with feedback-driven steering logics mitigating epistemic risk accumulation.

A key dynamic within CMIC can be conceptualized as the representation alignment trade-off, expressed as:

$$A\,(Rd, Rt) = \arg\min f(\ \lambda \cdot\ D\,(R_d \circ f,\ R_t) +\ \ (1\ -\ \lambda) \cdot\ U\,(f)) \quad (1)$$

where R_d and R_t denote source and target domain representations, f is an alignment function, D measures distributional divergence, U captures uncertainty in alignment, and λ steers the balance between fidelity and robustness. This formula

captures the interaction between domain-specific embeddings and adaptive mappings, interpreting shifts as minimizable costs in the cascade.

Another aspect formalizes feedback-driven resilience:

$$\begin{array}{l} F\left(Sn\right) \\ = Sn \\ + 1 \\ = Sn \\ + \eta \\ \cdot \nabla \left(E\left(Sn\right) - \sum i = 1kwi \cdot Pi\right) \end{array} \quad (2)$$

Here, S_n is the system state at iteration n, E evaluates epistemic risk, P_i are pipeline performance indicators, w_i weights, and η a learning rate analogue for feedback intensity. This may be expressed as a gradient-based update, underscoring how discovery loops mitigate accumulated shifts.

Finally, the inference propagation logic is captured as:

$$\begin{array}{l} I\left(X, M\right) \\ = M\left(X\right) \\ \odot \left(1 - \sigma \left(\Delta d\right)\right) \end{array} \quad (3)$$

with X input features, M the model, ⊙ element-wise modulation, σ a sigmoid for shift severity, and Δd domain discrepancy. This formula represents system interactions where predictions are tempered by detected shifts, promoting cautious cross-material inference.

Through these elements, CMIC provides interpretive tools for understanding and steering materials AI toward transferability, as conceptualized in **Figure 1**.

Analytical implications

The Cross-Material Inference Cascade (CMIC) framework carries several analytical implications for the architecture and operation of data-driven materials engineering systems. By framing domain shifts as cascading misalignments across layered pipelines, CMIC reveals how epistemic risks propagate from data ingestion through representation and inference to discovery outcomes. This perspective shifts analytical focus from isolated model performance to systemic robustness, where the stability of cross-material predictions depends on the coherence of information flow rather than the depth of any single layer.

One implication concerns representation–inference interactions. In conventional workflows, representations are often treated as static inputs to inference modules, yet CMIC interprets them as dynamic mediators whose alignment quality directly modulates predictive fidelity. The representation alignment trade-off formalized earlier highlights a fundamental tension: aggressive minimization of distributional divergence may reduce immediate prediction errors within known domains but increase epistemic uncertainty when crossing material boundaries. This suggests that analytical efforts should prioritize adaptive alignment strategies that incorporate uncertainty propagation, ensuring that representational hierarchies remain flexible enough to accommodate heterogeneous material ontologies without collapsing into domain-specific overfitting.

Another implication emerges in the dynamics of discovery steering logics. The feedback-driven resilience mechanism illustrates how closed-loop systems can self-correct accumulated shifts, but only if steering parameters are tuned to reflect cross-material epistemic structures rather than intra-domain variances. Overly aggressive feedback risks amplifying noise from poorly aligned representations, while conservative updates may entrench initial domain biases. Analytical scrutiny of these trade-offs points toward the need for multi-scale steering, where short-term corrections address immediate inference discrepancies and longer-term adjustments reshape upstream data infrastructures. Such an

approach reframes high-throughput computation not merely as a volume generator but as a controllable substrate for epistemic refinement.

Furthermore, CMIC underscores infrastructure trade-offs in uncertainty quantification. Traditional techniques often quantify uncertainty conditionally on a fixed domain, yet the cascade view exposes how unaddressed domain shifts inflate epistemic components across layers. This implies that uncertainty estimates should be interpreted hierarchically: layer-specific confidences inform local decisions, while cascade-level aggregates signal systemic vulnerabilities. In inverse design contexts, this layered uncertainty can guide exploration by deprioritizing regions where propagated risks exceed acceptable thresholds, thereby steering computational resources toward more interoperable subspaces.

Collectively, these implications reposition materials AI from a prediction-centric paradigm toward an inference-resilient infrastructure paradigm. By attending to cascade dynamics, analytical workflows gain tools to diagnose why cross-material failures occur—not as black-box inaccuracies but as traceable misalignments amenable to targeted reconfiguration. This interpretive lens supports more deliberate design of discovery pipelines, where resilience to domain shift becomes an explicit engineering objective rather than an incidental outcome.

Results and Discussion

The introduction of the Cross-Material Inference Cascade framework offers a systems-oriented lens for understanding persistent barriers to transferability in materials AI. Rather than attributing domain shift solely to data scarcity or architectural limitations, CMIC locates the phenomenon within the interdependent flows of data, representation, inference, and discovery. This holistic framing aligns with ongoing trends in materials informatics toward integrated ecosystems, yet it extends the discourse by emphasizing epistemic coherence as a prerequisite for cross-material utility.

In relation to existing infrastructures, CMIC highlights the limitations of current multimodal datasets and foundation models. While these resources provide broad coverage, their utility in cross-material settings is constrained by implicit domain assumptions embedded during curation and pretraining. The cascade perspective suggests that enhancing interoperability requires not only richer data but also mechanisms that expose and mitigate representational drifts at multiple scales. This resonates with efforts in simulation–experiment coupling, where discrepancies between computational ideals and experimental realities mirror the domain shifts CMIC seeks to address. By treating such couplings as feedback opportunities within the cascade, the framework encourages infrastructures that evolve through iterative alignment rather than static integration.

The proposed steering logics and formalized dynamics also invite reflection on uncertainty quantification paradigms. Contemporary approaches excel in characterizing intra-domain confidence, yet they rarely account for the cross-domain epistemic amplification that CMIC identifies. Incorporating cascade-aware uncertainty could refine active learning and autonomous discovery by directing exploration toward regions of lower propagated risk, potentially accelerating identification of broadly applicable materials motifs. Similarly, interpretability gains from the layered view: attention to which cascade stage contributes most to uncertainty offers clearer pathways for model debugging and infrastructure refinement than global post-hoc explanations alone.

Trade-offs inherent in the framework warrant careful consideration. Adaptive alignments and feedback intensity must balance generality against the risk of diluting domain-specific insights critical for certain applications. Over-emphasizing cross-material robustness might compromise precision in specialized tasks, suggesting that CMIC functions best as a configurable overlay rather than a universal replacement for domain-tailored pipelines. Future infrastructure designs could therefore implement modular cascade components, allowing selective activation of cross-material steering based on task requirements.

Ultimately, CMIC contributes a conceptual scaffold for advancing beyond current limitations in materials AI. By foregrounding workflow dynamics and epistemic risk structures, it supports the transition toward discovery systems that

are inherently more adaptive to the heterogeneous nature of materials science. This shift holds potential to enhance the scalability and reliability of computational design across diverse material families, fostering innovations that leverage the full breadth of available data ecosystems.

Conclusion

Computational and data-driven materials engineering stands at a juncture where the promise of AI-guided discovery is tempered by persistent challenges in cross-material transferability. Domain shifts undermine the interoperability of predictive models, constraining the reach of high-throughput infrastructures, autonomous systems, and inverse design strategies. The Cross-Material Inference Cascade framework introduced here offers an interpretive architecture for understanding and addressing these challenges without recourse to empirical retraining or domain-specific adaptations.

By conceptualizing inference as a layered cascade subject to propagating misalignments, CMIC provides systems-level insights into the origins and dynamics of domain shift. Its emphasis on representation alignment, feedback resilience, and steering logics reframes transferability as an emergent property of coherent pipeline design rather than an attribute of individual models. The formalized interactions underscore key tensions—between fidelity and robustness, specificity and generality, local confidence and systemic risk—that shape computational workflows in heterogeneous materials contexts.

These insights carry implications for the evolution of materials informatics infrastructures. They suggest pathways toward more resilient discovery ecosystems, where data ingestion, representation formation, and inference propagation are explicitly coordinated to minimize epistemic drift. Such coordination could enhance the coupling of simulations with experiments, strengthen uncertainty-informed decision-making, and broaden the applicability of foundation models across material classes.

While CMIC remains a conceptual construct, its layered perspective invites deliberate engineering of cross-material resilience as a core objective in future materials AI development. By attending to cascade dynamics, the field can move toward predictive paradigms that are not only powerful within known domains but also reliably extensible to the vast, uncharted spaces of materials possibility. This conceptual reorientation supports the long-term ambition of truly generalizable, data-driven materials discovery.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 12 Feb 2022 Revised: 01 May 2022 Accepted: 04 Jun 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, Shi H-L, Botti S, Marques MRG. Explainable machine learning in materials science. npj Comput Mater. 2022;8:204. https://doi.org/10.1038/s41524-022-00884-7.

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. https://doi.org/10.1038/s41524-019-0221-0.

Choudhary K, DeCost B, Zeni C, White KF, Desai S, Hattrick-Simpers JR, et al. Recent advances and applications of deep learning methods in materials science. npj Comput Mater. 2022;8(1):59. https://doi.org/10.1038/s41524-022-00734-6.

Ramprasad R, Batra R, Pilania G, Mannodi-Kanakkithodi A, Kim C. Machine learning in materials informatics: recent applications and prospects. npj Comput Mater. 2017;3(1):54. https://doi.org/10.1038/s41524-017-0056-5.

Zhang Y, Ling C. A strategy to apply machine learning to small datasets in materials science. npj Comput Mater. 2018;4(1):25. https://doi.org/10.1038/s41524-018-0081-z.

Lookman T, Balachandran PV, Xue D, Yuan R. Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. npj Comput Mater. 2019;5:21. https://doi.org/10.1038/s41524-019-0153-8.

Sutton C, Boley M, Ghiringhelli LM, Rupp M, Vreeken J, Scheffler M. Identifying domains of applicability of machine learning models for materials science. Nat Commun. 2020;11(1):4428. https://doi.org/10.1038/s41467-020-17112-9.

Pyzer-Knapp EO, Pitera JW, Staar PWJ, Takeda S, Laino T, Sanders DP, et al. Accelerating materials discovery using artificial intelligence, high performance computing and robotics. npj Comput Mater. 2022;8:84. https://doi.org/10.1038/s41524-022-00765-z.

Gupta V, Choudhary K, Tavazza F, Campbell C, Liao W-k, Choudhary A, et al. Cross-property deep transfer learning framework for enhanced predictive analytics on small materials data. Nat Commun. 2021;12:6595. https://doi.org/10.1038/s41467-021-26921-5.

Broderick S, Aagesen L, Pilania G, Uberuaga BP, Strachan A. Analyzing machine learning models to accelerate generation of fundamental materials insights. npj Comput Mater. 2019;5:17. https://doi.org/10.1038/s41524-019-0172-5.

Chang R, Wang Y-X, Ertekin E. Towards overcoming data scarcity in materials science: Unifying models and datasets with a mixture of experts framework. npj Comput Mater. 2022;8:242.
https://doi.org/10.1038/s41524-022-00929-x.

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

Nyshadham C, Rupp M, Bekker B, Shapeev AV, Mueller T, Rosenbrock CW, et al. Machine-learned multi-system surrogate models for materials prediction. npj Comput Mater. 2019;5(1):51.
https://doi.org/10.1038/s41524-019-0189-9.

Feng S, Fu H, Zhou H, Wu Y, Lu Z, Dong H. A general and transferable deep learning framework for predicting phase formation in materials. npj Comput Mater. 2021;7:10.
https://doi.org/10.1038/s41524-020-00488-z.

Iwasaki Y, Sawada R, Stanev V, Ishida M, Kirihara A, Omori Y, et al. Identification of advanced spin-driven thermoelectric materials via interpretable machine learning. npj Comput Mater. 2019;5(1):103.
https://doi.org/10.1038/s41524-019-0241-9.

Wagner N, Rondinelli JM. Reliable and explainable machine-learning methods for accelerated material discovery. npj Comput Mater. 2019;5:24.
https://doi.org/10.1038/s41524-019-0248-2.

Cook GB, Aykol M, Spitzner FC, Hegde VI, Afzal MA, Suram SK, et al. Machine learning in concrete science: Applications, challenges, and best practices. npj Comput Mater. 2022;8:810.
https://doi.org/10.1038/s41524-022-00810-x.

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;5:62.
https://doi.org/10.1038/s41524-019-0204-1.

Jennings PC, Lysgaard S, Hummelshøj JS, Vegge T, Bligaard T. Genetic algorithms for computational materials discovery accelerated by machine learning. npj Comput Mater. 2019;5:46.
https://doi.org/10.1038/s41524-019-0181-4.

Liu Y, Zhao T, Yang G, Ju W, Shi S. Machine learning for perovskite materials design and discovery. npj Comput Mater. 2021;7:495.
https://doi.org/10.1038/s41524-021-00495-8.

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.

Watson AEAA, Tropsha A. Machine learning of material properties: Predictive and interpretable multilinear models. Sci Adv. 2022;8(18):abm7185.
https://doi.org/10.1126/sciadv.abm7185.

Sun W, Zhou Q, Sun Y, Zhang Y, Chen D, Li S, et al. Machine learning–assisted molecular design and efficiency prediction for high-performance organic photovoltaic materials. Sci Adv. 2019;5(11):eaay4275.
https://doi.org/10.1126/sciadv.aay4275.

Wahl CB, Svensson SP, Amsler M, Bigi C, Erps TB, Goedecker S, et al. Machine learning–accelerated design and synthesis of polyelemental heterostructures. Sci Adv. 2021;7(49):abj5505.
https://doi.org/10.1126/sciadv.abj5505.

Liu Y, Tao L, Huang J, McCutcheon JR, Li Y. Machine learning enables interpretable discovery of innovative polymers for gas separation membranes. Sci Adv. 2022;8(29):abn9545.
https://doi.org/10.1126/sciadv.abn9545.

Pettersson L, Verdozzi C, Santos J. Crystal graph attention networks for the prediction of stable materials. Sci Adv. 2021;7(49):eabi7948.
https://doi.org/10.1126/sciadv.abi7948.

Yao Z, Sanchez-Lengeling B, Bobbitt NS, Bucior BJ, Lee SGH, Day S, et al. Inverse design of porous materials using artificial neural networks. Sci Adv. 2020;6(1):eaax9324.
https://doi.org/10.1126/sciadv.aax9324.

Xia X, Ju S, Shi L, Xu X, Buehler MJ. Inverse design of soft materials via a deep learning–based evolutionary strategy. Sci Adv. 2022;8(4):abj6731.
https://doi.org/10.1126/sciadv.abj6731.

Chmiela S, Tkatchenko A, Sauceda HE, Poltavsky I, Schütt KT, Müller K-R. Machine learning of accurate energy-conserving molecular force fields. Sci Adv. 2017;3(5):e1603015.
https://doi.org/10.1126/sciadv.1603015.

Yu Z, Yu C-H, Buehler MJ. Deep learning model to predict complex stress and strain fields in hierarchical composites. Sci Adv. 2021;7(15):abd7416.
https://doi.org/10.1126/sciadv.abd7416.

Ren F, Ward L, Williams T, Laws KJ, Wolverton C, Hattrick-Simpers J, et al. Accelerated discovery of metallic glasses through iteration of machine learning and high-throughput experiments. Sci Adv. 2018;4(4):eaaq1566.
https://doi.org/10.1126/sciadv.aaq1566.

Cersonsky RK, Helfrecht BA, Engel EA, Klym S, Anderson JA, Glotzer SC. Rapid discovery of stable materials by coordinate-free coarse graining. Sci Adv. 2022;8(30):abn4117.
https://doi.org/10.1126/sciadv.abn4117.