

ORIGINAL RESEARCH

Open access

Multimodal, Physics-Informed Machine Learning for Accelerated Materials Design and Discovery

Lucas Meyer^{1*}, Stefan Braun¹, Anna Schmid², David Keller¹

Abstract

In the evolving landscape of computational materials engineering, the integration of multimodal data sources with physics-informed machine learning paradigms promises to revolutionize the pace and precision of materials design and discovery. This conceptual manuscript explores the synergies between diverse data modalities—ranging from experimental spectra to simulation-derived properties—and machine learning models constrained by physical laws, aiming to address persistent challenges in data scarcity, model generalizability, and discovery efficiency within materials science. By synthesizing recent advancements in representation learning, graph neural networks, and autonomous systems, we identify a conceptual gap in holistic frameworks that unify multimodal inputs with physics-based priors for accelerated inverse design. We introduce a novel conceptual framework, termed the Multimodal Physics-Constrained Discovery Engine (MPCDE), which structures data-model-discovery pipelines through layered interactions, feedback mechanisms, and epistemic steering logics. This framework emphasizes computational workflows that balance representation fidelity with inference robustness, incorporating uncertainty quantification to mitigate risks in high-throughput settings. Implications for the field include enhanced coupling of simulation and experimentation, improved scalability of foundation models, and streamlined closed-loop discovery systems. Ultimately, this work posits interpretive insights into how such integrated approaches can transform materials informatics into a more predictive and autonomous discipline, fostering innovations in energy, electronics, and structural materials.

Keywords Materials informatics, Uncertainty quantification, Graph neural networks, Physics-informed machine learning, Inverse materials design, Multimodal data integration

*Correspondence:

Lucas Meyer

lucas.meyer@gmail.com

¹ Department of Materials Modeling and Simulation, Faculty of Engineering, ETH Zurich, Zurich, Switzerland

² Department of Data-Driven Materials Science, Faculty of Engineering, University of Bern, Bern, Switzerland

Introduction

The computational shift in materials engineering

The field of materials engineering has undergone a profound transformation with the advent of computational and data-driven methodologies, shifting from traditional trial-and-error approaches to systematic, algorithm-guided exploration [1, 2]. This evolution is driven by the exponential growth in computational power and the

accumulation of vast materials datasets, enabling researchers to model complex phenomena at atomic and mesoscopic scales. High-throughput computation, for instance, has facilitated the screening of thousands of candidate materials for applications in catalysis, energy storage, and semiconductors, reducing the time from conception to deployment [3, 4]. Yet, despite these advances, challenges persist in bridging the gap between computational predictions and real-world performance, particularly in handling heterogeneous data sources and incorporating domain-specific knowledge.

Machine learning has emerged as a cornerstone in this shift, offering tools to extract patterns from data that elude conventional physics-based simulations [5, 6]. Early applications focused on property prediction using supervised learning, but recent developments emphasize generative models and active learning strategies to navigate vast chemical spaces [7, 8]. The integration of physics-informed constraints ensures that models respect fundamental laws, such as conservation principles and thermodynamic consistency, enhancing their extrapolative capabilities [9, 10]. This computational paradigm not only accelerates discovery but also democratizes access to advanced materials research, allowing interdisciplinary teams to leverage shared infrastructures.

Challenges in data-driven materials discovery

Data scarcity remains a critical bottleneck in materials informatics, where experimental data is often sparse, noisy, and expensive to acquire [11, 12]. Multimodal datasets, combining structural, spectroscopic, and mechanical information, offer a pathway to richer representations, but their integration poses representational and computational challenges [13, 14]. Traditional machine learning models struggle with modality fusion, leading to suboptimal performance in tasks like inverse design, where the goal is to identify compositions and structures that meet specified properties [15, 16].

Moreover, the lack of physics-based priors in purely data-driven approaches can result in unphysical predictions, undermining trust in automated systems [17, 18]. Uncertainty quantification emerges as essential for managing these risks, providing measures of confidence that inform decision-making in discovery pipelines [19, 20]. The rise of autonomous experimentation systems highlights the need for closed-loop frameworks that iteratively refine models based on new data, yet current implementations often lack comprehensive multimodal handling [21, 22].

Opportunities from emerging technologies

Advancements in deep learning architectures, such as graph neural networks, have proven adept at capturing materials' topological and relational features, enabling more accurate simulations of crystal structures and defect behaviors [23, 24]. Foundation models for science,

pretrained on large corpora, offer transferable knowledge that can be fine-tuned for specific materials tasks, potentially alleviating data limitations [25, 26]. Coupling these with high-throughput computation allows for scalable exploration, while simulation-experiment coupling facilitates validation and iteration [27, 28].

The potential for accelerated design lies in harnessing these technologies within unified ecosystems, where data flows seamlessly from acquisition to inference [29, 30]. This requires conceptual frameworks that address not only technical integration but also epistemic considerations, such as how multimodal inputs influence discovery steering [31, 32]. By focusing on infrastructure-level analysis, we can envision systems that optimize trade-offs between computational cost and discovery yield.

In this manuscript, we position a novel conceptual framework that integrates multimodal data with physics-informed machine learning to accelerate materials design and discovery. This framework, the Multimodal Physics-Constrained Discovery Engine (MPCDE), provides a structured approach to pipeline dynamics, emphasizing interpretive insights into representation-inference interactions and epistemic risk management.

Theoretical Background & Literature Synthesis

Foundations of materials informatics and machine learning integration

Materials informatics has established itself as a discipline that leverages data science to inform materials research, drawing on databases like the Materials Project and AFLOW to fuel predictive modeling [1, 3]. The core tenet involves transforming raw data into actionable insights through statistical and machine learning techniques, with an emphasis on feature engineering to represent materials' intrinsic properties [5, 7]. Representation learning has advanced significantly, moving from simple descriptors like atomic radii to sophisticated embeddings that capture multi-scale interactions [9, 11].

Machine learning in materials science encompasses a spectrum of methods, from regression for property prediction to clustering for phase diagram exploration [2, 4]. Deep learning architectures, particularly convolutional and recurrent networks, have been adapted for spectral data

analysis, while graph neural networks excel in handling crystalline and molecular graphs [6, 8]. These models facilitate high-throughput computation by automating screenings that would otherwise require exhaustive ab initio calculations [10, 12].

Multimodal data handling in computational frameworks

Multimodal datasets in materials engineering integrate diverse sources, such as X-ray diffraction patterns, electronic structure calculations, and mechanical testing results, to provide a comprehensive view of material behavior [13, 15]. Challenges in fusion include aligning disparate data formats and scales, often addressed through embedding techniques that project modalities into a common latent space [17, 19]. Literature highlights the value of such integration in enhancing model robustness, as seen in applications to defect characterization and alloy design [21, 23].

Foundation models for science represent a burgeoning area, where large-scale pretraining on multimodal corpora enables zero-shot or few-shot learning for materials tasks [25, 27]. These models incorporate cross-modal attention mechanisms to weigh contributions from different data types, improving generalization across materials classes [29, 31]. However, ensuring physics compliance remains crucial, as unconstrained models may violate symmetry or energy principles [14, 16].

Physics-informed constraints and uncertainty management

Physics-informed machine learning embeds domain knowledge directly into model architectures, such as through loss functions that penalize deviations from governing equations [18, 20]. This approach is particularly relevant for materials discovery, where simulations must align with quantum mechanics or continuum theories [22, 24]. Uncertainty quantification complements this by estimating epistemic and aleatoric uncertainties, guiding active learning strategies to prioritize informative data points [26, 28].

In inverse design contexts, physics priors steer optimization towards feasible solutions, avoiding exploration of invalid chemical spaces [30, 32]. Literature synthesizes these elements in discussions of simulation-experiment coupling,

where machine learning bridges predictive gaps by iteratively refining models with experimental feedback [1, 3].

Autonomous and closed-loop discovery systems

Autonomous discovery systems automate the design-make-test-analyze cycle, employing robotics and AI to conduct experiments with minimal human intervention [5, 7]. Closed-loop experimentation integrates real-time data acquisition with model updates, accelerating convergence on optimal materials [9, 11]. Key enablers include active learning algorithms that select experiments based on information gain, often coupled with graph-based representations for efficient querying [13, 15].

High-throughput frameworks extend this to parallel processing, screening vast libraries while incorporating multimodal feedback [17, 19]. Epistemic considerations in these systems involve balancing exploration and exploitation, with uncertainty metrics serving as steering logics [21, 23]. Synthesis of recent works reveals a trend towards hybrid systems that combine computational steering with experimental validation, fostering resilient discovery pipelines [25, 27].

Gaps in current paradigms and conceptual needs

Despite progress, existing literature reveals gaps in unified frameworks that holistically address multimodal integration with physics-informed learning for accelerated discovery [2, 4]. Many approaches focus on isolated components, such as representation or inference, without considering end-to-end pipeline dynamics [6, 8]. There is a need for conceptual models that interpret trade-offs in computational infrastructure, such as between data fidelity and model complexity [10, 12].

Furthermore, epistemic risk structures—encompassing biases in multimodal datasets and propagation of uncertainties—are underexplored in synthesis efforts [14, 16]. By integrating these elements, new frameworks can provide systems-level insights into how machine learning can transform materials design from reactive to proactive paradigms [18, 20]. This synthesis underscores the opportunity for original conceptualizations that emphasize

workflow interactions over empirical validations [22, 24, 26, 28, 30, 32].

Proposed conceptual framework

Overview of the Multimodal Physics-Constrained Discovery Engine (MPCDE)

The proposed conceptual framework, named the Multimodal Physics-Constrained Discovery Engine (MPCDE), offers an integrative structure for accelerating materials design and discovery through layered computational workflows. MPCDE conceptualizes the discovery process as a dynamic system comprising data ingestion, model refinement, and inference steering, unified under physics-informed constraints. At its core, the framework organizes multimodal inputs—such as structural graphs, spectral signatures, and thermodynamic profiles—into a hierarchical representation layer that feeds into physics-constrained learning modules.

The structural layers of MPCDE include: (1) a multimodal fusion substrate that harmonizes diverse data streams via embedding alignments; (2) a physics-injection core that embeds conservation laws and symmetry operators into neural architectures; and (3) a discovery orchestration layer that manages feedback loops for iterative refinement. Data-model-discovery pipelines within MPCDE flow from raw inputs through fused representations to predictive outputs, with inverse mappings enabling design optimization.

The structural decomposition of MPCDE layers and their functional roles is summarized in **Table 1**.

Table 1. Framework Architecture Decomposition

| MPCDE Layer | Core Function | Key Computational Elements | Disc Contribution |
|---------------------------|---|--|--------------------------|
| Multimodal Data Substrate | Aggregates heterogeneous materials datasets | Spectra, graphs, thermodynamics, imaging | Expensive representation |
| Fusion & Alignment Layer | Harmonizes modality embeddings | Latent projections, weighting schemes | Improve model coherence |
| Physics-Injection | Embeds physical laws | Symmetry operators, | Ensures physical |

| | | | |
|-------------------------------|----------------------------------|--------------------------------------|-------------------------------------|
| Core | into models | conservation constraints | plaus |
| Discovery Orchestration Layer | Directs design and optimization | Inverse models, candidate evaluators | Accelerated material identification |
| Epistemic Steering Layer | Governs uncertainty and feedback | UQ metrics, error monitors | Stabilizes inference paths |

Feedback loops are integral, allowing discrepancies between predicted and physical behaviors to trigger adaptive updates, such as recalibrating embeddings or resampling uncertainties. Computational steering logics guide this process by prioritizing pathways that maximize information density while minimizing epistemic risks, such as overconfidence in sparse regimes.

Pipeline dynamics and representation-inference interactions

In MPCDE, data pipelines emphasize modular transformations, where multimodal sources are processed through graph-based encoders to capture relational dependencies [6, 23]. The interaction between representations and inference is captured through a conceptual formula that expresses the trade-off in fidelity versus generalizability. This can be conceptualized as the

$$\eta = \alpha \cdot C + \beta \cdot U \sum m w_m \cdot I_m$$

fusion efficiency metric, symbolized as

where w_m weights modality contributions, I_m denotes information content, C represents computational cost, U is aggregated uncertainty, and α, β are balancing factors. This formula captures the interaction between multimodal inputs and system constraints, highlighting how weighted integrations can optimize pipeline throughput without empirical tuning.

Inference steering in MPCDE leverages physics-informed priors to direct exploration, ensuring that discovery paths align with feasible physical manifolds. A second formula conceptualizes this steering as a constrained optimization dynamic: $\arg \max_{d \geq \theta} dV(d) \text{ subject to } P(d)$, where $V(d)$ is the value function for a design candidate d , $P(d)$ evaluates physical plausibility, and θ is a threshold for viability. This may be expressed as a mechanism to

balance exploratory breadth with constraint adherence, fostering efficient navigation of materials spaces.

Epistemic risk structures and feedback mechanisms

Epistemic risks in MPCDE are managed through layered uncertainty propagation, where aleatoric noise from data and epistemic gaps from models are quantified and fed back into the engine. This creates resilient loops that adapt to new insights, such as refining graph embeddings based on discrepancy signals [19, 26].

A third formula addresses feedback intensity:

$$F = \frac{\gamma \cdot \Delta E}{(1 - e^{-\kappa \cdot S})}$$

where ΔE is the error delta between predicted and constrained outcomes, S is the epistemic surprise measure, and γ, κ modulate response sensitivity.

This captures the interaction between error signals and surprise in driving loop iterations, providing interpretive insight into how MPCDE maintains stability in uncertain environments.

The MPCDE framework is conceptualized in **Figure 1**, which depicts a schematic of the layered architecture. At the base, multimodal data streams converge into a fusion hub, represented as interconnected nodes symbolizing modality alignments. Ascending layers show physics-injection via embedded operators, illustrated as constraining envelopes around neural pathways. The top orchestration layer features cyclic arrows indicating feedback loops, with steering logics as directional vectors guiding flows. Pipelines are shown as vertical channels linking layers, with uncertainty clouds overlaying inference points to highlight risk management. This textual description underscores the framework's emphasis on integrated dynamics, offering systems-level insights into accelerating materials discovery through balanced computational infrastructures. The layered architecture, feedback loops, and physics-steering dynamics of the proposed engine are schematically illustrated in **Figure 1**.




Figure 1. Conceptual architecture of the Multimodal Physics-Constrained Discovery Engine (MPCDE).

Multimodal data streams converge within a fusion substrate where cross-modal embeddings align heterogeneous representations. Physics-informed operators inject conservation laws and feasibility constraints into learning pathways, guiding inference within physically valid manifolds. The discovery orchestration layer enables inverse design and candidate optimization, while epistemic steering infrastructures regulate uncertainty, error propagation, and adaptive feedback. Bidirectional loops facilitate continuous recalibration across the pipeline, supporting resilient and accelerated materials discovery.

Analytical Implications

Interpretive dynamics in discovery pipelines

The MPCDE framework provides systems-level insights into how multimodal, physics-informed machine learning can reshape discovery pipelines in materials engineering. By structuring data flows through fused representations and constrained inferences, MPCDE highlights the potential for enhanced efficiency in navigating complex chemical spaces [1, 13]. Analytically, this implies a shift from linear prediction models to iterative ecosystems where representation-inference interactions drive adaptive exploration. For instance, in high-throughput settings, the framework's emphasis on modality weighting allows for prioritized processing of informative data streams, reducing computational overhead while maintaining fidelity to physical constraints [5, 17].

Epistemic risk structures within MPCDE offer interpretive lenses on uncertainty propagation, suggesting that integrated quantification can steer decisions towards robust outcomes [19, 26]. This has implications for inverse design tasks, where unphysical proposals are preemptively filtered, streamlining the path from conceptual ideation to

viable candidates [14, 30]. Computational workflow dynamics further imply trade-offs in scalability; as multimodal inputs grow, the engine's feedback loops can mitigate overfitting by recalibrating based on physics priors, fostering generalizable models across materials classes [9, 22].

A conceptual formula formalizes this scalability trade-off:

$$S = \delta \cdot (D + U) \lambda$$

where M is multimodal richness, P physics, M constraint density, D data volume, U uncertainty variance, and λ, δ are modulation coefficients. This may be expressed as capturing the balance between input complexity and system stability, providing insight into how MPCDE optimizes resource allocation in discovery infrastructures.

Infrastructure trade-offs and steering logics

Analytically, MPCDE underscores infrastructure trade-offs in coupling simulation and experimentation, where physics-informed elements bridge disparate scales [4, 18]. The framework's orchestration layer implies that steering logics can dynamically adjust exploration strategies, favoring exploitation in well-characterized domains while promoting diversity in sparse ones [8, 24]. This interpretive approach reveals how closed-loop mechanisms enhance autonomy, as feedback from epistemic surprises refines model behaviors without external intervention [11, 27].

In terms of representation learning, implications extend to graph neural architectures, where multimodal embeddings interact with physics operators to yield richer latent spaces [6, 23]. This suggests improved handling of disordered materials, where traditional methods falter, by incorporating uncertainty-aware fusions that preserve relational integrity [20, 31]. Systems-level insights point to epistemic benefits, such as reduced bias in dataset curation, enabling more equitable coverage of materials ecosystems [12, 28].

Systems-level analytical implications and associated infrastructure trade-offs are synthesized in **Table 2**.

Table 2. Analytical Implications & Infrastructure Trade-Offs

| Analytical Dimension | MPCDE Mechanism | Infrastructure Impact | Discovery Implications |
|-------------------------|------------------------------|------------------------------|------------------------------|
| Representation Fidelity | Multimodal fusion weighting | Increased compute demand | Higher predictive robustness |
| Physics Compliance | Constraint-embedded learning | Simulation coupling costs | Reduced unphysical outputs |
| Uncertainty Governance | Epistemic steering loops | Monitoring overhead | Risk-aware design filters |
| Scalability | Feedback recalibration | HPC dependency | Cross-domain generalization |
| Exploration Efficiency | Steering logics | Adaptive resource allocation | Accelerated inverse design |

| | | | |
|-------------------------|------------------------------|------------------------------|------------------------------|
| Representation Fidelity | Multimodal fusion weighting | Increased compute demand | Higher predictive robustness |
| Physics Compliance | Constraint-embedded learning | Simulation coupling costs | Reduced unphysical outputs |
| Uncertainty Governance | Epistemic steering loops | Monitoring overhead | Risk-aware design filters |
| Scalability | Feedback recalibration | HPC dependency | Cross-domain generalization |
| Exploration Efficiency | Steering logics | Adaptive resource allocation | Accelerated inverse design |

Another formula conceptualizes steering efficiency:

$$E = \eta \int (I - R) dt$$

with I as information gain, R as risk exposure, and η an efficiency scalar. This captures the interaction between temporal feedback and risk minimization, illustrating how MPCDE's logics accumulate value over discovery cycles.

Field-wide implications for materials innovation

The analytical implications of MPCDE extend to broader field transformations, particularly in accelerating design for sustainable applications like energy materials [2, 15]. By interpreting pipeline dynamics through physics-constrained lenses, the framework implies faster iteration in autonomous systems, potentially compressing discovery timelines [7, 21]. Uncertainty quantification's role in risk structures further implies resilient infrastructures, capable of adapting to evolving data landscapes without sacrificing accuracy [10, 25].

In computational terms, this fosters hybrid ecosystems where foundation models interface with domain-specific priors, implying scalable advancements in materials informatics [3, 16]. A final formula addresses innovation

$$V = \kappa \frac{F+L}{T}$$

where F is fusion depth, L loop iterations, T computational time, and κ a velocity constant.

This can be conceptualized as representing the acceleration derived from integrated workflows, offering interpretive guidance on optimizing materials discovery paradigms [29, 32].

Results and Discussion

Limitations of the conceptual approach

While the Multimodal Physics-Constrained Discovery Engine (MPCDE) offers a structurally integrative interpretive architecture for mapping discovery processes across multimodal and physics-informed materials ecosystems, it inevitably inherits limitations embedded within the paradigms it synthesizes. These constraints emerge not from weaknesses in the framework itself, but from epistemic, infrastructural, and translational tensions that characterize contemporary data-driven materials science.

At the representational level, MPCDE assumes that layered modality fusion can be operationalized through progressively harmonized alignment mechanisms.

Conceptually, this layered integration logic presumes that spectroscopic signals, microstructural imaging, thermomechanical measurements, and quantum-simulation outputs can be coherently integrated within unified representation spaces. However, real-world datasets often exhibit deep incompatibilities. Scale disparities, resolution asymmetries, modality-specific noise structures, and measurement artifacts introduce ontological discontinuities that resist seamless fusion. Consequently, the framework's fusion layers risk analytically smoothing over modality frictions that may carry scientifically meaningful distinctions [5, 19].

Epistemic constraints emerge more prominently within MPCDE's physics-steering infrastructures. The framework positions physics priors as stabilizing anchors that regulate inference trajectories and constrain discovery exploration. Yet this stabilizing role is contingent upon the fidelity of the priors themselves. In regimes characterized by sparse experimental grounding, extrapolative simulations, or approximated boundary conditions, physics constraints may function less as epistemic stabilizers and more as directional biases. Under such conditions, steering logics designed to regulate uncertainty could inadvertently amplify it—channeling discovery pathways toward locally coherent but globally suboptimal solution spaces [13, 26].

Infrastructure dependencies introduce additional limitations. MPCDE's recursive feedback loops, uncertainty monitoring layers, and multimodal ingestion pipelines presuppose sustained access to high-performance computational environments. Continuous model updating, simulation

coupling, and multimodal synchronization demand substantial storage, compute, and orchestration capacities. These requirements may restrict the framework's operational translation in resource-constrained laboratories or emerging research ecosystems, thereby creating asymmetries in accessibility and implementation scalability [8, 22].

Finally, the framework's interpretive orientation constitutes both a strength and a limitation. MPCDE is deliberately conceptual, designed to structure discovery logics rather than empirically validate predictive performance. While this abstraction enables cross-platform theoretical generality, it constrains immediate experimental translation. Practical deployment would require methodological instantiation, validation pipelines, and benchmarking infrastructures beyond the present conceptual scope [1, 14].

Synergies with existing ecosystems

Notwithstanding these limitations, MPCDE demonstrates strong structural compatibility with contemporary materials research infrastructures and computational discovery ecosystems.

One domain of synergy lies in high-throughput experimentation platforms. Autonomous synthesis, rapid characterization, and simulation-driven screening pipelines generate continuous multimodal data streams. Integrating MPCDE's steering and interpretive layers into such infrastructures could enhance discovery governance—enabling not only accelerated exploration but epistemically guided experimentation trajectories [4, 11].

The framework also aligns closely with advances in graph neural networks and topology-aware representation systems. While such architectures excel at encoding structural and relational material features, they often operate with limited explicit physics integration. MPCDE's physics-constrained steering layers can function as interpretive overlays that guide representational evolution toward physically plausible regions of materials space, thereby enhancing both predictive reliability and mechanistic interpretability [6, 23].

Closed-loop discovery ecosystems present another fertile convergence zone. In these systems, simulation outputs inform experimental design, and experimental feedback recursively refines models. MPCDE enriches this paradigm by embedding epistemic monitoring and uncertainty-

responsive steering within the loop itself. Feedback thus becomes not merely corrective but discovery-shaping—actively modulating search trajectories and inference priorities [18, 27].

Further synergies emerge in uncertainty quantification and foundation-scale materials modeling. By positioning uncertainty as an infrastructural signal rather than a peripheral metric, MPCDE supports hybrid modeling architectures that integrate multimodal datasets with large pretrained materials representations. Such integration is particularly consequential for inverse design, where physics-conditioned feasibility constraints can guide generative exploration toward synthesizable and functional material candidates [10, 25]. In this capacity, MPCDE bridges data-driven inference with physically grounded design governance [15, 30].

Future directions and extensions

The conceptual elasticity of MPCDE opens multiple avenues for theoretical expansion and infrastructural evolution.

An immediate extension involves adaptive layering architectures capable of accommodating emergent data modalities. As robotic laboratories and cyber-physical experimentation systems mature, real-time sensor streams, in-situ diagnostics, and dynamic process signals will become integral to discovery ecosystems. Extending MPCDE to integrate temporally streaming modalities would transform it from a static fusion engine into a temporally responsive discovery infrastructure [7, 21].

Another promising direction concerns multi-fidelity epistemic hierarchies. Materials discovery increasingly operates across resolution gradients—from coarse surrogate predictors to high-precision quantum simulations. Embedding fidelity-aware layering within MPCDE would allow interpretive mapping of trade-offs between computational efficiency, predictive resolution, and epistemic reliability, enabling dynamically optimized discovery allocation strategies [12, 28].

Bias mitigation and representational inclusivity constitute further expansion pathways. Existing materials datasets disproportionately represent stable or industrially prioritized compounds, leaving vast chemical territories underexplored. Future MPCDE architectures could incorporate bias-diagnostic steering layers that redirect

discovery flows toward epistemically marginalized materials domains, fostering more inclusive and exploratory innovation ecosystems [3, 16].

Beyond discrete pipelines, integration with digital materials twins offers transformative infrastructural potential. Embedding MPCDE within twin architectures could enable predictive co-evolution between computational models and physical materials systems, supporting real-time performance forecasting and lifecycle optimization. However, such integration necessitates advances in synchronization protocols, scalable computation, and real-time data assimilation infrastructures.

Ultimately, these forward trajectories position MPCDE not as a static conceptual artifact but as an evolving interpretive engine—capable of scaling alongside autonomous laboratories, multimodal AI systems, and physics-grounded discovery infrastructures. In doing so, it contributes to the broader transformation of materials engineering into an integrated, intelligent, and reflexive innovation ecosystem.

Conclusion

In synthesizing the conceptual landscape of multimodal, physics-informed machine learning, this manuscript has introduced the Multimodal Physics-Constrained Discovery Engine (MPCDE) as a novel framework for accelerating materials design and discovery. Through structured layers, pipeline dynamics, and epistemic steering, MPCDE offers interpretive insights into balancing representation fidelity with inference robustness, addressing key challenges in data scarcity and model reliability.

Analytical implications highlight enhanced workflow efficiencies and infrastructure trade-offs, with conceptual formulas illustrating system interactions that foster resilient discovery paradigms. Discussions of limitations underscore the need for adaptive extensions, while synergies with existing ecosystems point to transformative potential in autonomous systems and inverse design.

Ultimately, MPCDE positions materials informatics at the cusp of a more integrated, computationally steered future, promising advancements in sustainable and high-performance materials across sectors. This conceptual contribution invites further exploration into unified frameworks that harness multimodal synergies with

physical principles, driving the next wave of innovation in computational materials engineering.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 30 Sep 2022 Revised: 03 Nov 2022 Accepted: 29 Dec 2022
Published online: 18 March 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

- 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(1):84. <https://doi.org/10.1038/s41524-022-00765-z>.
- Choudhary K, DeCost B, Chen C, Jain A, Tavazza F, Cohn R, 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>.
- Chang R, Wang YX, Ertekin E. Towards overcoming data scarcity in materials science: Unifying models and datasets with a mixture of experts framework. *npj Comput Mater.* 2022;8(1):242.
- Thiagarajan JJ, Venkatesh B, Anirudh R, Bremer P-T, Gaffney J, Anderson G, et al. Designing accurate emulators for scientific processes using calibration-driven deep models. *Nat Commun.* 2020;11(1):5622.
- de Pablo JJ, Jackson NE, Webb MA, Chen LQ, Moore JE, Morgan D, et al. New frontiers for the materials genome initiative. *npj Comput Mater.* 2019;5(1):41. <https://doi.org/10.1038/s41524-019-0173-4>.
- Gupta V, Choudhary K, DeCost B, Tavazza F, Campbell C, Liao WK, et al. Structure-aware graph neural network based deep transfer learning framework for enhanced predictive analytics on diverse materials datasets. *npj Comput Mater.* 2024;10(1):1.
- Oviedo F, Ferres JL, Agarwal TM, Cai J, Zhang T, Ganguli S, et al. Fast and interpretable classification of small X-ray diffraction datasets using data augmentation and deep neural networks. *npj Comput Mater.* 2019;5(1):60. <https://doi.org/10.1038/s41524-019-0196-x>.
- Lu S, Zhou Q, Guo Y, Zhang Y, Wu Y, Wang J. Coupling a crystal graph multilayer descriptor to active learning for rapid discovery of 2D ferromagnetic semiconductors/heterostructures. *Nat Commun.* 2020;11(1):3714.
- Chen C, Zuo Y, Ye W, Li J, 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>.
- Saidi P, Kaminskyj R, Boyd PG, Freitas Dos Santos D, Woo TK. A machine learning approach for rapid and robust data mining of gases adsorbed in metal-organic frameworks. *Digit Discov.* 2022;1(5):655-68.
- Connolly AB, Sutherland DR, Amsler M, Akhlaghi S, Gans MA, Chang MC, et al. Autonomous materials synthesis via hierarchical active learning of nonequilibrium phase diagrams.

Sci Adv. 2021;7(51):eabg4930.

<https://doi.org/10.1126/sciadv.abg4930>.

Li Y, Wen Y, Wang D, Zhang L. Composite materials property prediction with uncertainty quantification using a multi-task deep learning framework. *Comput Mater Sci.* 2023;216:111847.

Batra R, Song L, Ramprasad R. Emerging materials intelligence ecosystems propelled by machine learning. *Nat Rev Mater.* 2021;6(8):655-78.

<https://doi.org/10.1038/s41578-020-00255-y>.

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 design of inorganic materials. *npj Comput Mater.* 2020;6(1):84.

<https://doi.org/10.1038/s41524-020-00352-0>.

Wang AY-T, Murdock RJ, Kauwe SK, Bukhvalov A, Hu A, Salim S, et al. Machine learning for materials scientists: An introductory guide toward harnessed simulations. *Chem Mater.* 2020;32(12):4954-65.

<https://doi.org/10.1021/acs.chemmater.0c01907>.

Huang H, Mojumder S, Suarez D, Al Amin A, Fleming M, Liu WK. Knowledge database creation for design of polymer matrix composite. *Comput Mater Sci.* 2022;214:111703.

Rohilla T, Singh N, Krishnan NC, Mahajan DK. Designing sulfonated polyimide-based fuel cell polymer electrolyte membranes using machine learning approaches. *Comput Mater Sci.* 2023;219:111974.

Wang XQ, Chen P, Chow CL, Lau D. Artificial-intelligence-led revolution of construction materials: From molecules to Industry 4.0. *Matter.* 2023;6(6):1831-59.

Galan EA, Zhao H, Wang X, Dai Q, Huck WTS, Ma S. Intelligent microfluidics: The convergence of machine learning and microfluidics in materials science and biomedicine. *Matter.* 2020;3(6):1893-922.

Bengio E, Jain M, Korablyov M, Precup D, Bengio Y. GFlowNets for AI-driven scientific discovery. *Digit Discov.* 2023;2(4):557-77.

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.

<https://doi.org/10.1038/s43246-022-00315-6>.

Butler KT, Oviedo F, Balachandran PV, McMahon EJ. Machine learning for materials science: Barriers to broader adoption. *Matter.* 2023;6(6):1856-65.

Kalinin SV, Ziatdinov MA, Hinkle J, Ghosh A, Kelley KP, Lupini AR, et al. Teaching machine learning to materials scientists: Lessons from hosting tutorials and competitions. *Matter.* 2022;5(5):1312-5.

Oviedo F, Ren Z, Sun X, Zhang C, Liu Z, Paggiaro A, et al. Leveraging machine learning in the innovation of functional materials. *Matter.* 2023;6(8):2486-8.

Maruyama B, Hatrick-Simpers JR, Brown KA, Anasori B, Abolhasani M, Kalinin S, et al. Autonomous experimentation systems for materials development: A community perspective. *Matter.* 2021;4(9):2702-26.

<https://doi.org/10.1016/j.matt.2021.06.036>.

Kalinin SV, Mukherjee D, Roccapriore K, Blaiszik BJ, Ghosh A, Ziatdinov MA, et al. Machine learning for automated experimentation in scanning transmission electron microscopy. *npj Comput Mater.* 2023;9(1):227.

Tao Q, Xu P, Li M, Lu W. Machine learning for perovskite materials design and discovery. *npj Comput Mater.* 2021;7(1):23.

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.

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

Jiang L, Zhang Z, Hu H, He X, Fu H, Xie J. A rapid and effective method for alloy materials design via sample data transfer machine learning. *npj Comput Mater.* 2023;9(1):79.

<https://doi.org/10.1038/s41524-023-00979-9>.

Frey NC, Wang J, Bellon G, Akinwande D. Prediction of the electron density of states for crystalline compounds with deep convolutional neural networks. *npj Comput Mater.* 2022;8(1):235.

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

Pilania G, Gubernatis JE, Lookman T. Multi-fidelity machine learning models for accurate bandgap predictions of solids. *Comput Mater Sci.* 2017;129:156-63.

<https://doi.org/10.1016/j.commatsci.2016.12.004>.