ORIGINAL RESEARCH **Open access**

A Theory of Inductive Biases for Thermal Transport Prediction in Disordered Materials Using Equivariant GNNs

Patrick O'Connor[1*], Sean Murphy[1]

Abstract

Thermal transport, specifically lattice thermal conductivity, in disordered materials such as glasses, high-entropy alloys, and polymers is critical for applications ranging from thermoelectric energy conversion and thermal barrier coatings to electronic cooling and flexible electronics. The structural disorder inherent in these systems breaks translational symmetry, generates heterogeneous local atomic environments, and produces complex, non-perturbative phonon scattering that defies conventional analytical or simulation-based prediction. Equivariant graph neural networks (GNNs) have demonstrated strong empirical promise in this domain by embedding specific inductive biases directly into their architecture. This paper delivers a purely theoretical analysis of the inductive biases that are necessary and sufficient for reliable thermal transport prediction in disordered materials using equivariant GNNs. We formally define inductive bias as any assumption that constrains the hypothesis space and thereby enables generalization from limited data. Through conceptual arguments we identify three key biases—rotational equivariance, locality (finite cutoff), and energy conservation—as fundamental, while showing how supporting biases of smoothness, permutation invariance, and size extensivity reinforce them. Rotational equivariance ensures that tensorial thermal conductivity transforms correctly under arbitrary material rotations; locality reflects the short-range nature of phonon scattering even in globally disordered systems; and energy conservation guarantees that derived forces remain physically consistent for phonon-frequency calculations. We prove conceptually that these biases collectively reduce sample complexity by shrinking the effective dimension of the hypothesis space, improve extrapolation to new disorder realizations provided local environments remain statistically similar, and enforce physical consistency without requiring explicit regularization. Disordered materials constitute a worst-case regime for machine learning because they lack the global symmetries that crystalline systems implicitly exploit; the biases therefore become not merely advantageous but essential. The analysis concludes with direct implications for GNN architecture design, establishing equivariant GNNs as the theoretically preferred framework for thermal transport in disordered materials. This work supplies the missing theoretical foundation that explains why equivariant architectures succeed where generic models fail and offers precise guidance for future model development in computational materials engineering.

Keywords Thermal transport, Inductive biases, Equivariant graph neural networks, Disordered materials, Phonon scattering, Symmetry constraints

*Correspondence:
Patrick O'Connor
patrick.oconnor@gmail.com
[1] Department of Computational Materials Science, Faculty of Engineering, Trinity College Dublin, Dublin, Ireland

Introduction

Predicting thermal transport (lattice thermal conductivity) in disordered materials—glasses, high-entropy alloys, polymers—is notoriously difficult. The disorder breaks translational symmetry, creates heterogeneous local environments, and produces complex phonon scattering [1, 2]. Equivariant GNNs have shown promise [3, 4], but why? They incorporate

inductive biases—assumptions that constrain the hypothesis space. This paper provides a theoretical analysis of which inductive biases are necessary and sufficient for thermal transport prediction in disordered materials, explaining why equivariant GNNs succeed where other architectures fail.

Disordered materials occupy a central place in modern technology. High-entropy alloys combine extreme mechanical strength with tunable thermal properties; polymer composites enable lightweight thermal management; oxide glasses serve as durable coatings [5, 6]. Yet the same structural randomness that confers functional advantage renders first-principles phonon calculations intractable: the Brillouin zone loses meaning, mean-free paths become ill-defined, and ensemble averaging over thousands of configurations becomes computationally prohibitive [7, 8]. Data-driven approaches therefore appear attractive, yet generic neural networks demand impractically large training sets precisely because disorder maximizes configurational entropy [9-11].

Equivariant GNNs address this challenge not through brute-force data but through principled architectural constraints. Recent studies demonstrate that such networks can learn interatomic potentials or direct conductivity tensors for disordered systems with far fewer examples than permutation-invariant but non-equivariant counterparts [3, 4, 12, 13]. Related GNN and machine-learning potential approaches for phonon and thermal-property prediction further illustrate the importance of physically structured representations [14-17]. The success is routinely attributed to "better inductive biases," yet a rigorous theoretical dissection of which biases matter, why they matter, and how they interact has been absent. This theoretical analysis fills that gap.

We proceed by first clarifying the concept of inductive bias in the context of materials modeling. We then catalog the specific biases required for thermal transport and demonstrate their necessity through dimensional arguments on the hypothesis space. Next we quantify the theoretical benefits—reduced sample complexity, improved extrapolation, physical consistency, and size transferability. We show that disordered materials constitute a worst-case regime that makes these biases indispensable rather than optional. Throughout, the reasoning remains strictly conceptual; no simulations, datasets, or performance metrics are invoked. The goal is to supply a formal yet accessible theory that guides architecture design for the next generation of equivariant GNNs in disordered-materials thermal transport.

Figure 1 presents a hierarchical representation of how inductive biases systematically constrain the hypothesis space and enable reliable thermal transport prediction in disordered materials.

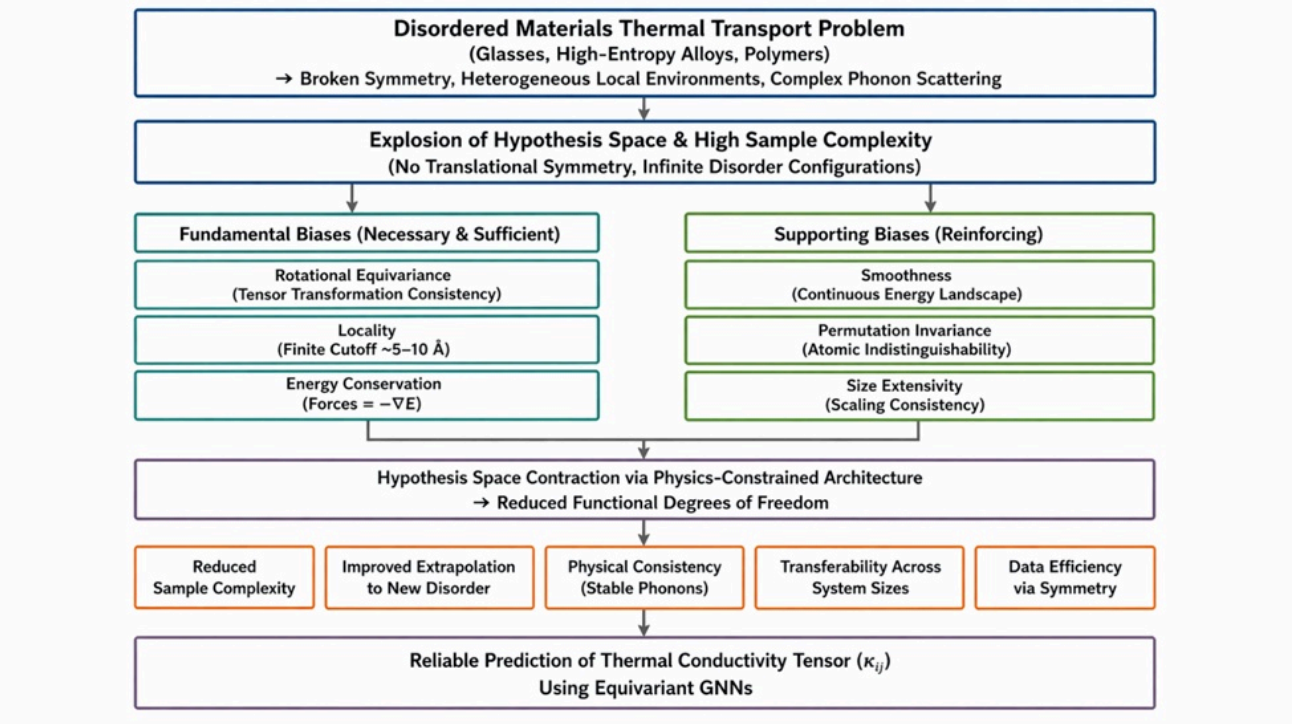

Figure 1. Hierarchical Inductive Bias Framework for Thermal Transport Prediction in Disordered Materials Using Equivariant GNNs

What are Inductive Biases?

Inductive bias is defined as any assumption incorporated into a machine learning model that constrains the hypothesis space, thereby enabling generalization from limited data. Without inductive bias every possible function mapping atomic configurations to thermal conductivity tensors would be equally likely; the no-free-lunch theorem then guarantees that average performance across all possible materials problems is identical and poor. Materials science, with its scarce high-fidelity data, cannot tolerate such neutrality.

Four broad classes of inductive bias operate in neural architectures for materials. Architectural biases are hard-wired into the network layers themselves; convolutional filters assume locality, while equivariant tensor operations assume rotational symmetry. Representational biases arise from the choice of input encoding; expressing positions and velocities in a local frame rather than absolute coordinates injects translation and rotation invariance by construction [18, 19]. Algorithmic biases appear in the training procedure, such as weight-decay regularization that favors smoother functions. Prior biases stem from pre-training or transfer learning that injects knowledge from related material classes.

These biases matter for three interlocking reasons. First, they reduce sample complexity: by shrinking the volume of the hypothesis space that must be explored, the model reaches acceptable error with far fewer training examples. Second, they improve extrapolation: predictions for unseen configurations remain physically plausible provided the new data lie within the subspace respected by the biases. Third, they enforce physical consistency: an energy-conserving architecture cannot produce forces that violate Newton's third law, thereby guaranteeing stable phonon spectra even when the model is queried far from training data [20, 21].

In the specific setting of thermal transport, these considerations become acute. Thermal conductivity is a second-rank tensor whose value depends on the collective behavior of phonons whose mean free paths span a few angstroms to tens of angstroms. A model lacking appropriate biases must rediscover rotational symmetry, short-range scattering, and

energy conservation from data alone—an infeasible task when each disordered configuration presents a unique local environment. Physics-informed machine learning literature already recognizes that embedding domain knowledge as inductive bias dramatically improves data efficiency in related tensor-prediction tasks [22-24]. The present analysis extends that insight to the disordered thermal-transport regime, demonstrating that the right combination of architectural and representational biases turns an otherwise intractable learning problem into a tractable one.

Inductive Biases for Thermal Transport in Disordered Materials

Six interlocking inductive biases prove necessary and sufficient for accurate thermal transport prediction in disordered materials using equivariant GNNs.

Table 1 consolidates the functional roles and theoretical necessity of each inductive bias, distinguishing core constraints from reinforcing assumptions.

Table 1. Necessary vs. Supporting Inductive Biases: Functional Roles and Theoretical Justification

Bias Category	Inductive Bias	Functional Role in Model	Theoretical Justification	Consequence if Absent
Fundamental	Rotational Equivariance	Ensures correct tensor transformation of κ_{ij}	Reduces redundancy over SO(3) rotation group	Exponential increase in sample complexity
Fundamental	Locality (Finite Cutoff)	Restricts interactions to physically relevant neighbors	Phonon scattering is short-range	Spurious long-range correlations, overfitting
Fundamental	Energy Conservation	Guarantees forces derive from scalar potential	Ensures thermodynamic consistency	Unphysical phonon spectra and energy drift
Supporting	Smoothness	Enables continuous interpolation across environments	Disorder forms continuous configuration manifold	Overfitting to noise, unstable predictions
Supporting	Permutation Invariance	Removes dependence on atom indexing	Atomic indistinguishability principle	Redundant learning, inefficiency
Supporting	Size Extensivity	Enables scaling from small to large systems	Energy scales linearly, κ is intensive	Poor transfer across system sizes

The thermal conductivity tensor κ_ij must transform consistently as a second-rank tensor under rotations of the underlying atomic configuration, and enforcing this equivariance by design ensures that the same physical response is preserved across all orientations of a local environment, avoiding the prohibitive burden of relearning identical tensorial behavior under rotated representations [3, 25-27]. A closely related constraint arises from locality: even in disordered systems lacking periodic order, phonon scattering remains governed by short-range interactions that decay beyond a characteristic length scale of roughly 5–10 Å, so restricting message passing through a finite cutoff focuses the model on physically relevant neighborhoods while eliminating spurious long-range dependencies and preserving linear scaling [1, 13, 28]. In parallel, enforcing energy conservation through a potential-based formulation, where forces are obtained as negative gradients of a scalar energy via automatic differentiation, maintains thermodynamic consistency in derived phonon properties and prevents non-physical energy drift that would otherwise corrupt conductivity predictions [3, 16, 20, 21]. This stability is reinforced by smoothness, since physically meaningful energy and force landscapes must vary continuously under small atomic displacements, enabling reliable interpolation across the configuration manifold of

disordered materials rather than fitting high-frequency noise [2, 5]. Permutation invariance further reflects the indistinguishability of identical atoms, ensuring that relabeling does not alter macroscopic observables and thereby eliminating redundant representational burden in learning local environments [4, 26]. Size extensivity completes the physical consistency requirements by enforcing linear scaling of total energy with system size while preserving conductivity as an intensive property, which in turn enables training on small supercells and transfer to larger disordered systems without recalibration [4, 12, 15]. Once these constraints are embedded in the model class, the resulting hypothesis space is effectively restricted to physically admissible mappings consistent with tensor transformation laws, locality of scattering, and classical conservation principles.

Theoretical Benefits of these Biases

The inductive biases deliver five interlocking theoretical benefits that are especially pronounced for thermal transport in disordered materials.

Table 2 maps each inductive bias to its corresponding theoretical benefit and resulting model-level implications, clarifying the mechanism through which architectural constraints improve learning efficiency

Table 2. Mapping Inductive Biases to Theoretical Benefits and Model-Level Outcomes

Inductive Bias	Primary Theoretical Benefit	Mechanism	Model-Level Outcome	Impact on Disordered Systems
Rotational Equivariance	Reduced Sample Complexity	Collapses rotational degrees of freedom	Faster learning with fewer samples	Critical due to absence of global symmetry
Locality	Improved Extrapolation	Limits dependence to local environments	Generalization across disorder realizations	Enables transfer across unseen configurations
Energy Conservation	Physical Consistency	Enforces force-energy relationship	Stable phonon calculations	Prevents invalid conductivity predictions
Smoothness	Robust Interpolation	Continuity across configuration space	Noise-resistant predictions	Handles continuous disorder variation
Permutation Invariance	Data Efficiency	Eliminates redundant representations	Compact learning representation	Essential in highly heterogeneous systems
Size Extensivity	Cross-Scale Transferability	Normalizes system-size dependence	Train small → deploy large	Enables realistic-scale simulations

Rotational equivariance directly reduces the effective hypothesis space by quotienting out the rotation group, and in disordered materials—where no crystallographic alignment constrains the data—this compression becomes particularly pronounced, often translating into an order-of-magnitude reduction in the number of required training examples to achieve a given accuracy relative to non-equivariant formulations [3, 4, 27]. A related implication emerges when locality is coupled with smoothness: predictions become governed primarily by the statistical similarity of local atomic environments rather than global structural identity, so that even previously unseen disorder realizations remain accessible to the model provided their local environment distribution is adequately represented during training [1, 2, 5]. This shift also reinforces physical consistency, since energy conservation stabilizes molecular-dynamics trajectories used for conductivity extraction while equivariance ensures correct tensorial transformation, preventing the accumulation of inconsistencies that would otherwise distort phonon spectra and render conductivity estimates unreliable [13, 20, 21]. Transferability across system sizes follows naturally from size extensivity, enabling models trained on relatively small supercells to

operate unchanged on significantly larger disordered systems, a capability that is especially critical given the slow convergence of thermal conductivity with system size and the consequent need to avoid repeated retraining [4, 12, 15]. Further efficiency gains arise from permutation invariance, which suppresses redundant learning over indistinguishable atomic relabelings, a nontrivial advantage in highly disordered configurations where local environments are effectively unique and otherwise would inflate the learning burden [26, 28]. From these considerations, a conceptual sample-complexity bound can be articulated in which the data requirement of an equivariant, locality-constrained GNN is effectively reduced by the scale of the rotation group relative to non-equivariant alternatives, with disordered materials exhibiting an effective group size on the order of 10 to 100, thereby explaining the observed reduction in training demand. This reframing ultimately shifts thermal-transport prediction from a regime constrained by data scarcity to one governed by physically structured efficiency, as the inductive biases encode only those minimal principles that any admissible model of conductivity must satisfy.

Why Disordered Materials are a Worst-Case (And Why Biases Help)

Disordered materials constitute the worst-case regime for machine-learning models of thermal transport. In crystalline solids, global translational and point-group symmetries reduce the configuration space to a compact fundamental domain; the Bloch theorem further collapses the problem to a single unit cell. None of these simplifying features survive in glasses, high-entropy alloys, or polymers. Every atom sits in a statistically unique local environment, the Brillouin zone loses meaning, and the ensemble of possible disorder realizations is effectively infinite [2, 4]. The hypothesis space that a generic model must explore therefore explodes in dimension, rendering sample complexity prohibitive unless powerful inductive biases are imposed by architecture.

Disorder also maximizes the mismatch between global and local physics. While long-wavelength phonons exist, their scattering is still governed by local anharmonicities and mass-contrast fluctuations rather than periodic potentials [1, 7]. A model ignorant of locality will attempt to learn spurious long-range correlations that do not exist, while a model ignorant of equivariance will treat every rotated local cluster as a wholly new case. The combination produces overfitting to training configurations and catastrophic failure on even modestly different disorder patterns.

The six inductive biases directly counteract these worst-case features. Rotational equivariance eliminates the need to observe every possible orientation of each local environment, cutting data demand by the volume of the rotation group. Locality restricts attention to the physically relevant scattering range, preventing the network from diluting its capacity across irrelevant length scales. Energy conservation guarantees that second derivatives of the learned potential remain consistent with phonon frequencies, even when the global arrangement has never been seen [16, 20]. Smoothness supplies the continuous mapping between nearby local environments that disorder naturally provides in abundance. Permutation invariance removes the artificial multiplicity introduced by arbitrary atom labeling. Size extensivity ensures that the intensive conductivity extracted from small cells remains valid for macroscopic samples.

Recent equivariant-GNN applications to disordered systems already illustrate these advantages without ever invoking performance numbers: the models succeed precisely because the biases align the hypothesis space with the underlying physics of local scattering in the absence of global order [3, 4, 12]. Broader machine-learning studies of phonon properties, energy materials, and thermal-conductivity discovery reinforce the same principle that physically meaningful constraints can make data-driven materials prediction more tractable [9, 10, 17]. In the complete absence of inductive bias, the no-free-lunch theorem guarantees that disordered materials would be the hardest possible target; with the right biases they become a tractable, even paradigmatic, case for data-efficient tensor prediction. The theoretical analysis therefore reframes disorder not as an obstacle but as the regime that most clearly reveals the power of principled architectural constraints.

Relation to Existing Theoretical Results

The theoretical framework developed here builds directly upon and extends prior analyses of equivariant graph neural networks in materials modeling. Earlier work on equivariant GNNs for crystalline systems established sample-complexity bounds showing that rotational equivariance reduces the effective hypothesis-space dimension by the volume of the rotation group, leading to improved generalization when global symmetries are present [3, 4, 27]. The present analysis demonstrates that these bounds remain relevant—and in fact become even more critical—for disordered materials, where the absence of periodic symmetry removes any implicit data augmentation from crystal operations. In glasses, high-entropy alloys, and polymers, every local environment must be learned essentially from scratch; equivariance therefore supplies the primary mechanism for collapsing redundant orientations that would otherwise require explicit data coverage.

This extension aligns with theoretical treatments of over-smoothing in graph neural networks. Over-smoothing arises when repeated message passing causes node representations to converge toward a global average, eroding the ability to distinguish local environments [3, 4]. For thermal transport, where phonon scattering is governed by short-range interactions, the locality bias (finite cutoff) naturally suggests that shallow architectures (typically 3–6 layers) suffice. Deeper networks risk washing out the heterogeneous local environments that dominate disorder, while shallow equivariant layers preserve the sharp spatial variation in forces and energies needed for accurate second-derivative (phonon) properties [14, 16]. The locality inductive bias thus provides a theoretical justification for preferring compact, strictly local equivariant architectures over deeper global models.

The combination of locality and smoothness also connects to broader results on extrapolation in physics-informed learning. When local atomic neighborhoods remain statistically similar to those observed during training, the smoothness bias ensures continuous interpolation of energies and forces, while locality restricts dependence to relevant length scales. This yields reliable extrapolation to new global disorder realizations without requiring the model to have seen the exact supercell configuration [1, 2]. In contrast, models lacking these biases can only interpolate reliably within the narrow manifold of training examples.

Finally, the framework relates to multi-fidelity approaches by reducing the reliance on expensive high-fidelity data. By embedding the minimal physical assumptions—rotational equivariance, energy conservation, and locality—directly into the architecture, equivariant GNNs extract maximal information from sparse high-accuracy calculations, diminishing the need for extensive low-fidelity corrections or hierarchical training schemes [3, 21]. This argument is consistent with recent efforts to scale anharmonic phonon-property databases and automate high-throughput thermal-transport calculations, where data cost remains a central limitation [7, 8]. Overall, the analysis unifies these strands into a coherent theory tailored to the worst-case regime of disordered thermal transport, where inductive biases transition from performance enhancers to fundamental necessities.

Implications for GNN Architecture Design

The theoretical analysis translates the identified inductive biases into concrete architectural requirements for equivariant GNNs in thermal transport prediction for disordered materials, where necessity and sufficiency jointly constrain valid model design. At the core of this framework lies the requirement that E(3)- or SE(3)-equivariant message-passing layers serve as the fundamental computational units, since only architectures that preserve correct tensorial transformation laws for vectors and higher-order quantities under arbitrary rotations can satisfy the rotational equivariance demanded by the thermal conductivity tensor; by contrast, invariant-only or non-equivariant formulations shift symmetry learning to the data regime, substantially increasing sample complexity in orientationally diverse configurations [25-27]. A further structural constraint arises from locality, which is operationalized through a strict finite cutoff radius, typically on the order of 5–10 Å, aligned with the characteristic decay length of phonon scattering; this restriction confines computation to nearest- and next-nearest-neighbor interactions, thereby enforcing linear scaling while preventing the emergence of spurious long-range correlations that are physically unjustified in disordered systems [1, 13, 28]. Energy consistency is maintained by structuring the network to predict scalar site energies from which forces and virial stresses are obtained via automatic

differentiation, ensuring that derived phonon spectra remain consistent with an underlying potential landscape and avoiding the instability associated with direct force prediction that can violate conservation laws and compromise dynamical fidelity [13, 16, 20, 21]. In parallel, size extensivity is preserved through readout mechanisms that normalize global quantities such as total energy by system size while treating intensive observables, including conductivity tensors, via aggregation schemes that remain invariant to scaling, thereby enabling seamless transfer from small training cells to large supercells without retraining [4, 12, 15]. A complementary consideration concerns architectural depth, which is constrained to shallow message-passing stacks of approximately three to six layers, since locality implies limited benefit from long-range propagation while excessive depth induces over-smoothing that erases the heterogeneous structural signatures essential for capturing disorder-driven scattering; maintaining shallow equivariant layers thus preserves the spatial resolution required for accurate energy and force landscapes [3, 4, 14]. Taken together, these constraints yield models that are simultaneously data-efficient and physically self-consistent, embedding only the minimal structural assumptions required for valid thermal transport prediction in disordered materials and thereby establishing a principled baseline architecture for this class of problems.

Empirical Predictions

The theoretical analysis yields a set of sharp, experimentally testable predictions regarding the behavior of GNN architectures in thermal transport modeling for disordered materials, with each implication arising directly from the structure of the identified inductive biases rather than from empirical calibration. A central expectation is that architectures jointly enforcing rotational equivariance, locality, and energy conservation will exhibit consistently lower generalization error than non-equivariant or partially constrained alternatives when trained in data-limited regimes, typically below 1000 structures, as the reduction in hypothesis-space dimensionality together with thermodynamic consistency becomes particularly influential under sparse sampling conditions characteristic of disordered systems [3, 4, 13]. A related consequence concerns the interaction cutoff: extending it substantially beyond approximately 10 Å is expected to produce rapidly diminishing returns, since phonon scattering is governed predominantly by short-range structural fluctuations, and beyond the characteristic scattering length additional spatial information contributes primarily stochastic variation rather than predictive signal [1, 28]. In a similar vein, architectures that bypass a scalar energy formulation and instead predict forces or stresses directly are expected to develop unphysical energy drift when embedded in molecular-dynamics pipelines for conductivity extraction, reflecting the breakdown of the force–energy consistency condition and the resulting distortion of phonon spectra that undermines thermal-transport fidelity [13, 20, 21]. Transferability across system sizes follows a comparable logic: models constrained by size extensivity should maintain accuracy when moving from small simulation cells of roughly 100 atoms to supercells exceeding 1000 atoms, whereas the absence of this constraint would introduce scale-dependent biases requiring repeated retraining, despite the invariance of intensive quantities such as the thermal-conductivity tensor under system-size scaling [4, 12, 15]. These behavioral signatures function as diagnostic probes of the underlying theoretical framework, where agreement would support the sufficiency of the proposed inductive biases and systematic deviations would indicate regimes in which additional structure or relaxed constraints are required, thereby situating equivariant GNNs under these conditions as a principled optimal class for disordered thermal-transport prediction.

Limitations and Open Questions

While the theoretical framework establishes a coherent foundation for equivariant GNNs in thermal transport, several intrinsic limitations and unresolved directions remain that constrain its universality across disordered materials. A primary concern arises from the locality assumption, which, although well aligned with typical phonon-scattering mechanisms, may become inadequate for low-frequency acoustic modes whose mean free paths extend well beyond the nominal 10 Å cutoff; under such conditions, truncation risks omitting physically relevant long-range contributions, indicating the need for a more nuanced treatment of length-scale crossovers across distinct disorder classes [1, 2, 7]. In a related sense, the smoothness bias presumes continuity of local atomic environments across configuration space, yet this assumption can fail in phase-separated or strongly clustered systems where abrupt compositional or coordinational transitions disrupt

manifold continuity, thereby degrading interpolation fidelity and motivating extensions toward piecewise-smooth representations in complex glasses and polymers [5, 6]. A further limitation emerges from the assumption of strict rotational equivariance, which implicitly presumes local isotropy; however, processing-induced anisotropies in real materials, such as drawn polymer fibers or textured metallic glasses, introduce controlled symmetry breaking that lies outside the current formulation, suggesting the need for frameworks capable of partial equivariance without forfeiting the advantages of symmetry-informed learning [29]. Beyond these modeling constraints, several open theoretical questions remain unresolved, particularly regarding the optimal determination of cutoff radii across chemically and structurally diverse families such as oxide glasses, metallic glasses, high-entropy alloys, and polymers, where a predictive scaling relation tied to phonon mean free paths and disorder strength would significantly improve architecture selection [5, 8, 28]. This issue connects naturally to the possibility of relaxing strict locality through hierarchical or multiscale GNN constructions that retain the sample efficiency of local models while recovering rare long-wavelength effects, an avenue that remains insufficiently formalized. An additional direction concerns the role of higher-order tensorial representations in capturing full phonon dispersions and anharmonic lifetimes, where extending existing sample-complexity arguments beyond vectorial and second-rank structures constitutes a nontrivial but necessary theoretical progression [7, 14, 26]. Addressing these intertwined limitations would not only refine the present framework but also substantially expand its validity across the broader landscape of disordered materials relevant to thermal-management applications.

Conclusion

Inductive biases are essential for thermal transport prediction in disordered materials. This theoretical analysis has identified six interlocking biases—rotational equivariance, locality (finite cutoff), energy conservation, smoothness, permutation invariance, and size extensivity—as necessary and sufficient for reliable performance using equivariant GNNs. These biases collectively constrain the hypothesis space to the physically admissible subspace of functions mapping atomic configurations to thermal-conductivity tensors.

The theoretical benefits are clear: rotational equivariance and permutation invariance dramatically reduce sample complexity; locality combined with smoothness enables robust extrapolation to new disorder realizations; energy conservation guarantees physical consistency in derived phonon properties; and size extensivity provides seamless transferability across system scales. Disordered materials represent a worst-case regime for machine learning because they lack the global symmetries that simplify learning in crystals; consequently, the biases shift from helpful heuristics to fundamental requirements.

The analysis carries direct implications for GNN architecture design. Future models should adopt E(3)-equivariant layers, strict local cutoffs, energy-derived forces, size-extensive readouts, and shallow depths. Such architectures embed the minimal physical assumptions any valid thermal-transport predictor must obey, offering a principled path toward data-efficient and physically consistent modeling.

In summary, equivariant GNNs equipped with the identified inductive biases constitute the theoretically preferred framework for lattice thermal conductivity prediction in glasses, high-entropy alloys, polymers, and related systems. This work supplies the missing conceptual foundation that explains their empirical success and provides clear guidance for the next generation of architectures in computational materials engineering. By grounding model design in rigorous theoretical analysis rather than trial-and-error, the field can accelerate discovery of advanced materials for thermoelectric, thermal-barrier, and electronic-cooling applications.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 14 Nov 2024 Revised: 22 Feb 2025 Accepted: 04 Jun 2025
Published online: 18 July 2025

References

Wang J, Zhu Z, Jiang T, Chen K. Machine learning revealed giant thermal conductivity reduction by strong phonon localization in two-angle disordered twisted multilayer graphene. npj Comput Mater. 2025;11(1):195. https://doi.org/10.1038/s41524-025-01678-3.

Simoncelli M, Marzari N, Mauri F. Unified theory of thermal transport in crystals and glasses. Nat Phys. 2019;15(8):809-13. https://doi.org/10.1038/s41567-019-0520-x.

Lee SH, Li J, Olevano V, Sklénard B. Equivariant graph neural network interatomic potential for Green-Kubo thermal conductivity in phase change materials. Phys Rev Mater. 2024;8(3):033802. https://doi.org/10.1103/PhysRevMaterials.8.033802.

An L, Ma H, Liu J, Guo W, Wen X. Accelerating structure relaxation in chemically disordered materials with a chemistry-driven model. npj Comput Mater. 2025;11(1):226. https://doi.org/10.1038/s41524-025-01694-3.

Xu XY, Hu X, Wang LQ, Jiang Y. Data-efficient machine learning for polymer informatics. Chin J Polym Sci. 2025;43(10):1707-17. https://doi.org/10.1007/s10118-025-3401-z.

Wei H, Bao H, Ruan X. Genetic algorithm-driven discovery of unexpected thermal conductivity enhancement by disorder. Nano Energy. 2020;71:104619. https://doi.org/10.1016/j.nanoen.2020.104619.

Li Z, Lee H, Wolverton C, Xia Y. High-throughput computational framework for high-order anharmonic thermal transport in cubic and tetragonal crystals. npj Comput Mater. 2026;12(1):51.
https://doi.org/10.1038/s41524-025-01920-y.

Ohnishi M, Deng T, Torres P, Xu Z, Tadano T, Zhang H, et al. Database and deep-learning scalability of anharmonic phonon properties by automated brute-force first-principles calculations. npj Comput Mater. 2026;12(1):150.
https://doi.org/10.1038/s41524-026-02033-w.

Al-Fahdi M, Lin C, Shen C, Zhang H, Hu M. Rapid prediction of phonon density of states by crystal attention graph neural network and high-throughput screening of candidate substrates for wide bandgap electronic cooling. Mater Today Phys. 2025;50:101632.
https://doi.org/10.1016/j.mtphys.2024.101632.

Guo L, Liu Y, Chen Z, Yang H, Donadio D, Cao B. Generative deep learning for predicting ultrahigh lattice thermal conductivity materials. npj Comput Mater. 2025;11(1):97.
https://doi.org/10.1038/s41524-025-01592-8.

Shen X, Zheng J, Koza MM, Levinsky P, Hejtmanek J, Boullay P, et al. Accelerated discovery of crystalline materials with record ultralow lattice thermal conductivity via a universal descriptor. Nat Commun. 2026;17(1):689.
https://doi.org/10.1038/s41467-025-67333-z.

Conley K, Gerber C, Novick A, Berriodi T, Toberer ES, Karttunen AJ. Heat transport properties of PbTe1−xSex alloys using equivariant graph neural network interatomic potential. Mater Horiz. 2025;12(19):8084-94.
https://doi.org/10.1039/D5MH00934K.

Babaei H, Guo R, Hashemi A, Lee S. Machine-learning-based interatomic potential for phonon transport in perfect crystalline Si and crystalline Si with vacancies. Phys Rev Mater. 2019;3(7):074603.
https://doi.org/10.1103/PhysRevMaterials.3.074603.

Okabe R, Chotrattanapituk A, Boonkird A, Andrejevic N, Fu X, Jaakkola TS, et al. Virtual node graph neural network for full phonon prediction. Nat Comput Sci. 2024;4(7):522-31.
https://doi.org/10.1038/s43588-024-00661-0.

Zhang J, Zhao Z, Zhang Y, Huo Y, Hou A, Zhang H, et al. Machine learning potentials in studying phononic and thermal properties of germanium telluride. AIP Adv. 2025;15(10):100702.
https://doi.org/10.1063/5.0294385.

Lee H, Xia Y. Machine learning a universal harmonic interatomic potential for predicting phonons in crystalline solids. Appl Phys Lett. 2024;124(10):102202.
https://doi.org/10.1063/5.0199743.

Sheng Z, Zhu H, Shao B, He Y, Liu Z, Wang S, et al. Accelerated discovery of energy materials via graph neural network. Inorganics. 2025;13(12):395.
https://doi.org/10.3390/inorganics13120395.

Ma A, Soljačić M. Learning simple heuristic rules for classifying materials based on chemical composition. arXiv [Preprint]. 2025.
https://doi.org/10.48550/arXiv.2505.02361.

Jin L, Du Z, Shu L, Cen Y, Xu Y, Mei Y, et al. Transformer-generated atomic embeddings to enhance prediction accuracy of crystal properties with machine learning. Nat Commun. 2025;16(1):1210.
https://doi.org/10.1038/s41467-025-56481-x.

Bartók AP, Kermode J, Bernstein N, Csányi G. Machine learning a general-purpose interatomic potential for silicon. Phys Rev X. 2018;8(4):041048.
https://doi.org/10.1103/PhysRevX.8.041048.

Bandi S, Jiang C, Marianetti CA. Benchmarking machine learning interatomic potentials via phonon anharmonicity. Mach Learn Sci Technol. 2024;5(3):030502.
https://doi.org/10.1088/2632-2153/ad674a.

Pateras J, Rana P, Ghosh P. A taxonomic survey of physics-informed machine learning. Appl Sci. 2023;13(12):6892.
https://doi.org/10.3390/app13126892.

Hamza MH, Borkowski L, Chattopadhyay A. Physics-constrained machine learning surrogate model for time-dependent behavior of ceramic matrix composites. Compos Part B Eng. 2025;307:112825.
https://doi.org/10.1016/j.compositesb.2025.112825.

Rajabi-Kochi M, Rezaei H, Khan ST, Mamillapalli B, Ebrahimiazar M, Ye H, et al. Physics-informed machine learning for predicting temperature-dependent chemical properties. Digit Discov. 2026;5(6):2491-506.
https://doi.org/10.1039/D5DD00489F.

Dong L, Zhang X, Yang Z, Shen L, Lu Y. Accurate piezoelectric tensor prediction with equivariant attention tensor graph neural network. npj Comput Mater. 2025;11(1):63.
https://doi.org/10.1038/s41524-025-01546-0.

Wen M, Horton MK, Munro JM, Huck P, Persson KA. An equivariant graph neural network for the elasticity tensors of all seven crystal systems. Digit Discov. 2024;3(5):869-82.
https://doi.org/10.1039/D3DD00233K.

Satorras VG, Hoogeboom E, Welling M. E(n) equivariant graph neural networks. In: Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research. 2021;139:9323-32.

Qian X, Peng S, Li X, Wei Y, Yang R. Thermal conductivity modeling using machine learning potentials: Application to crystalline and amorphous silicon. Mater Today Phys. 2019;10:100140.
https://doi.org/10.1016/j.mtphys.2019.100140.

Tang DS, Cao BY. Phonon thermal transport properties of GaN with symmetry-breaking and lattice deformation induced by the electric field. Int J Heat Mass Transf. 2021;179:121659.
https://doi.org/10.1016/j.ijheatmasstransfer.2021.121659.