

ORIGINAL RESEARCH

Open access

Perspective: Embedding Domain Knowledge Is Not Optional — A Position on Physics-Constrained Materials GNNs

Carlos Ramirez^{1*}, Elena Torres², Pablo Ortega¹, Sofia Mendes³

Abstract

Embedding domain knowledge into materials graph neural networks (GNNs) is not an optional enhancement for performance tuning; it is a fundamental requirement for models that must extrapolate reliably, operate with acceptable sample efficiency, and deliver physically consistent predictions. Purely data-driven architectures, which treat materials systems as generic graphs without built-in physical inductive biases, consistently fail when deployed beyond their training distributions — precisely the regime in which materials discovery operates. This Perspective articulates a clear position: the continued reliance on unconstrained, black-box GNNs represents a dead end for computational materials engineering. Three interlocking reasons compel this stance. First, materials discovery demands extrapolation to unseen compositions, structures, and conditions; unconstrained models offer no guarantees beyond interpolation. Second, density-functional theory (DFT) data remain scarce and computationally prohibitive, rendering sample-inefficient architectures unsustainable. Third, predictions must obey conservation laws, symmetry requirements, and thermodynamic limits; violations render long-term simulations unstable and untrustworthy. Four classes of domain knowledge must be embedded as hard constraints: (1) symmetry (E(3) equivariance, permutation invariance, space-group symmetries), (2) conservation laws (energy conservation, momentum balance), (3) physical scales and units (bounded energies, forces, and lengths), and (4) locality and smoothness principles (finite cutoffs, hierarchical many-body interactions). Recent literature provides compelling evidence that physics-constrained models — notably equivariant architectures — achieve comparable accuracy with an order-of-magnitude reduction in training data while maintaining physical consistency. Conversely, ignoring these constraints leads to extrapolation collapse, massive data waste, unphysical molecular-dynamics drift, and models that cannot be interpreted or transferred. The materials community must therefore treat physics-constrained design as the default, not an afterthought. We recommend concrete standards for model development, benchmark construction, and peer review. Only by making domain knowledge non-optional can machine learning accelerate, rather than merely decorate, the discovery of next-generation materials.

Keywords Materials graph neural networks, Physics-constrained machine learning, Domain knowledge embedding, Equivariant architectures, Inductive biases, Extrapolation failure

*Correspondence:

Carlos Ramirez
carlos.ramirez@gmail.com

¹ Department of Data-Driven Materials Science, Faculty of Engineering, University of Barcelona, Barcelona, Spain

² Department of Computational Materials Modeling, Faculty of Technology, University of Lisbon, Lisbon, Portugal

³ Department of Intelligent Materials Systems, Faculty of Engineering, University of Porto, Porto, Portugal

Introduction

The field of artificial intelligence for materials science stands at a crossroads. After a decade of impressive leaderboard gains, it is time to confront a structural limitation that purely data-driven approaches cannot overcome.

Position Statement Embedding domain knowledge into materials GNNs is not an optional “add-on” for improving performance. It is a requirement for building models that extrapolate reliably, achieve acceptable sample efficiency, and produce physically consistent predictions. The continued pursuit of purely data-driven architectures without physics constraints is a dead end for materials discovery.

This position is not a call to abandon data or neural networks. It is a call to stop pretending that materials are simply another domain for generic graph learning. Materials data are scarce, expensive, and heavily biased toward stable, ordered structures [1-5]. Real discovery requires predicting metastable phases, disordered alloys, extreme conditions, and entirely new chemical spaces — regimes where interpolation guarantees evaporate. Pure black-box models, no matter how deep or how many parameters they contain, cannot supply the inductive biases that physics provides for free.

The urgency is immediate. High-throughput DFT campaigns have generated impressive but still limited datasets. Generating additional data for every new material family is economically and environmentally unsustainable. At the same time, industry and national laboratories demand force fields and property predictors that remain stable over nanoseconds of molecular dynamics, transfer across chemistries, and respect thermodynamic limits. Unconstrained GNNs repeatedly fail these tests. They achieve low test-set errors on random splits yet produce negative formation energies, non-conservative forces, or exploding simulations when deployed.

The community has already seen the warning signs. Early successes with message-passing networks were hailed as universal approximators, yet subsequent studies revealed that these same networks collapse outside their training distributions. The solution is not more data or larger models; it is the deliberate injection of known physics as architectural constraints [6-11]. Equivariant networks that respect $E(3)$ symmetry, energy-conserving architectures that guarantee forces derive from a scalar potential, and locality-aware designs that mirror the finite-range nature of interatomic interactions all demonstrate superior generalization with dramatically fewer training points [12-19].

This Perspective therefore rejects the narrative that domain knowledge is merely “helpful.” It is essential. The remainder of the paper develops three reasons why constraints are non-optional, identifies the precise types of knowledge that must be embedded, marshals evidence from the recent literature, and details the concrete failures that arise when physics is ignored. The goal is not incremental improvement but a paradigm shift: physics-constrained materials GNNs must become the community standard.

Figure 1 maps the manuscript’s central argument as a directional decision structure, showing how the three core failures of unconstrained materials GNNs motivate four mandatory classes of embedded domain knowledge and culminate in concrete standards for trustworthy materials ML.




Figure 1. Why Domain Knowledge Is Non-Optional in Materials GNNs: From Structural Failures of Black-Box Models to Physics-Constrained Design Standards

Three Reasons Why Domain Knowledge Is Not Optional

The case for embedding domain knowledge rests on three interlocking failures of purely data-driven materials GNNs.

Materials discovery is intrinsically an out-of-distribution problem, as the search for novel compositions, polymorphs, and operating regimes necessarily extends beyond the support of existing datasets [10, 20]. Models trained on equilibrium structures therefore lack any principled mechanism to enforce physical validity once they depart from the training manifold. While interpolation within the convex hull of observed data may remain well behaved, predictions in extrapolative regimes often become physically implausible. The imposition of symmetry, conservation laws, and bounded physical scales constrains the hypothesis space to admissible configurations, thereby anchoring extrapolation; absent such constraints, the model remains susceptible to generating nonphysical artifacts such as negative energies or non-conservative force fields [6, 11].

This limitation becomes more acute under the stringent data regimes characteristic of atomistic modeling. Although density functional theory remains the benchmark for ground-truth data, its computational cost renders large-scale dataset expansion impractical. Within this context, the efficiency gains reported by Batzner *et al.* [21], where E(3)-equivariant architectures match invariant baselines using an order of magnitude fewer samples, underscore the structural advantage of embedding physical priors. Subsequent results by Batatia *et al.* [22] further indicate that increasing the expressivity of equivariant message passing compounds these gains. For complex material classes such as high-entropy alloys, amorphous phases, and interfacial systems, where data acquisition is severely constrained, unconstrained models effectively incur a prohibitive computational burden [18, 20]. Incorporating invariances and conservation principles thus substitutes learned regularities with encoded structure, directly reducing data requirements.

The implications extend beyond efficiency to the internal consistency of predicted physical quantities. Reliable deployment in downstream simulations requires strict adherence to foundational laws: forces must derive as gradients of scalar energies, total momentum must remain conserved in isolated systems, and predicted magnitudes must align with chemically and thermodynamically plausible scales. Data-driven models that lack these constraints frequently violate such conditions, leading to energy drift, distorted bonding geometries, and thermodynamic inconsistencies that

compromise molecular dynamics trajectories. Embedding these constraints at the architectural level precludes such failures, avoiding reliance on corrective strategies applied after training.

These limitations are deeply coupled rather than independent. Deficient extrapolation reflects underlying inefficiencies in data utilization, which in turn magnify violations of physical consistency. Treating domain knowledge as optional perpetuates this feedback loop, whereas its integration restructures the learning problem itself, aligning model behavior with the governing principles of the physical systems under study.

Table 1 clarifies that the weaknesses of unconstrained materials GNNs are not isolated technical defects but structurally linked failures that directly undermine discovery-oriented deployment.

Table 1. Why Unconstrained Materials GNNs Fail on Discovery-Relevant Tasks

Discovery requirement	What the task demands	Behavior of unconstrained GNNs	Why the failure occurs	Practical consequence for materials discovery	Why embedded domain knowledge resolves it
Extrapolation to unseen compositions, structures, and conditions	Reliable prediction outside the empirical support of training data	Performance degrades sharply once inputs move beyond equilibrium-like or previously observed regions	The model has no built-in restriction to physically admissible responses outside the training manifold	Novel chemistries, metastable phases, and extreme-condition predictions become unreliable	Symmetry, conservation, bounded scales, and locality shrink the hypothesis space to physically plausible behaviors
Learning under DFT scarcity	Competitive accuracy with limited and expensive supervision	Requires far larger datasets to infer regularities that physics already specifies	Invariances and admissibility conditions must be learned indirectly from data rather than encoded directly	Data-generation cost becomes computationally and economically unsustainable	Hard-coded inductive biases reduce sample complexity and improve learning efficiency
Force-field usability in simulation	Energy-force consistency and long-horizon dynamical stability	Low held-out error may coexist with unstable trajectories and drift	Outputs are not structurally tied to scalar potentials or conservation rules	Molecular dynamics becomes untrustworthy despite strong benchmark scores	Energy-conserving architectures enforce consistency by construction
Cross-chemistry transfer	Adaptation across related but nonidentical material families	Fine-tuning is brittle and often requires extensive retraining	Learned representations are dataset-specific rather than physically grounded	Transfer learning becomes expensive, uncertain, and domain-fragmented	Physics-based constraints provide reusable structure across chemical families
Mechanistic auditability	Ability to diagnose why the model	Errors remain opaque and difficult to localize	The architecture does not separate physically	Model debugging becomes guesswork and	Constraint-aware designs expose violation modes

	succeeds or fails		meaningful sources of failure	benchmark gains are hard to interpret	such as symmetry error or conservation inconsistency
--	-------------------	--	-------------------------------	---------------------------------------	--

Types of Domain Knowledge That Must be Embedded

Four classes of domain knowledge must be injected into materials GNNs as hard architectural constraints rather than soft regularizers.

Crystalline materials are governed by E(3) equivariance, wherein energies remain invariant and forces transform covariantly under rotations, translations, and inversions [16, 17, 19]. Enforcing these symmetries has been shown to substantially improve data efficiency, as demonstrated by Batzner *et al.* and Batatia *et al.* [21, 22], by aligning model structure with known physical invariances. At the same time, permutation invariance among indistinguishable atoms must be embedded directly within the message-passing mechanism, while discrete symmetries such as time-reversal in non-magnetic systems and magnetic space group constraints in spin-polarized materials further restrict the admissible function space [13]. Neglecting these conditions forces the model to expend capacity rediscovering symmetries that are already prescribed by physical law.

A related constraint emerges from conservation principles, which impose strict relationships between predicted quantities. Energy conservation requires forces to be derivable as exact gradients of a scalar potential, $F_i = -\nabla_i E$ while momentum conservation enforces a vanishing net force in isolated systems, and thermodynamic consistency couples free energies, entropies, and temperature-dependent observables. Architectures that fail to encode these relations yield force fields that are fundamentally incompatible with molecular dynamics [6, 11]. By structuring the network to predict energies and derive forces analytically, such constraints are satisfied by construction rather than approximated through auxiliary objectives.

Beyond formal invariances, physically meaningful predictions must respect characteristic scales and units intrinsic to chemical systems. Bond lengths are bounded, formation energies for stable compounds lie within a narrow energetic window, and force magnitudes reflect underlying interatomic potentials. In the absence of such priors, neural networks readily produce outputs that deviate by orders of magnitude, particularly in extrapolative regimes. Incorporating scale-aware normalization, unit-consistent parameterization, and bounded representations constrains the output space and stabilizes optimization [9, 10].

These structural considerations extend naturally to the spatial organization of interactions. Interatomic forces decay rapidly beyond a few angstroms, motivating finite cutoff radii, while potential energy surfaces remain smooth except in regions associated with bond rearrangement or electronic degeneracy. Many-body interactions exhibit a hierarchical structure, with pairwise, three-body, and higher-order contributions dominating at distinct length scales. Encoding locality through cutoff-aware graph construction and reflecting this hierarchy in message passing reduces unnecessary complexity while preserving essential physics [12, 14, 15].

Taken together, these constraints delineate the admissible hypothesis space for graph neural networks in materials modeling. Omitting them does not merely degrade performance but alters the problem formulation itself, yielding models that are misaligned with the governing structure of the underlying physical system.

Table 2 consolidates the manuscript's normative claim by specifying which forms of domain knowledge should be treated as architectural requirements, what each constraint governs, and what failure mode each prevents.

Table 2. Mandatory Classes of Domain Knowledge for Physics-Constrained Materials GNNs

Domain-knowledge class	What must be enforced architecturally	Functional role in the model	Typical implementation logic	Failure mode when omitted	Expected modeling benefit
Symmetry constraints	E(3) equivariance, translational invariance, permutation invariance, and where relevant crystal or magnetic symmetry structure	Ensures predictions transform correctly under physically valid coordinate operations	Equivariant message passing; invariant scalar outputs for energies; vector-consistent force representations	Capacity is wasted relearning guaranteed symmetries; extrapolation becomes fragile; outputs may transform incorrectly	Stronger generalization, higher sample efficiency, and physically meaningful coordinate behavior
Conservation laws	Forces derived from scalar energy; momentum-compatible interactions; thermodynamic consistency where relevant	Guarantees internal coherence of predictions used in downstream simulation	Energy-first architectures with analytical gradients; constraint-preserving formulations	Long-term drift, non-conservative dynamics, and simulation instability	Stable molecular dynamics and trustworthy force-field deployment
Physical scales and units	Valid ranges, unit-consistent normalization, chemistry-aware magnitudes for energies, forces, and lengths	Prevents scale pathologies and anchors outputs to realistic material regimes	Physically informed normalization, bounded outputs, and unit-aware preprocessing or parameterization	Extreme or nonsensical outputs under distribution shift	Faster training, safer extrapolation, and more interpretable predictions
Locality and smoothness	Finite interaction cutoffs, smooth potentials, and hierarchical many-body structure	Aligns model complexity with actual interatomic physics	Cutoff-based graph construction; local neighborhoods; layered interaction orders	Overly global or noisy representations, poor efficiency, and unstable behavior near perturbed configurations	Better inductive fit to materials physics with lower computational burden
Constraint integration strategy	Constraints treated as hard design principles rather than optional regularizers	Determines whether physics is guaranteed or merely encouraged	Architectural enforcement preferred over penalty-only training	Constraint violations persist despite improved loss values	More robust and auditable behavior across training and deployment settings

Evidence from the Literature

The literature now contains decisive evidence that physics-constrained GNNs outperform unconstrained counterparts on precisely the metrics that matter for materials discovery.

Empirical evidence consistently indicates that embedding physical structure into graph neural networks fundamentally alters their learning dynamics [5, 7, 8]. The introduction of E(3)-equivariant architectures by Batzner *et al.* demonstrated that NequIP attains parity with invariant models such as SchNet while requiring an order of magnitude fewer training samples, a result subsequently extended by Batatia *et al.* through the MACE framework, which further improves both accuracy and extrapolative robustness [21, 22]. Such gains are not attributable to incremental scaling but to the direct incorporation of symmetry constraints, a pattern replicated across crystal-graph models and higher-order equivariant formulations in diverse material settings [3, 4, 12–14, 16–19, 23, 24]. This efficiency advantage is mirrored in the behavior of energy-conserving force fields, where enforcing consistency between energies and forces suppresses the accumulation of numerical drift that otherwise leads to unphysical thermal artifacts in molecular dynamics [6, 11]. When forces are learned independently, even minor inconsistencies compound over time, whereas architectures that derive forces from scalar potentials eliminate this instability at its origin, enabling stable trajectories over extended timescales. The limitations of unconstrained approaches become particularly visible under realistic evaluation protocols: random train–test splits obscure generalization failures, as models trained on ordered crystals often collapse when applied to disordered alloys or high-temperature phases. In such regimes, the absence of physical priors prevents reliable identification of metastable configurations or extrapolation across compositional space. Although narrowly defined classification tasks, such as space-group prediction, may show limited sensitivity to these constraints, the broader literature on property prediction and interatomic potential learning leaves little ambiguity regarding their necessity.

The practical consequences of neglecting domain knowledge follow directly from these empirical patterns and manifest as a cascade of failures that compromise both reliability and scientific utility [20, 25–29]. Apparent success under interpolation masks severe degradation when models encounter novel material families, where extrapolative predictions become unreliable despite strong benchmark performance. This fragility is compounded by pronounced data inefficiency, as unconstrained architectures demand orders of magnitude more density functional theory calculations to achieve comparable accuracy, significantly inflating computational cost and extending discovery timelines. The absence of enforced physical consistency further leads to pathological behaviors, including energy drift, unstable bonding geometries, and simulations that either diverge or stagnate due to a lack of thermodynamic grounding. Interpretability deteriorates in parallel, as purely data-driven predictions offer no mechanism to disentangle errors arising from violated symmetries, broken conservation laws, or mis-scaled representations, rendering systematic debugging infeasible [25]. Transfer learning similarly degrades, with models failing to adapt across chemical domains because they lack invariant structural priors that would otherwise guide generalization. Under these conditions, unconstrained graph neural networks are not merely inefficient but fundamentally misaligned with the epistemic and operational demands of materials design.

Common objections to physics-informed modeling do not withstand scrutiny when evaluated against these constraints. Concerns that domain knowledge restricts flexibility overlook the fact that appropriate constraints, such as E(3) equivariance and energy conservation, eliminate nonphysical regions of parameter space while preserving all admissible configurations, thereby improving rather than limiting generalization, as evidenced by NequIP and MACE [21, 22]. Arguments invoking incomplete physical understanding conflate partial knowledge with absence of structure; incorporating established principles such as symmetry, locality, and conservation still yields substantial reductions in search complexity and consistently outperforms unconstrained baselines [21, 22]. Analogies to image recognition further obscure domain-specific requirements, since materials datasets are inherently sparse, costly, and extrapolation-driven, unlike the large-scale, interpolation-dominated regimes of computer vision; equivariant architectures succeed precisely because they encode this distinction [21, 22]. Implementation challenges likewise fail to justify omission, given the availability of validated libraries and the comparatively minor engineering overhead relative to the downstream cost of unstable simulations or excessive data generation [21, 22]. Finally, benchmark-centric claims of success rest on evaluation protocols that privilege interpolation and neglect physical consistency, thereby overstating model reliability. Under closer inspection, each objection ultimately reduces to a mischaracterization of the same underlying requirements—robust extrapolation, efficient data usage, and strict adherence to physical law—that render domain knowledge indispensable in materials machine learning.

Relation to Other Positions

This Perspective aligns with, yet sharpens, several existing strands in the literature.

Versus pure data-driven ML

We are not anti-data. Data remain indispensable for fitting the unknown residuals after physics is embedded. The position simply rejects the claim that data alone suffice. Purely data-driven GNNs treat materials as generic graphs; the result is models that scale poorly and generalize narrowly.

Versus full physics-based modeling

This is not a call to replace DFT with ML. First-principles methods will always anchor the most accurate predictions. The argument is that physics should constrain the ML layer rather than serve as its substitute. Hybrid approaches gain the speed of neural networks while inheriting the correctness guarantees of known physical laws.

Versus the broader “inductive bias” literature

The position is fully consistent with the wider machine-learning consensus that appropriate inductive biases improve generalization [26-28]. Equivariant networks are simply the materials-specific instantiation of this principle. The community has already accepted permutation invariance and cutoff locality; extending the same logic to E(3) symmetry and energy conservation is a natural, not revolutionary, step.

Versus interpretability research

Physics constraints often serve as built-in interpretability mechanisms. When a model is forced to output energies from which forces are derived analytically, the causal chain becomes transparent [25]. Symmetry violations can be audited layer by layer. In contrast, black-box force predictions offer no mechanistic explanation when they fail. Embedding domain knowledge therefore advances both predictive power and scientific understanding simultaneously.

Taken together, the position does not reject any of these neighboring fields; it integrates them under a single guiding imperative: domain knowledge must be architecturally enforced, not merely encouraged.

Recommendations for Practice

To translate this position into community norms, we offer concrete, actionable recommendations.

Table 3 translates the Perspective’s position into an actionable evaluation standard by linking model design, benchmark construction, and review criteria to the properties that discovery-ready materials ML must actually satisfy.

Table 3. A Field Standard for Evaluating Physics-Constrained Materials GNNs

Stakeholder domain	Required evaluation question	What should be reported	Minimum acceptable evidence	What should no longer count as sufficient on its own	Intended field-level effect

Model development	Which physical constraints are embedded, and why are they appropriate for the task?	Explicit statement of symmetry handling, conservation design, scale treatment, and locality assumptions	Architectural description plus ablation showing the effect of each constraint	Reporting predictive accuracy alone on random splits	Shifts model design from generic flexibility toward principled physical admissibility
Benchmark construction	Does the benchmark test interpolation only, or true discovery-relevant extrapolation?	Separate results for in-distribution and out-of-distribution settings; learning curves versus data size	Novel composition, disorder, or extreme-condition test sets with transparent split logic	Single random train-test split with aggregate error metrics only	Prevents overstated generalization claims
Physical consistency auditing	Are predictions usable in downstream simulation and decision contexts?	Force-energy consistency, symmetry-violation magnitude, long-horizon drift, and stability diagnostics	Demonstrated low violation rates and stable simulation behavior under deployment-like conditions	Low MAE without any physical-consistency reporting	Makes trustworthiness a first-class evaluation target
Transfer and reuse	Does the model adapt across chemistry families without brute-force retraining?	Cross-domain fine-tuning or transfer results with data budgets stated explicitly	Evidence that physically grounded structure supports adaptation under limited new supervision	Claiming transferability without out-of-family tests	Rewards reusable scientific representations rather than benchmark-specific tuning
Peer review and publishing	Can reviewers verify that the architecture is discovery-ready rather than merely leaderboard-optimized?	Constraint rationale, validation protocol, benchmark design, and code availability	Reproducible implementation and explicit answers to physics-consistency questions	Vague references to being “physics-informed” without operational evidence	Raises publication standards and improves comparability across studies

For model developers (a) Begin every new architecture with the strongest physics constraints that the task permits—E(3) equivariance for force fields, energy conservation by construction, and locality-aware graph construction. Only after these are satisfied should additional data-driven flexibility be introduced [9, 10]. (b) Adopt equivariant message-passing layers as the default backbone for interatomic potentials [21, 22]. (c) Mandate reporting of physical-consistency metrics: force–energy consistency error, symmetry-violation magnitude, and long-term energy drift in molecular-dynamics runs.

For benchmark designers (a) Include explicit extrapolation tasks—compositionally novel alloys, disordered structures, high-pressure polymorphs—that lie outside the convex hull of training data. (b) Publish learning curves that plot error against training-set size so that sample-efficiency gains from constraints become visible. (c) Require physical-consistency tests as first-class evaluation criteria rather than optional appendices.

For journal editors and reviewers (a) Add a standard referee question: “What physics constraints are embedded in the architecture, and how were they validated?” (b) Treat purely data-driven submissions for extrapolation-critical tasks (force fields, property prediction under novel conditions) with heightened scrutiny. (c) Encourage or require open-source release

of constrained reference implementations so that best practices can propagate rapidly.

Adopting these standards will shift the field from leaderboard chasing to trustworthy engineering. The cost is modest; the benefit is a generation of models that accelerate, rather than merely decorate, materials discovery.

Conclusion

Embedding domain knowledge into materials GNNs is not optional. It is a requirement. Purely data-driven architectures fail on extrapolation, sample efficiency, and physical consistency—the three pillars that materials discovery demands. Four classes of knowledge must be embedded as hard constraints: symmetry (E(3) equivariance and permutation invariance), conservation laws (energy and momentum), physical scales and units, and locality with hierarchical many-body interactions. Evidence from the recent literature, most notably the order-of-magnitude data-efficiency gains of NequIP and MACE, demonstrates that these constraints deliver exactly the generalization and stability that unconstrained models cannot achieve. Ignoring them produces extrapolation collapse, unsustainable data demands, unphysical simulations, opaque predictions, and poor transfer—outcomes the community can no longer tolerate.

The path forward is clear. Model developers must start with physics; benchmark designers must test extrapolation and consistency; journals must enforce these standards. Funding bodies must reward sample-efficient, physically grounded methods. Only by making domain knowledge architecturally non-optional can artificial intelligence fulfill its promise in materials science. The era of black-box GNNs for materials is over. The era of physics-constrained, trustworthy, and discovery-ready models has begun.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 27 Oct 2022

Revised: 08 Feb 2023

Accepted: 29 Apr 2023

Published online: 18 July 2023

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- 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).
<https://doi.org/10.1126/sciadv.1701816>.
- Schütt KT, Sauceda HE, Kindermans PJ, Tkatchenko A, Müller KR. SchNet: A deep learning architecture for molecules and materials. *J Chem Phys.* 2018;148(24):241722.
<https://doi.org/10.1063/1.5019779>.
- Xie T, Grossman JC. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys Rev Lett.* 2018;120(14):145301.
<https://doi.org/10.1103/PhysRevLett.120.145301>.
- Chen C, Ye W, Zuo Y, Zheng C, Ong SP. Graph networks as a universal machine learning framework for molecules and crystals. *Chem Mater.* 2019;31(9):3564-72.
<https://doi.org/10.1021/acs.chemmater.9b01294>.
- 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>.
- Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. *Nat Rev Phys.* 2021;3(6):422-40.
<https://doi.org/10.1038/s42254-021-00314-5>.
- Morgan JP, Paiement A, Klinken C. Domain-informed graph neural networks: A quantum chemistry case study. *Neural Netw.* 2023;165:938-52.
<https://doi.org/10.1016/j.neunet.2023.06.030>.
- Wang R, Zou Y, Zhang C, Wang X, Yang M, Xu D. Combining crystal graphs and domain knowledge in machine learning to predict metal-organic frameworks performance in methane adsorption. *Microporous Mesoporous Mater.* 2022;331:111666.
<https://doi.org/10.1016/j.micromeso.2021.111666>.
- Bødker ML, Bauchy M, Du T, Mauro JC, Smedskjaer MM. Predicting glass structure by physics-informed machine learning. *NPJ Comput Mater.* 2022;8(1):192.
<https://doi.org/10.1038/s41524-022-00882-9>.
- Khatamsaz D, Neuberger R, Roy AM, Zadeh SH, Otis R, Arróyave R. A physics informed bayesian optimization approach for material design: Application to NiTi shape memory alloys. *NPJ Comput Mater.* 2023;9(1):221.
<https://doi.org/10.1038/s41524-023-01173-7>.
- Wu JL, Xiao H, Paterson E. Physics-informed machine learning approach for augmenting turbulence models: A comprehensive framework. *Phys Rev Fluids.* 2018;3(7):074602.
<https://doi.org/10.1103/PhysRevFluids.3.074602>.
- Choudhary K, DeCost B. Atomistic line graph neural network for improved materials property predictions. *NPJ Comput Mater.* 2021;7(1):185.
<https://doi.org/10.1038/s41524-021-00650-1>.
- Jørgensen PB, Garijo del Río E, Schmidt MN, Jacobsen KW. Materials property prediction using symmetry-labeled graphs as atomic-position independent descriptors. *Phys Rev B.* 2019;100(10):104114.
<https://doi.org/10.1103/PhysRevB.100.104114>.

Cheng J, Zhang C, Dong L. A geometric-information-enhanced crystal graph network for predicting properties of materials. *Commun Mater.* 2021;2(1):92.

<https://doi.org/10.1038/s43246-021-00194-3>.

Gong S, Yan K, Xie T, Shao-Horn Y, Gomez-Bombarelli R, Ji S, et al. Examining graph neural networks for crystal structures: Limitations and opportunities for capturing periodicity. *Sci Adv.* 2023;9(45).

<https://doi.org/10.1126/sciadv.adi3245>.

Kaba SO, Ravanbakhsh S. Equivariant networks for crystal structures. *Adv Neural Inf Process Syst.* 2022;35:4150-64.

Pakornchote T, Ektarawong A, Chotibut T. StrainTensorNet: Predicting crystal structure elastic properties using SE(3)-equivariant graph neural networks. *Phys Rev Res.* 2023;5(4):043198.

<https://doi.org/10.1103/PhysRevResearch.5.043198>.

Landes FP, Furtlehner C. Equivariant graph neural networks for amorphous materials [Internet]. Gif-sur-Yvette: LISN, Université Paris-Saclay; 2021 [cited 2026 Jun 21]. Available from:

https://www.lri.fr/~gcharpia/LISN_INRIA_Equivariant_GNN_for_amorphous_materials.pdf

Satorras VG, Hoogeboom E, Welling M. E(n) equivariant graph neural networks. In: Proceedings of the 38th International conference on machine learning. *Proc Mach Learn Res.* 2021;139:9323-32.

Sudmanns M, Bach J, Weygand D, Schulz K. Data-driven exploration and continuum modeling of dislocation networks. *Model Simul Mater Sci Eng.* 2020;28(6):065001.

Batzner S, Musaelian A, Sun L, Geiger M, Mailoa JP, Kornbluth M, et al. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat Commun.* 2022;13(1):2453.

<https://doi.org/10.1038/s41467-022-29939-5>.

Batatia I, Kovács DP, Simm G, Ortner C, Csányi G. MACE: Higher order equivariant message passing neural networks for fast and accurate force fields. *Adv Neural Inf Process Syst.* 2022;35:11423-36.

Huo H, Rupp M. Unified representation of molecules and crystals for machine learning. *Mach Learn Sci Technol.* 2022;3(4):045017.

Atz K, Grisoni F, Schneider G. Geometric deep learning on molecular representations. *Nat Mach Intell.* 2021;3(12):1023-32.

<https://doi.org/10.1038/s42256-021-00418-8>.

Zhong X, Gallagher B, Liu S, Kaikhura B, Hiszpanski A, Han TY. Explainable machine learning in materials science. *NPJ Comput Mater.* 2022;8(1):204.

<https://doi.org/10.1038/s41524-022-00884-7>.

Oliva M, Banik S, Josifovski J, Knoll A. Graph neural networks for relational inductive bias in vision-based deep reinforcement learning of robot control. In: 2022 International Joint Conference on Neural Networks (IJCNN); 2022. p. 1-9.

<https://doi.org/10.1109/IJCNN55064.2022.9892101>.

Bishnoi S, Bhattoo R, Ranu S, Krishnan NMA. Enhancing the inductive biases of graph neural ODE for modeling dynamical systems [Preprint]. *arXiv*; 2022. arXiv:2209.10740.

Ringsquandl M, Sellami H, Hildebrandt M, Beyer D, Henselmeyer S, Weber S, et al. Power to the relational inductive bias: Graph neural networks in electrical power grids. In: Proceedings of the 30th ACM International conference on information & knowledge management. New York: Association for Computing Machinery; 2021. p. 1538-47.

<https://doi.org/10.1145/3459637.3482464>.

Angione C, Silverman E, Yaneske E. Using machine learning as a surrogate model for agent-based simulations. *PLoS One*. 2022;17(2).
<https://doi.org/10.1371/journal.pone.0263150>.