ORIGINAL RESEARCH **Open access**

Physics as Constraint, Not Input: A Conceptual Reframing of Physics-Guided Machine Learning in Materials Science

Kenji Mori[1], Yuki Sato[1*], Rina Okabe[2], Takashi Ito[1]

Abstract

Physics-guided machine learning (PGML) has emerged as a hybrid paradigm in materials science, integrating domain knowledge with data-driven methods to enhance predictive accuracy and generalizability. Conventional approaches typically embed physical principles as soft inputs—either through loss-function regularization or auxiliary features—allowing violations during optimization. This manuscript advances a conceptual reframing in which physics operates as a hard constraint on the model's hypothesis space rather than as an additive input. By restricting permissible functional forms, symmetries, and conservation relations a priori, the framework enforces physical consistency at the architectural level, altering the interaction dynamics between data and prior knowledge. The reframing yields systems-level insights into epistemic trade-offs: reduced reliance on large datasets, improved extrapolation beyond training regimes, and inherent satisfaction of thermodynamic or mechanical invariants critical to materials behavior. Analytical implications include feedback structures that couple data refinement to constraint satisfaction, revealing emergent robustness in multiscale modeling. This perspective addresses persistent challenges in materials science, such as sparse experimental data and complex microstructure-property relationships, without resorting to empirical validation. The contribution lies in reinterpreting PGML's epistemic foundation, steering future developments toward constraint-centric designs that prioritize physical fidelity over post-hoc penalization.

Keywords Epistemic constraints, Theoretical framework, Materials science, Physics-guided machine learning, Constraint-based modeling, Hypothesis space restriction

*Correspondence:
Yuki Sato
yuki.sato@gmail.com
[1] Department of Intelligent Materials Engineering, Faculty of Engineering, University of Tokyo, Tokyo, Japan
[2] Department of AI-Driven Materials Discovery, Faculty of Information Science, Kyoto University, Kyoto, Japan

Introduction

The integration of machine learning (ML) into materials science has transformed approaches to property prediction, microstructure analysis, and materials discovery. Traditional physics-based modeling, rooted in continuum mechanics, quantum mechanics, or statistical thermodynamics, provides a mechanistic understanding but often demands intensive computation or simplifying assumptions that limit applicability to complex systems. Purely data-driven ML, conversely, excels at pattern recognition in high-dimensional spaces yet frequently yields physically implausible predictions, especially under data scarcity or extrapolation conditions common in materials research [1, 2].

Hybrid paradigms, collectively termed physics-guided machine learning (PGML) or physics-informed neural networks (PINNs), seek to bridge this divide. These methods incorporate physical laws—typically expressed as partial differential equations (PDEs), conservation principles, or symmetry relations—directly into the learning process. In materials science,

PGML has been applied to tasks ranging from constitutive modeling of viscoelastic materials to inverse design of metamaterials and fatigue life prediction in additively manufactured alloys [3, 4].

The dominant paradigm treats physics as an input: either as features augmenting the input space or, more commonly, as soft constraints enforced via additional terms in the loss function. For instance, residual losses derived from governing equations penalize deviations during training, allowing the network to approximate solutions while remaining flexible [5]. This input-centric approach has demonstrated value in scenarios with limited data, as the physical regularization guides the model toward plausible regions of parameter space [6].

However, treating physics as input introduces subtle epistemic tensions. Soft enforcement permits transient or final violations of physical laws, particularly when data gradients dominate or when the optimization landscape favors data fidelity over constraint satisfaction. In materials science, where phenomena span multiple length and time scales—from atomic lattices to macroscopic deformation—such violations can undermine interpretability and reliability. Microstructural heterogeneities, phase transformations, and nonlinear responses amplify these issues, as models may learn spurious correlations that contradict conservation laws or thermodynamic principles [7].

Recent literature highlights these limitations. Applications to polymer mechanics, cellular materials, and defect identification reveal that while PGML improves accuracy over pure data-driven baselines, generalization to unseen loading conditions or compositions remains challenging [8, 9]. Inverse problems, such as parameter identification from full-field measurements, further highlight sensitivities to the relative weighting of data and physics terms, often requiring ad hoc tuning [10].

The present work proposes a conceptual reframing: physics as constraint, not input. Rather than augmenting the loss or input layer, physical principles define the allowable structure of the model itself—restricting the hypothesis space to functions that inherently satisfy key invariants. This shift alters the fundamental dynamics of interaction between data and knowledge. Data no longer competes with physics in a weighted sum but refines parameters within a physically admissible manifold.

Such a reframing draws inspiration from broader theoretical developments in constrained optimization and symmetry-aware architectures but applies them specifically to the epistemic structure of PGML in materials contexts. By enforcing constraints at the architectural level—through equivariant layers, projection operators, or embedded conservation modules—the framework ensures physical fidelity by construction, reducing the epistemic burden on data and enhancing systems-level coherence [11].

This perspective addresses core challenges in materials science without relying on empirical demonstrations. In multiscale modeling, where homogenization theories must preserve energy dissipation or momentum balance, hard constraints prevent unphysical energy creation or violation of thermodynamic inequalities. In design tasks, the constrained hypothesis space naturally embeds manufacturability or stability criteria, steering optimization toward feasible outcomes.

The reframing also illuminates trade-offs in knowledge integration. Traditional input-based PGML balances flexibility and fidelity through hyperparameters; constraint-based approaches trade flexibility for guaranteed consistency, potentially at the cost of expressivity in highly uncertain regimes. Yet this trade-off yields analytical insights into robustness: models become less sensitive to distributional shifts in data, as the constraint manifold provides an intrinsic regularizer.

Epistemically, the approach reframes the role of theory in ML. Physics ceases to be an external advisor and becomes the boundary condition of the learning system itself. This aligns with deeper philosophical considerations in scientific modeling, where laws define the space of possible worlds rather than merely penalizing deviations [12].

Subsequent sections synthesize the theoretical background, highlighting how current literature implicitly or explicitly navigates these input-constraint distinctions. The proposed framework is then articulated through its core components, interaction dynamics, and systems insights, with a textual description of an accompanying conceptual figure. The analysis remains interpretive, focusing on conceptual implications rather than predictive claims.

Theoretical Background and Literature Synthesis

Evolution of physics-informed approaches in machine learning

The integration of physical knowledge into machine learning has evolved through distinct methodological phases, reflecting a gradual shift from loosely coupled hybrid models toward deeply embedded physics-aware learning frameworks. Early physics-guided approaches relied primarily on domain-specific feature construction, using outputs from physical simulations or analytically derived descriptors to augment conventional data-driven models [13]. In these formulations, physics served as an external guide to representation design rather than as an intrinsic component of the learning process.

By the early 2020s, this paradigm underwent a substantive transformation with the emergence of physics-informed neural networks (PINNs), which introduced a tighter coupling between physical laws and optimization objectives. Instead of treating physics as preprocessed input, PINNs embed governing equations—typically expressed as partial differential equations (PDEs)—directly into the loss function, thereby enabling simultaneous satisfaction of data-fidelity and physical-consistency constraints [5]. This approach allowed neural networks to address both forward and inverse problems within a unified framework, thereby reframing learning as a constrained optimization problem grounded in first principles.

In materials science, the adoption of physics-informed learning accelerated markedly after 2020, driven by structural limitations endemic to the field. Experimental data are often sparse, costly, and heterogeneous, while strongly coupled, multiscale physical processes govern material behavior. Under such conditions, purely data-driven models tend to suffer from poor extrapolation and limited interpretability. Physics-informed approaches emerged as a response to these challenges, offering a means to stabilize learning, reduce data dependence, and ensure physically admissible predictions across regimes that are underrepresented in available datasets.

Physics as input: Feature engineering and loss regularization

Within the broader category of physics-guided machine learning (PGML), two dominant strategies can be distinguished based on how physical knowledge is incorporated: physics as input and physics as regularization. In the former, physical insight informs feature engineering, with models trained on descriptors that encode invariances, symmetries, or scaling relations intrinsic to the underlying system. Examples include symmetry-preserving descriptors derived from crystallographic principles or dimensionless groups constructed using the Buckingham Π theorem, which compress physical relationships into reduced representations suitable for learning [14].

Despite the conceptual appeal of feature-level integration, loss-based regularization has become the prevailing strategy in contemporary PGML. In this formulation, the learning objective combines a data-mismatch term with an explicit physics residual, often expressed as a weighted norm of the violations of the governing equations. The weighting coefficient λ mediates the trade-off between empirical accuracy and adherence to physical laws, allowing practitioners to tune the relative influence of data and theory during optimization [6]. This flexibility is particularly valuable in regimes where data are limited or noisy, and where strict enforcement of physics may otherwise impede convergence.

In materials science applications, physics-regularized learning has demonstrated tangible benefits across a range of problem settings. In viscoelastic modeling, constitutive stress–strain relations derived from continuum mechanics are incorporated into the loss function to supplement sparse experimental measurements, improving both predictive stability and physical plausibility [15]. Similarly, fatigue prediction in additively manufactured materials leverages physics-guided

loss terms to enforce crack initiation and propagation laws, mitigating the risk of unphysical extrapolation over long fatigue lifetimes [3]. Collectively, these studies illustrate how physics-guided regularization constrains optimization trajectories toward physically meaningful minima, enhancing generalization in low-data regimes.

Physics-informed neural networks and variants in materials science

Physics-informed neural networks represent the most explicit instantiation of physics-aware learning, embedding governing equations directly into the training process through automatic differentiation. In materials science, PINNs and related architectures—often described as physics-guided neural networks (PGNNs)—have been applied to a wide array of problems, including defect identification, microstructure reconstruction, and property characterization [3, 8]. Their appeal lies in their ability to recover latent fields or material parameters from indirect or incomplete observations by enforcing physical consistency throughout the solution domain.

A prominent application of PINNs involves inverse problems, where internal material states must be inferred from sparse boundary or sensor data. In such cases, embedding PDEs governing heat transfer, elastic wave propagation, or diffusion enables the model to resolve ambiguities that would otherwise remain underdetermined in purely data-driven settings [16]. Beyond the canonical PINN formulation, several methodological variants have been proposed to address practical challenges, such as slow convergence and imbalance in the loss function. These include co-training multiple networks for coupled variables. These adaptive sampling strategies concentrate training points in regions of high residual error, and hybrid architectures that separate state estimation from parameter identification [17].

In polymeric and composite materials, PINNs have been employed to model curing kinetics, dynamic wave control in metabeams, and multifunctional responses under coupled mechanical and thermal loading [2, 18]. While these studies demonstrate the expressive power of physics-informed architectures, they also expose persistent limitations, including sensitivity to constraint weighting, optimization stiffness, and scalability to high-dimensional domains. These challenges underscore the need for principled frameworks to manage trade-offs between data fit, physical fidelity, and computational tractability.

Applications to multiscale modeling and inverse problems

Physics-guided machine learning has proven particularly valuable in materials science due to its capacity to operate across multiple length and time scales. At the continuum scale, PGML frameworks have been used to predict large-deformation viscoelastic behavior and to enable inverse design of cellular metamaterials, where mechanical objectives are mapped back to admissible geometries under governing mechanical constraints [9, 19]. By embedding continuum mechanics into learning objectives, these models preserve physical coherence while navigating high-dimensional design spaces.

At the microscale, physics-informed approaches facilitate the reconstruction of otherwise ill-posed problems, such as inferring internal structures from diffraction patterns or full-field displacement measurements. Incorporating physical constraints helps disambiguate multiple candidate solutions and improves robustness to measurement noise [20]. Inverse design paradigms increasingly employ paired generator–forward architectures, where the forward model acts as a physics-constrained surrogate simulator, enforcing feasibility during the design process [9]. While these approaches have shown promise for complex systems such as three-dimensional cellular materials, they also highlight the sensitivity of inverse solutions to constraint weighting and model calibration, reinforcing the importance of systematic trade-off management in physics-informed learning.

Challenges and limitations in current paradigms

Despite notable empirical successes, prevailing physics-guided and physics-informed machine learning paradigms exhibit fundamental conceptual and practical limitations. A central weakness lies in their reliance on soft constraint enforcement, in which physical laws are incorporated as penalty terms rather than as inviolable structural conditions. In data-scarce

regimes, common in materials science, optimization dynamics often prioritize empirical loss terms over physics-based regularization, leading to non-negligible violations of governing principles [21]. Such violations are not merely numerical artifacts; in materials systems, breaches of thermodynamic consistency—such as negative stiffness, energy non-conservation, or violation of dissipation inequalities—directly undermine scientific credibility and engineering reliability.

Generalization remains a persistent challenge. Models trained on narrowly sampled microstructures, compositions, or processing regimes frequently fail under compositional shifts or novel morphologies, even when physics-informed features or residuals are included [7]. This limitation reflects a deeper issue: physics incorporated as input or loss regularization does not structurally bind the hypothesis space. Consequently, the learned function may satisfy physical constraints locally or approximately, yet deviate globally when extrapolating beyond the training distribution.

Computational barriers further constrain applicability. Evaluating high-dimensional PDE residuals—particularly for nonlinear, multiscale, or coupled systems—introduces stiffness into the optimization landscape, substantially increasing training costs [22]. These challenges scale unfavorably with dimensionality, rendering many physics-informed approaches impractical for realistic three-dimensional materials systems or time-dependent processes.

Moreover, insights from the explainable machine learning literature in materials science highlight that the inclusion of physics terms does not inherently resolve the black-box problem. Even when models obey physical laws in aggregate, their internal representations often remain opaque, offering limited mechanistic insight into structure–property–processing relationships [23]. Physics-guided losses constrain outputs, but they do not guarantee interpretability at the architectural or representational level.

Conceptual gaps and emerging perspectives

Recent integrative reviews and syntheses converge on a critical observation: in most existing frameworks, physics is treated as an auxiliary consideration rather than as a foundational organizing principle of learning [7]. Comprehensive surveys of PINNs in materials modeling similarly note that while loss augmentation improves data efficiency, it falls short of delivering deep integration between physical theory and model structure [24]. Physics remains external to the representational core, appended rather than internalized.

Several emerging works gesture toward more structural integration, exploring symmetry-preserving layers, equivariant architectures, or projection-based correction schemes that enforce invariances or constraints more directly [25]. However, these efforts largely remain within the input-centric or penalty-based paradigm, offering incremental rather than transformative departures from conventional practice. Critically, the distinction between soft enforcement (penalization) and hard enforcement (structural constraint) remains insufficiently theorized, particularly with respect to learning dynamics, expressivity, and epistemic validity.

Taken together, the literature reveals a conceptual opening: reframing physics from a corrective signal into a defining constraint on admissible representations. Such a shift has the potential to address persistent shortcomings in consistency, scalability, and the generation of scientific insight, while aligning machine learning more closely with the epistemic norms of materials science.

Proposed conceptual framework

The proposed framework advances this reframing by conceptualizing physics not as supplementary input, but as a definitional constraint on the model's representational capacity. Learning is thus reinterpreted as exploration within a physically admissible function space, rather than optimization over an unconstrained hypothesis class with penalized violations. Four core components structure this framework.

1. Constraint manifold definition

Fundamental physical principles—including conservation laws, symmetry groups, and thermodynamic inequalities—are used to define a constraint manifold that delimits the space of permissible functions. Instead of minimizing residuals, model outputs are parameterized or projected to reside strictly within this manifold. This ensures that physical admissibility is preserved globally and by construction, rather than approximately enforced during training.

2. Architectural embedding of physics

Physical constraints are embedded directly into network architecture through mechanisms such as equivariant layers, conserved quantity trackers, or operator-theoretic constraints. In materials science, this may involve encoding crystal symmetry groups, enforcing positive-definite energy functionals, or embedding dissipation mechanisms into network topology. Architecture thus becomes an expression of physical law, not merely a computational scaffold.

3. Data–constraint feedback loop

Within the constrained manifold, data serve to refine parameters rather than arbitrate between physics and empirical fit. Violations of physical principles trigger structural adaptation—such as reparameterization or manifold refinement—rather than being absorbed as residual loss. This creates an iterative feedback loop in which data and physics co-evolve, reinforcing consistency while retaining adaptability.

4. Epistemic steering and trade-off analysis

The framework explicitly prioritizes physical consistency over representational flexibility. Trade-offs among expressivity, accuracy, and robustness are analyzed using information-geometric perspectives on constrained spaces, enabling a principled assessment of model capacity under physical constraints. Learning thus becomes an epistemically guided search within a physically meaningful domain.

Implications

By shifting physics from input to constraint, this framework fundamentally alters learning dynamics. Optimization is no longer a negotiation between competing loss terms, but a structured search within physically bounded domains. This reorientation enhances robustness under sparse, noisy, or biased data conditions endemic to materials experimentation—while improving generalization, interpretability, and scientific legitimacy. The proposed perspective lays conceptual groundwork for next-generation materials AI systems that are not merely physics-aware, but physics-defined. As illustrated in **Figure 1**, the proposed constraint-centric paradigm shifts physics from a regularizing signal to a structural boundary on representational capacity.

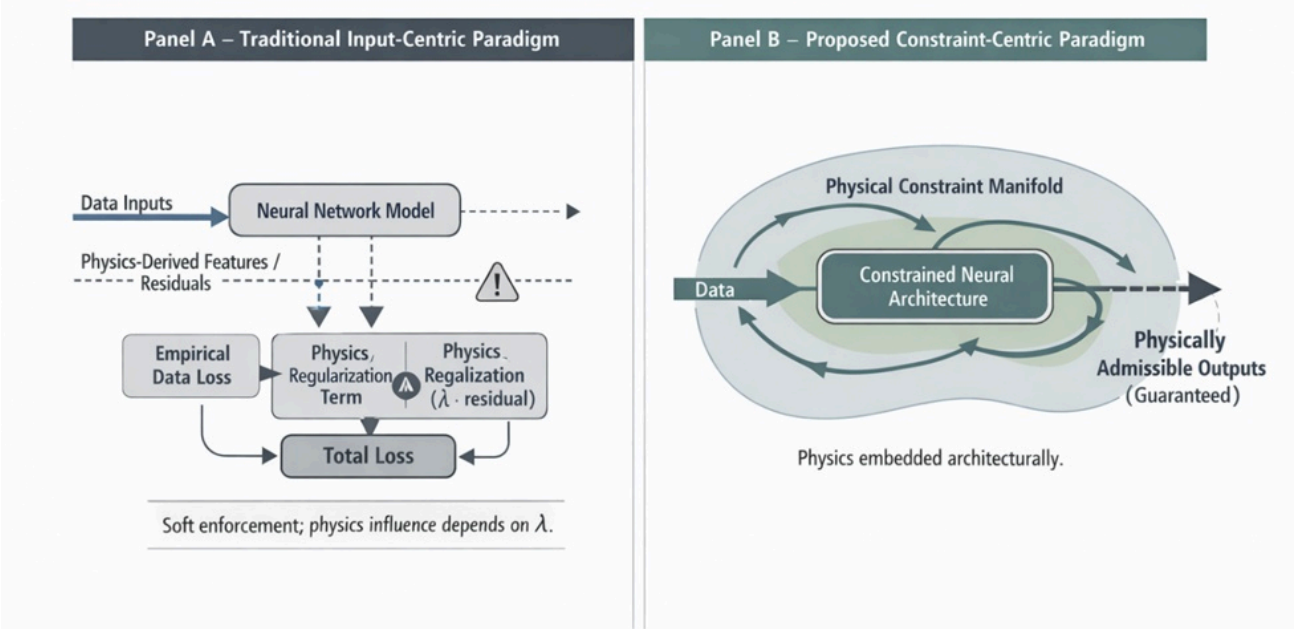

Figure 1. Conceptual contrast between input-centric and constraint-centric physics integration in machine learning for materials science.

Table 1 synthesizes the core epistemic and architectural distinctions between conventional input-centric physics-guided machine learning and the proposed constraint-centric paradigm. The comparison highlights how shifting physics from a regularizing signal to a structural boundary redefines hypothesis spaces, learning dynamics, uncertainty behavior, and scientific trustworthiness.

Table 1. Epistemic and architectural contrast between input-centric and constraint-centric physics-guided machine learning in materials science

Dimension	Input-centric PGML (Physics as input/regularization)	Constraint-centric PGML (Physics as structural constraint)
Role of physics	Auxiliary guidance applied through features or loss penalties	Definitional boundary that delimits admissible representations
Epistemic status of physics	External advisor competing with data during optimization	Constitutive principle defining the space of possible models
Mode of enforcement	Soft enforcement via weighted residuals (λ-dependent)	Hard enforcement via architectural embedding or manifold restriction
Hypothesis space	Unconstrained; physics violations permitted during and after training	Restricted a priori to physically admissible functions
Optimization dynamics	Trade-off between data fidelity and physics consistency	Search and refinement within a physically bounded domain
Sensitivity to hyperparameters	High; performance and consistency depend on λ tuning	Low; consistency guaranteed independently of loss weighting
Behavior under data scarcity	Improved over purely data-driven models, but is still vulnerable to violations	Robust; physical constraints stabilize learning with minimal data

Generalization and extrapolation	Limited; may fail under compositional or regime shifts	Enhanced; extrapolation confined to physically plausible regimes
Multiscale consistency	Requires ad-hoc coupling or additional regularization	Emerges naturally from constraint propagation across scales
Interpretability	Physics-constrained outputs, but internal representations remain opaque	Architecture reflects physical structure, enabling mechanistic insight
Uncertainty characteristics	Includes risk of structural (physical) violations	Confined to parametric uncertainty within an admissible manifold
Inverse problem stability	Sensitive to loss balancing and initialization	Stabilized by a priori admissible solution space
Expressivity	High flexibility, including non-physical solutions	Reduced flexibility but guaranteed physical fidelity
Failure modes	Physically implausible predictions with good numerical fit	Potential underfitting where physics is incomplete or emergent
Ethical and scientific trust	Conditional; requires post-hoc validation	High; violations of core principles are structurally impossible
Best-suited use cases	Exploratory modeling, weakly understood phenomena	Safety-critical modeling, thermodynamics, stability, and design feasibility
Epistemic trade-off	Flexibility prioritized over consistency	Consistency prioritized over flexibility
Conceptual framing	Physics as a corrective signal	Physics as a boundary condition of learning
Long-term scientific role	Accelerates approximation	Supports theory-aligned, trustworthy AI for materials science

Analytical implications

The reframing of physics as a hard constraint rather than an input generates several analytical implications for the epistemic structure of machine learning in materials science. At its core, this approach restricts the hypothesis space to a manifold in which physical principles are satisfied by construction, rather than approximately through optimization. This structural embedding shifts the learning process from a competition between empirical fidelity and physical regularization to a refinement within a bounded domain of physically admissible functions.

One key implication concerns epistemic uncertainty. In conventional PGML, uncertainty arises from both data noise and the potential for constraint violation, as soft penalties permit solutions that deviate from governing laws under strong data influence [1, 26-29]. Constraint-centric designs mitigate this by eliminating the possibility of unphysical predictions within the model's representational capacity. The manifold acts as an intrinsic regularizer, channeling uncertainty toward parametric variations within physically consistent regimes rather than structural violations. This yields greater robustness in extrapolation scenarios, where materials data are often sparse or confined to narrow regimes due to experimental limitations.

Systems-level insights emerge from the interaction dynamics between data and constraints. Data-driven updates occur along trajectories confined to the constraint manifold, creating feedback loops that couple parameter refinement to the preservation of invariants. For instance, in modeling multiscale material behavior, constraints on energy conservation or symmetry properties ensure that coarse-scale predictions remain consistent with fine-scale physics without requiring

explicit bridging terms. This structural coherence reduces the epistemic load on data, enabling models to infer plausible behaviors from limited observations while maintaining thermodynamic consistency [5, 27].

Trade-offs in expressivity and consistency become analytically transparent. Input-based methods offer high flexibility, allowing the model to approximate complex relations beyond known physics, yet at the risk of overfitting to artifacts or violating fundamental principles [7, 28]. Constraint-based approaches sacrifice some expressivity in regions where physics is incomplete or approximate—such as phenomenological models of plasticity or fracture—but gain guaranteed adherence to core invariants. This trade-off manifests in information-geometric terms: the effective dimensionality of the hypothesis space contracts to the manifold's intrinsic dimension, reducing sensitivity to distributional shifts in training data [11, 29].

In inverse problems prevalent in materials science, such as parameter identification from microstructural imaging or full-field measurements, the framework alters inference dynamics. Traditional methods balance data mismatch and physics residuals, often requiring careful weighting to avoid ill-posedness [10, 30]. By embedding constraints architecturally, the inference confines solutions to the admissible set a priori, transforming the problem from a penalized optimization to a constrained search. This reduces sensitivity to initialization and enhances convergence in ill-conditioned settings, where data alone may not uniquely determine parameters.

For multiscale modeling, the constraint paradigm supports hierarchical integration. Lower-scale constraints—such as lattice symmetries or conservation laws—propagate upward through equivariant or projection-based architectures, ensuring consistency across scales without ad-hoc homogenization assumptions [19]. Emergent properties arise from the interplay: models naturally satisfy macroscopic balance laws as consequences of microscopic constraints, thereby fostering interpretability in complex systems such as composites or metamaterials [9].

Ethical and epistemic considerations also surface. By prioritizing physical fidelity, the approach aligns model outputs with established scientific principles, reducing risks of deploying predictions that contradict foundational knowledge. This is particularly salient in materials design, where unphysical solutions could lead to infeasible or unsafe materials. The framework thus promotes epistemic humility, acknowledging that data-driven flexibility must be bounded by theoretical limits [12].

Overall, these implications reorient PGML toward architectures that embed physics as a foundational structure, yielding systems in which consistency is not negotiated but inherent.

Results and Discussion

The proposed reframing directly addresses persistent conceptual challenges in physics-guided machine learning (PGML) by repositioning physics from an auxiliary heuristic to a definitional boundary on admissible model behavior. Conventional input-centric paradigms have undeniably advanced predictive performance in data-scarce regimes, particularly in materials systems where experiments are expensive and slow. However, their reliance on soft enforcement mechanisms—typically through loss regularization—permits physical violations when optimization dynamics favor empirical fit, especially under data imbalance or noise [6, 21]. In extrapolation-heavy tasks, such violations undermine scientific trust, as predictions may achieve numerical accuracy while violating fundamental physical principles.

Constraint-centric designs offer a qualitatively different epistemic stance. By enforcing physical laws intrinsically through architecture or parameterization, these approaches shift learning from balancing heterogeneous loss terms toward structured exploration within a physically bounded hypothesis space. This reorientation transforms physics from a corrective signal into a constitutive element of representation, ensuring fidelity by construction rather than by approximation. As a result, consistency is no longer contingent on tuning trade-off parameters, but is instead guaranteed across the entire domain of inference.

This shift also illuminates a broader tension in knowledge integration strategies. Input-centric approaches emphasize flexibility, enabling discovery in regimes where governing physics is incomplete, uncertain, or poorly characterized. Constraint-centric paradigms, by contrast, prioritize reliability and coherence where physical laws are well-established. In materials science, many critical phenomena—such as thermodynamic stability, elastic response, and transport processes—are governed by conservation laws and symmetry constraints. For such tasks, including thermodynamic modeling, stability analysis, and safety-critical design, constraint-centric formulations offer clear conceptual advantages by preventing a priori physically inadmissible predictions [15, 25].

Nevertheless, the limitations of the proposed reframing are primarily conceptual rather than empirical. Hard architectural constraints may restrict expressivity in systems where governing principles are approximate, scale-dependent, or emergent rather than fundamental. In these contexts, overly rigid enforcement of constraints risks excluding valid but unconventional behaviors. This motivates the exploration of hybrid strategies that combine hard constraints on well-established principles with controlled flexibility in poorly understood regimes. Additionally, scalability remains a concern: enforcing constraints via manifold projection or operator-level embedding can increase computational complexity, particularly in high-dimensional or multiscale settings [22, 24].

Looking forward, several avenues emerge for advancing constraint-centric paradigms. One promising direction involves adaptive constraints that evolve as data accumulates, allowing the admissible manifold itself to be refined in response to new evidence. Such approaches could reconcile rigidity with learning, enabling progressive incorporation of newly discovered physical relationships. Integration with symbolic regression or theory-discovery methods presents another opportunity, potentially enabling the identification of novel invariants or governing relations within constrained spaces [7, 27]. Together, these directions suggest that constraint-centric learning need not be static, but can support iterative refinement of physical understanding.

Overall, the reframing contributes a conceptual lens rather than a prescriptive algorithm. It complements existing PGML approaches by clarifying when and why intrinsic physical enforcement is epistemically advantageous, guiding architectural design toward deeper physical coherence rather than incremental regularization.

Conclusion

This manuscript advances a conceptual reframing of physics-guided machine learning in materials science by positioning physics as a hard constraint on the hypothesis space rather than as an additive input or penalty term. By embedding invariants, conservation laws, and symmetries directly into model architecture, the proposed framework ensures physical consistency by construction, fundamentally altering the interaction dynamics between data and prior knowledge.

The analytical implications of this reframing include enhanced epistemic robustness, reduced dependence on large labeled datasets, and improved coherence across scales in complex materials systems. By constraining representational capacity to physically admissible domains, the framework clarifies trade-offs between expressivity and fidelity, offering system-level insight into uncertainty management, extrapolation reliability, and inference under constraint.

While input-centric methods remain valuable for exploratory modeling and discovery in poorly characterized regimes, constraint-centric designs are better suited to domains governed by well-established physical principles. This perspective steers future theoretical and architectural developments toward learning systems that inherently respect physical law, fostering greater interpretability, trustworthiness, and scientific legitimacy in materials modeling and design.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 13 Jun 2024 Revised: 10 Jul 2024 Accepted: 16 Aug 2024
Published online: 18 January 2025

References

Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. Nat Rev Phys. 2021;3(6):422-40.

Lu L, Meng X, Mao Z, Karniadakis GE. Deepxde: A deep learning library for solving differential equations. SIAM Rev. 2021;63(1):208-28.

Chen Y, Lu L, Karniadakis GE, Dal Negro L. Physics-informed neural networks for inverse problems in nano-optics and metamaterials. Opt Express. 2020;28(8):11618-33.

Haghighat E, Juanes R. Sciann: A keras/tensorflow wrapper for scientific computations and physics-informed deep learning using artificial neural networks. Comput Methods Appl Mech Eng. 2021;373:113552.

Zhang E, Dao M, Karniadakis GE, Suresh S. Analyses of internal structures and defects in materials using physics-informed neural networks. Sci Adv. 2022;8(7):eabk0644.

Miao ZW, Chen Y. Physics-informed neural networks method in high-dimensional integrable systems. Mod Phys Lett B. 2022;36(01):2150531.

Hussain A, Sakhaei AH, Shafiee M. Machine learning-based constitutive modelling for material non-linearity: A review. Mech Adv Mater Struct. 2024:1-9.

Zhang Z, Gu GX. Physics-informed deep learning for digital materials. Theor Appl Mech Lett. 2021;11(1):100220.

Liu X, Qin J, Zhao K, Featherston CA, Kennedy D, Jing Y, et al. Design optimization of laminated composite structures using artificial neural network and genetic algorithm. Compos Struct. 2023;305:116500.

Hamel CM, Long KN, Kramer SL. Calibrating constitutive models with full-field data via physics informed neural networks. Strain. 2023;59(2):e12431.

Faroughi SA, Pawar N, Fernandes C, Raissi M, Das S, Kalantari NK, et al. Physics-guided, physics-informed, and physics-encoded neural networks in scientific computing. arXiv preprint arXiv:2211.07377. 2022.

Wang S, Teng Y, Perdikaris P. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. SIAM J Sci Comput. 2021;43(5):A3055-81.

Lee S, Popovics J. Applications of physics-informed neural networks for property characterization of complex materials. RILEM Tech Lett. 2022;7:178-88.

Khatib O, Ren S, Malof J, Padilla WJ. Informed deep learning in metamaterials. In: 2023 International Applied Computational Electromagnetics Society Symposium (ACES); 2023; IEEE. p. 1-2.

Yang S, Peng S, Guo J, Wang F. A review on physics-informed machine learning for monitoring metal additive manufacturing process. Adv Manuf. 2024;1(2):1-2.

Diao Y, Yang J, Zhang Y, Zhang D, Du Y. Solving multi-material problems in solid mechanics using physics-informed neural networks based on domain decomposition technology. Comput Methods Appl Mech Eng. 2023;413:116120.

Jeong I, Cho M, Chung H, Kim DN. Data-driven nonparametric identification of material behavior based on physics-informed neural network with full-field data. Comput Methods Appl Mech Eng. 2024;418:116569.

Jeong H, Batuwatta-Gamage C, Bai J, Xie YM, Rathnayaka C, Zhou Y, et al. A complete physics-informed neural network-based framework for structural topology optimization. Comput Methods Appl Mech Eng. 2023;417:116401.

Han J, Lee Y. Hierarchical learning to solve partial differential equations using physics-informed neural networks. arXiv preprint arXiv:2112.01254. 2021.

Khatib O, Ren S, Malof J, Padilla WJ. Metamaterial design with physics informed neural networks. In: Photonic and Phononic Properties of Engineered Nanostructures XIII; 2023; SPIE. p. PC124310O.

Anton D, Wessels H. Physics-informed neural networks for material model calibration from full-field displacement data. arXiv preprint arXiv:2212.07723. 2022.

Zhou G, Hu Y, Cao Z, Lee MG, Li D. A physics-constrained neural network for crystal plasticity modelling of fcc materials. Scr Mater. 2024;241:115861.

Cheng Z, Wang K, Tanvir AM, Shang W, Luo T, Zhang Y, et al. Bayesian optimization of low-temperature nonthermal plasma jet sintering of nanoinks. ACS Appl Mater Interfaces. 2024;16(35):46897-908.

Zhao C, Zhang F, Lou W, Wang X, Yang J. A comprehensive review of advances in physics-informed neural networks and their applications in complex fluid dynamics. Phys Fluids. 2024;36(10).

Cai S, Mao Z, Wang Z, Yin M, Karniadakis GE. Physics-informed neural networks (pinns) for fluid mechanics: A review. Acta Mech Sin. 2021;37(12):1727-38.

Shukla K, Jagtap AD, Karniadakis GE. Parallel physics-informed neural networks via domain decomposition. J Comput Phys. 2021;447:110683.

Portela CM, Vidyasagar A, Krödel S, Weissenbach T, Yee DW, Greer JR, et al. Extreme mechanical resilience of self-assembled nanolabyrinthine materials. Proc Natl Acad Sci U S A. 2020;117(11):5686-93.

Cai S, Wang Z, Wang S, Perdikaris P, Karniadakis GE. Physics-informed neural networks for heat transfer problems. J Heat Transfer. 2021;143(6):060801.

Zhang E, Yin M, Karniadakis GE. Physics-informed neural networks for nonhomogeneous material identification in elasticity imaging. arXiv preprint arXiv:2009.04525. 2020.

Haghighat E, Raissi M, Moure A, Gomez H, Juanes R. A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics. Comput Methods Appl Mech Eng. 2021;379:113741.