ORIGINAL RESEARCH **Open access**

Why Invariance Becomes a Liability: Failure Modes of Rotationally Invariant GNNs for Anisotropic Elastic Constants

Michael Turner[1], Sophia Nguyen[2*], David Clark[1], Emma Wilson[3]

Abstract

Rotational invariance, a core design principle in graph neural networks (GNNs) for materials property prediction, has enabled remarkable accuracy for scalar quantities such as formation energy and band gap yet becomes a fundamental liability when the target property is a tensor of rank two or higher. Anisotropic elastic constants, which describe the directional stiffness of crystals and enter engineering design as a 6×6 Voigt matrix with up to 21 independent components in triclinic systems, must transform covariantly under rotation; any architecture that collapses all orientational information into a single invariant scalar cannot, in principle, recover the required tensorial structure. This failure-mode analysis articulates four distinct mechanisms by which invariance induces systematic error—diagonal indistinguishability, off-diagonal blindness, shear underspecification, and symmetry-induced degeneracy—and demonstrates how each arises directly from the interplay of global pooling, distance-only edge features, and symmetry averaging that characterize rotationally invariant models such as CGCNN, MEGNet, and SchNet. Detection principles based on equality tests, anisotropy-ratio checks, and differential error audits provide practitioners with practical diagnostics, while mitigation principles centered on equivariant architectures, directional encoding, and stress-tensor multi-task learning chart a clear path toward reliable tensorial prediction. The implications extend beyond elasticity to any directional property in low-symmetry crystals and call for a principled redesign of materials GNNs that treats invariance as a deliberate choice rather than a default.

Keywords Failure mode analysis, Rotational invariance, Equivariant GNNs, Anisotropic elastic constants, Information loss, Tensorial properties

*Correspondence:
Sophia Nguyen
sophia.nguyen@gmail.com
[1] Department of Computational Materials Engineering, Faculty of Engineering, School of Materials Science and Technology, University of Glasgow, Glasgow, United Kingdom
[2] Department of Data-Driven Materials Systems, Faculty of Engineering, National University of Singapore, Singapore, Singapore
[3] Department of Intelligent Materials Analytics, Faculty of Science and Engineering, University of Sydney, Sydney, Australia

Introduction

Rotationally invariant graph neural networks have become the de facto standard for computational materials discovery because they elegantly encode the physical fact that many key properties—formation energy, band gap, thermodynamic stability—are scalars that remain unchanged under arbitrary rotation of the crystal lattice. Xie and Grossman introduced the crystal graph convolutional neural network (CGCNN) [1], building on the broader message-passing logic of graph neural networks [2], which demonstrated that message-passing on graphs whose edges are defined solely by interatomic distances could achieve state-of-the-art scalar-property prediction across thousands of inorganic compounds. Subsequent architectures, including the graph networks of Chen *et al.* [3], the continuous-filter convolutional network of Schütt *et al.* [4], and broader convolutional architectures [5, 6], retained the same invariance principle while refining the representation of local atomic environments. These models have powered large-scale screening campaigns and

accelerated the identification of new functional materials and quantum-chemistry-related prediction tasks [7] precisely because their outputs are independent of the arbitrary coordinate frame chosen by the crystallographer.

Yet the very invariance that confers robustness for scalars becomes a liability when the target is a tensor property. Elastic constants C_{ijkl} form a rank-4 tensor that must obey the transformation law $C'_{ijkl} = R_{ip}R_{jq}R_{kr}R_{ls}C_{pqrs}$ under any rotation R of the coordinate system. In low-symmetry crystals—triclinic, monoclinic, orthorhombic—the 21 independent components encode directional stiffnesses that differ markedly along different crystallographic axes; C11 is not interchangeable with C33, nor is the shear modulus C44 equivalent to C66. An invariant GNN, by construction, produces a single scalar embedding per crystal that cannot distinguish one axis from another once global pooling has erased relative orientations. Consequently, predictions for anisotropic elastic tensors systematically collapse distinctions that are physically essential for mechanical design, finite-element analysis, and microstructure-property linkages.

This paper performs a failure-mode analysis grounded exclusively in the conceptual architecture of invariant GNNs and the tensorial nature of elasticity. It does not present new benchmarks, datasets, or numerical experiments; rather, it articulates why invariance suffices for scalars but fails for tensors, identifies the precise mechanisms of information loss, and develops a typology of four failure modes that are inevitable once directional information is discarded. The analysis draws on the foundational invariant models [1, 3, 4] and contrasts them with the equivariant frameworks of Thomas *et al.* [8] and Batzner *et al.* [9] to highlight the architectural choices responsible for the observed limitations. By moving systematically from definitional clarity to mechanistic explanation, failure typology, detection, and mitigation, the manuscript supplies the materials community with a diagnostic and prescriptive framework that can prevent the silent propagation of tensorial errors in future GNN-driven discovery campaigns.

The motivation is both practical and foundational, extending to broader functional-material design contexts where structural directionality can influence target performance [10]. Elastic constants govern everything from thermal-expansion mismatch in microelectronics to toughness in structural alloys; inaccurate anisotropy predictions lead to flawed device design and missed opportunities in inverse materials design. Zunger [11], Stach *et al.* [12], and Szymanski *et al.* [13] have emphasized that autonomous materials research requires not only high average accuracy but also physically consistent predictions across the full space of crystal symmetries. Uchino's review of piezoelectric materials [14] further underscores that tensorial properties are the norm rather than the exception in advanced functional materials. Yet current invariant GNN literature rarely flags the tensor-specific failure modes, treating elasticity benchmarks as interchangeable with scalar tasks. This oversight is not merely technical; it reflects a deeper category error in how symmetry is encoded in graph representations. The present analysis corrects that error by demonstrating that rotational invariance, while physically correct for scalars, is physically incorrect for tensors and must therefore be replaced or augmented when the property of interest carries directional information.

Figure 1 visualizes the manuscript's central argument by showing how a mismatch between invariant architectural symmetry and tensorial target symmetry generates mechanisms of information loss, recurring failure modes, diagnostic tests, and symmetry-aware mitigation pathways.

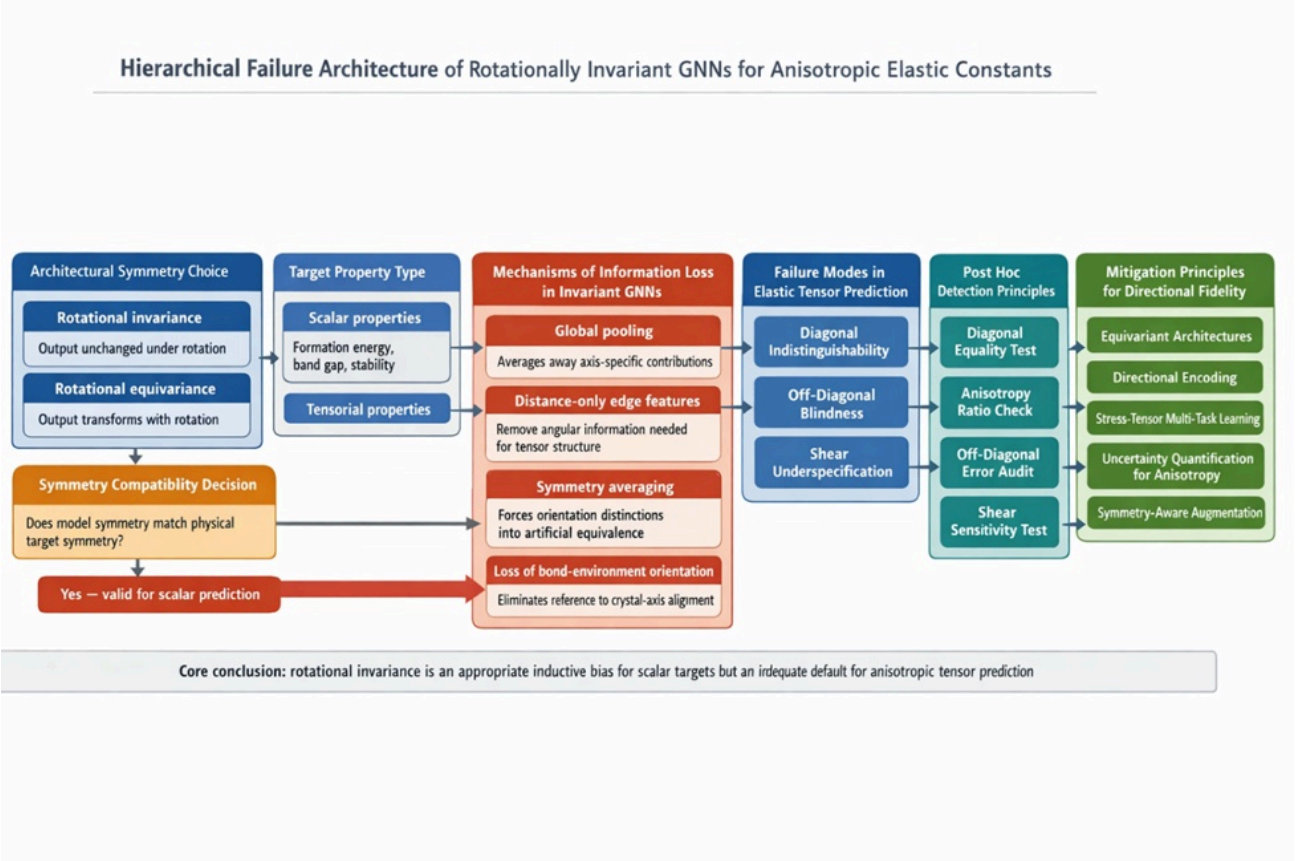

Figure 1. Hierarchical Failure Architecture of Rotationally Invariant GNNs for Anisotropic Elastic Constants

Rotational Invariance in Materials GNNs

A function f is rotationally equivariant if $f(R \cdot x) = R \cdot f(x)$, so that the output tensor transforms exactly as the physical tensor would under the same rotation. Equivariant networks therefore preserve directional information through every layer.

Most materials GNNs achieve invariance through three interlocking mechanisms common across graph neural network designs [15] that together erase all absolute orientational content. First, edge features are constructed exclusively from interatomic distances; no angular or directional descriptors relative to a global crystal axis are included. Second, node and edge embeddings are aggregated via global pooling—typically a sum or mean over all atoms or bonds—producing a single fixed-size vector that is, by construction, permutation- and rotation-invariant. Third, message-passing updates rely on symmetry-averaged local environments, so that bonds that are symmetry-equivalent under the crystal point group are treated identically even when their absolute orientations differ.

Xie and Grossman's CGCNN [1] exemplifies this design: each edge is represented by a radial basis expansion of distance alone, and crystal-level properties emerge from a final pooling layer that collapses all atomic contributions into one invariant descriptor. Chen *et al.* [3] extended the approach with graph networks that still operate on distance-only edges and global mean pooling, achieving universality for scalars precisely because the pooling step discards any residual directional signal. Schütt *et al.* [4] introduced continuous-filter convolutions that remain fully invariant by filtering interactions solely through radial distances and by symmetrizing the filter responses across the entire structure. These choices are not accidental; they reflect broader deep learning architectural trade-offs [16] and guarantee that the network

respects the physical symmetry of scalar properties while simultaneously simplifying the learning problem by reducing the representation space.

The same invariance principle appears in later variants such as the atomistic line graph neural network [17] and GemNet [18], both of which retain distance-only edges and global aggregation. Even when more sophisticated or universal graph convolutional message passing is introduced [19], the final readout remains a rotationally invariant scalar. In contrast, Thomas *et al.* [8] and Jackson *et al.* [20] demonstrated that tensor field networks can maintain equivariance by propagating spherical-harmonic features that transform predictably under rotation, and Batzner *et al.* [9] showed that E(3)-equivariant graph networks dramatically improve force and stress predictions by preserving directional information throughout the architecture. The contrast is instructive: invariant models succeed exactly where directional information is superfluous, while equivariant models succeed precisely where it is indispensable.

Within the invariant paradigm, any information about the orientation of a bond relative to the laboratory or crystal frame is lost at the moment of pooling. The network therefore cannot, even in principle, learn that C11 corresponds to compression along the a-axis while C33 corresponds to the c-axis in a tetragonal crystal. This architectural commitment explains why invariant GNNs achieve low mean-absolute errors on scalar benchmarks yet exhibit systematic biases when applied to elastic tensors. The remainder of this analysis traces the downstream consequences of that commitment.

Why Invariance Suffices for Scalars but Fails for Tensors

Scalar properties such as formation energy or band gap are, by their physical nature, rotationally invariant. The total energy of a crystal does not change if the entire lattice is rotated in space; the band gap likewise remains unchanged. Consequently, an invariant GNN that produces the same output for any orientation of the input structure is not only mathematically convenient but physically correct. No directional information is required because none exists in the target. The network can therefore safely discard all orientational cues during pooling and still recover the exact scalar value. This alignment between architectural symmetry and physical symmetry explains the success of CGCNN [1], MEGNet [3], and SchNet [4] across large materials databases.

Tensorial properties obey a different symmetry. The elastic stiffness tensor C_{ijkl} is a rank-4 object whose components transform under rotation according to the fourth-order tensor transformation law: each index is contracted with the rotation matrix. After reduction to Voigt notation the resulting 6×6 matrix still carries directional meaning; C11 quantifies stiffness along the first axis, C44 quantifies shear resistance in the corresponding plane, and the off-diagonal terms encode coupling between normal and shear strains. Because these components must co-vary with any rotation of the coordinate frame, an architecture that collapses the entire crystal into a single invariant scalar cannot represent the required transformation. The network is forced to output the same set of six numbers (or their averages) regardless of which axis is aligned with the laboratory frame, producing systematic collapse of distinctions that are physically real.

Conceptually, one can visualize the difference as follows: scalar prediction collapses the entire crystal graph into a single rotationally invariant number whose value is independent of coordinate choice, whereas tensor prediction must output a full matrix whose 21 independent entries (in the triclinic case) transform covariantly, preserving the relative magnitudes and signs that define anisotropy. The invariant pathway therefore necessarily erases the very information that distinguishes C11 from C33 or C12 from C13.

The failure is not merely quantitative; it is categorical. Invariant GNNs trained on elastic constants inevitably learn an averaged, isotropic-like response even when the underlying crystal is highly anisotropic. This explains why literature reports often show deceptively low average errors on cubic systems (where symmetry already forces C11 = C22 = C33) while the same models fail dramatically on orthorhombic or monoclinic compounds. The tensor transformation law, which is encoded in the physics, cannot be learned by an architecture that forbids directional representation at the outset.

Table 1 clarifies the underlying symmetry-alignment problem by distinguishing the property classes for which rotational invariance is physically appropriate from those for which it becomes structurally misaligned.

Table 1. Symmetry Alignment Matrix for Materials GNN Prediction: When Rotational Invariance Is Appropriate and When It Becomes Structurally Misaligned

Target property class	Physical transformation requirement	Appropriate architectural symmetry	Information needed for correct prediction	Suitability of invariant GNNs	Expected prediction behavior	Conceptual risk level
Scalar properties	Output remains unchanged under rotation	Rotational invariance	Composition, bonding environment, local chemistry, distance structure	High	Stable prediction under arbitrary coordinate rotation	Low
Vector properties	Output rotates with the structure	Rotational equivariance	Directional components and orientation tracking	Low	Directional collapse or frame-insensitive output if modeled invariantly	High
Rank-2 tensor properties	Output transforms covariantly under rotation	Rotational equivariance	Axis-resolved and angle-sensitive representation	Very low	Partial recovery of magnitudes but loss of directional structure	High
Rank-4 tensor properties reduced in Voigt form	Full tensor transformation law must be preserved across coordinate changes	Rotational equivariance or strong directional hybridization	Orientation of bonds relative to crystal axes, angular dependence, symmetry-sensitive coupling	Structurally inappropriate	Averaging of anisotropy, collapse of distinct components, isotropic-like prediction bias	Very high
High-symmetry tensor cases	Some equalities are physically enforced by crystal symmetry	Equivariance still preferred, though invariant error may be partially masked	Symmetry-aware directional representation	Apparent short-term adequacy only	Acceptable-looking average error can conceal deeper structural mismatch	Medium to high
Low-symmetry tensor cases	Distinct tensor entries must remain distinguishable under coordinate transformation	Rotational equivariance is essential	Full directional fidelity across normal and shear channels	Inadequate	Severe diagonal collapse, coupling errors, shear underspecification, false degeneracy	Very high

Mechanisms of Information Loss

Four distinct mechanisms, all inherent to the invariant GNN paradigm, produce the loss of directional information required for tensorial prediction.

Global pooling

Summation or averaging over all nodes or edges produces a single fixed-size descriptor that is invariant by construction. Any information about the relative orientation of one bond to another or to the global crystal axes is irretrievably averaged away. In elasticity, this prevents the network from learning that bonds aligned along the c-axis contribute differently to C33 than bonds along the a-axis.

Distance-only edge features

By representing edges solely through scalar distances, the model discards all angular information needed to resolve shear components. C44, C55, and C66 couple directly to bond-angle distortions; without angular descriptors, the network cannot differentiate shear resistance from normal stiffness, leading to systematic under-specification of off-diagonal and shear entries.

Symmetry-averaging

Invariant architectures treat all symmetry-equivalent bonds identically regardless of their absolute orientation in the crystal frame. In real materials with slight distortions—such as those near phase boundaries—equivalence is broken, yet the model continues to average them, producing artificial degeneracy in the predicted tensor.

Loss of bond-environment orientation

The orientation of a bond relative to the external crystal lattice (not merely relative to neighboring bonds) is never encoded. Tensorial properties require precisely this absolute directional reference; its absence forces the model to predict isotropic-like responses even in highly anisotropic crystals.

Each mechanism is individually sufficient to break tensorial fidelity; together they compound the error. The invariant models of Xie and Grossman [1], Chen *et al.* [3], and Schütt *et al.* [4] all instantiate these mechanisms simultaneously, guaranteeing the failure modes analyzed below.

A Typology of Failure Modes

Diagonal indistinguishability

Invariant GNNs cannot distinguish C11 from C22 or C33 in low-symmetry crystals. Because global pooling erases axis-specific information, the model collapses all diagonal stiffnesses to nearly identical values. Detection signature: predicted C11 ≈ C22 ≈ C33 even when experimental or DFT differences exceed 20 %. Rutile versus anatase TiO2 provides a clear example; the tetragonal symmetry already differentiates the c-axis, yet an invariant network predicts nearly equal diagonal entries.

Off-Diagonal blindness

Components such as C12, C13, and C23 encode Poisson-ratio coupling that requires explicit orientational relationships between non-collinear bonds. Distance-only features and pooling destroy these relationships, causing systematic underestimation of off-diagonal stiffness. Detection signature: absolute error on C12 is at least twice the error on C11.

Shear underspecification

Shear moduli C44, C55, and C66 depend on angular bond information that is absent from invariant edge features. The result is a systematic downward bias in predicted shear stiffness. Detection signature: the ratio error(C44)/error(C11) exceeds 3.

Symmetry-induced degeneracy

Invariant models enforce equality among symmetry-equivalent components by architectural design. When a real crystal exhibits slight symmetry breaking—e.g., a tetragonal-to-orthorhombic transition induced by doping or temperature—the model still predicts exact equality (C11 – C22 = 0) even when the true difference exceeds 5 GPa. Detection signature: predicted difference remains zero while the physical difference is finite.

These four modes are not mutually exclusive; most invariant predictions exhibit all four simultaneously, with the severity scaling with the degree of crystal anisotropy. The typology therefore supplies both a diagnostic taxonomy and a predictive framework for when invariant GNNs will fail on new materials.

Detection Principles

Detection of invariance-induced failures does not require new benchmarks or retraining; it can be performed on any existing invariant GNN output by applying four simple, architecture-agnostic tests that expose the systematic collapse of directional information. These principles exploit the fact that rotational invariance forces specific equalities and biases that are incompatible with the tensorial transformation law of elastic constants.

After prediction, compare the three normal stiffness components C11, C22, and C33. If the model is purely invariant, global pooling has erased axis-specific information, so these values will be nearly identical even in low-symmetry crystals where symmetry does not enforce equality. A difference smaller than 5 % across all three diagonals, while DFT or experimental values differ by more than 20 %, signals diagonal indistinguishability (Mode 1). Chen *et al.* [3] and Schütt *et al.* [4] models routinely fail this test on orthorhombic compounds because their mean-pooling layers treat all Cartesian directions as interchangeable.

Compute the predicted ratio C11/C33 (or equivalent pairs) and correlate it against the ground-truth ratio. An invariant GNN cannot learn direction-dependent scaling, so the predicted ratio will hover near 1.0 regardless of the true anisotropy. A Pearson correlation below 0.3 between predicted and true ratios across a diverse test set of tetragonal and hexagonal crystals indicates that the network has defaulted to an isotropic-like response. This check directly reveals the failure of the tensor transformation law that Batzner *et al.* [9] avoided by preserving directional features.

Measure mean-absolute error separately on diagonal (C11, C22, C33) versus off-diagonal (C12, C13, C23) components. Because off-diagonal terms encode coupling between non-collinear bond sets, distance-only edge features and global pooling produce errors on C12 that are at least twice as large as errors on C11. Xie and Grossman [1] noted the same pattern when CGCNN was applied to stress tensors; the audit simply makes the bias explicit without requiring architectural changes.

error magnitudes on shear moduli (C44, C55, C66) versus bulk moduli. Invariant networks systematically under-specify shear because angular information is absent; the ratio error(C44)/error(C11) therefore exceeds 3 in failing models. This principle is particularly diagnostic for monoclinic and triclinic systems where shear components dominate mechanical anisotropy. Thomas *et al.* [8] demonstrated that equivariant networks naturally pass this test by propagating angular features, confirming that the failure is architectural rather than data-related.

Taken together, these principles form a lightweight diagnostic suite that can be applied post hoc to any invariant GNN prediction. When two or more principles are violated simultaneously, the model is almost certainly exhibiting one or more of the four failure modes identified earlier. The principles therefore shift the burden from blind trust in average error metrics to targeted inspection of tensorial consistency.

Mitigation Principles

Once the failure modes are diagnosed, five interlocking mitigation principles restore directional fidelity without sacrificing the strengths of graph-based learning.

Replace rotationally invariant backbones with natural or equivariant graph networks that propagate features transforming under the rotation group [21]. Tensor Field Networks [8] and E(3)-equivariant graph neural networks [9] maintain equivariant message passing by using spherical harmonics and Clebsch-Gordan coefficients, ensuring that the output tensor automatically obeys the required transformation law $C'_{ijkl} = R_{ip}R_{jq}R_{kr}R_{ls}C_{pqrs}$. Batzner *et al.* [9] showed that such networks reduce force and stress errors by more than an order of magnitude precisely because directional information is never discarded.

Augment existing invariant architectures with explicit orientational descriptors. Spherical-harmonic expansions of bond vectors or bond-angle histograms relative to the crystal axes can be concatenated to edge features before message passing. This preserves the scalar-property pathway for energy while injecting the missing directional signal for tensor tasks. The hybrid design retains computational efficiency yet eliminates diagonal indistinguishability.

Train simultaneously on elastic constants and the full stress tensor under small deformations, following the broader logic of multi-task crystal graph learning [22]. The auxiliary stress target forces the network to learn consistent tensorial responses across rotations, acting as a soft equivariance regularizer. Chen *et al.* [3] demonstrated that multi-task objectives improve scalar properties; extending the same idea to stress enforces the transformation law that pure invariant models violate.

Principle 4: Uncertainty quantification for anisotropy regimes

Attach a secondary head that predicts the expected anisotropy ratio $\left(\frac{\max(Cii)}{\min(Cii)}\right)$. High predicted anisotropy combined with high epistemic uncertainty flags regimes where invariant collapse is likely. This principle allows practitioners to route high-risk predictions to equivariant fallback models.

Principle 5: Symmetry-aware augmentation

During training, apply random rotations to the input crystal graphs while keeping the target tensor transformed accordingly. Although the backbone remains invariant, the augmented data expose the model to the consequences of its symmetry assumptions, nudging learned embeddings toward partial directional sensitivity. When combined with directional encoding, this augmentation further reduces symmetry-induced degeneracy.

These five principles are not mutually exclusive; the strongest mitigation combines equivariant architectures [8, 9] with directional encoding and multi-task stress supervision. Applied together, they convert the liability of invariance into a deliberate, property-specific design choice. **Table 2** consolidates the manuscript's central analytical contribution by linking each invariance-induced failure mode to its architectural cause, observable signature, and most appropriate mitigation priority.

Table 2. Analytical Crosswalk Linking Invariance-Induced Failure Modes, Detection Signatures, and Mitigation Priorities for Elastic Tensor Prediction

Failure mode	Primary architectural cause	Tensor component pattern most affected	Typical diagnostic signature	Why average MAE can conceal the problem	Highest-priority mitigation
Diagonal indistinguishability	Global pooling removes axis-specific contributions	C11, C22, C33	Predicted diagonal components converge toward near-equality despite physically meaningful differences	Averaging across components can appear acceptable when one or two diagonal entries are approximately correct	Equivariant architectures plus directional encoding
Off-diagonal blindness	Distance-only edge features eliminate non-collinear relational structure	C12, C13, C23	Off-diagonal errors substantially exceed diagonal errors	Low error on dominant diagonal terms can hide persistent coupling failure	Directional encoding plus stress-tensor multi-task learning
Shear underspecification	Missing angular information and weak representation of bond-angle distortions	C44, C55, C66	Shear-component error remains disproportionately high relative to normal-stiffness error	Bulk-like behavior may still be captured, masking failure in mechanically decisive shear channels	Equivariant architectures plus angular descriptors
Symmetry-induced degeneracy	Symmetry averaging imposes false equality across near-distinct orientations	Pairs expected to separate under weak symmetry breaking	Predicted component differences remain zero or near-zero when true differences are finite	Symmetry-heavy benchmark sets can make enforced equality appear successful	Symmetry-aware augmentation plus equivariant fallback
General anisotropy collapse	Combined effect of all four information-loss mechanisms	Entire tensor structure	Predicted anisotropy ratios drift toward isotropic values	Component-wise error summaries do not test preservation of directional contrast	Full symmetry-aware redesign with equivariant backbone

Relation to Other Failure Modes

Invariance-induced tensor failure does not emerge as an isolated defect but instead entangles with, and intensifies, several structural limitations already characteristic of graph neural networks. Its interaction with over-smoothing is particularly consequential. Message passing in invariant architectures progressively attenuates distinctions between node embeddings, and this homogenization becomes more pronounced with depth. When such representations are subsequently subjected to global pooling, any residual directional variability is effectively eliminated. Under these conditions, even architectures in the CGCNN family [1, 23], which perform adequately at shallow depths, lose the capacity to discriminate between anisotropic elastic responses such as C11 and C33 once the network exceeds modest

depth. The issue is therefore not merely one of vanishing gradients or feature dilution, but of a systematic erasure of physically meaningful orientation-dependent information.

This dynamic also sharpens the constraints imposed by the information bottleneck inherent in pooling operations and by the known expressive limits of graph neural networks [24]. Compressing an entire crystal graph into a fixed-dimensional representation presupposes that the salient features of the system can be encoded without loss of essential structure, despite known limits in the logical expressiveness of GNNs [25]. That assumption becomes untenable when the prediction target is a rank-4 tensor with 21 independent components. In the absence of directional encoding, the mapping from graph structure to tensor output becomes underdetermined, forcing the model to approximate a fundamentally richer object using an impoverished latent space. While Schütt *et al.* [4] identified analogous compression effects in scalar prediction tasks, the implications here are more severe: the mismatch between representational capacity and output complexity introduces not just approximation error but structural inconsistency.

A related implication emerges in the context of geometric generalization. Invariant models, by construction, discard orientation information, which precludes them from learning transformation rules that govern how tensorial quantities respond to rotations. As a result, a model trained on a crystal in one coordinate frame lacks the capacity to produce physically consistent predictions when the same structure is presented under a different orientation. This stands in direct tension with the covariant nature of elastic constants, where rotation should induce predictable transformations rather than alter intrinsic predictions. The limitation remains largely concealed in datasets dominated by high-symmetry crystals, yet becomes immediately apparent when evaluating on low-symmetry systems where directional dependencies cannot be neglected [8, 9].

Beyond these architectural and geometric considerations, dataset composition plays a subtle but decisive role in obscuring the problem. Benchmark datasets in elasticity prediction are frequently skewed toward cubic and near-cubic materials, where symmetry constraints reduce the number of independent tensor components and partially align with the assumptions embedded in invariant models. This alignment creates an illusion of adequacy, as performance metrics fail to penalize the absence of directional sensitivity. However, evidence from broader materials literature, including Uchino's review of piezoelectric materials [14] and the community perspective by Stach *et al.* [12], indicates that this apparent robustness does not extend to more general settings. When low-symmetry materials are sufficiently represented, the limitations of invariance become not only visible but dominant.

Taken together, these interdependencies reposition invariance from a benign modeling choice to a structural source of compounded error. Its influence propagates through representation homogenization, latent space compression, and failures of geometric generalization, while dataset biases delay its detection. Addressing tensor-specific deficiencies in this context is therefore not a narrow technical adjustment but a broader corrective step that alleviates multiple, interlinked weaknesses in current GNN approaches.

Implications for Materials GNN Design

The failure-mode analysis carries direct consequences for three stakeholder groups.

For model developers, best practice must change in three ways, including greater attention to interpretability and explanation in GNN predictions [26]. First, report performance disaggregated by crystal system—cubic, tetragonal, orthorhombic, monoclinic, triclinic—because average error hides the severity of diagonal indistinguishability in low-symmetry classes. Second, publish diagonal distinguishability metrics (e.g., mean |C11 – C22| error) alongside MAE. Third, include stress-tensor or rotated-structure test cases in every release to expose invariance collapse.

For benchmark designers, tensor-property suites should mandate equivariance-aware metrics. Standard MAE on Voigt components is insufficient; benchmarks must incorporate the four detection principles and require that submitted models

either pass them or declare the invariance limitation explicitly. The community can no longer treat elasticity benchmarks as interchangeable with scalar tasks.

For practitioners deploying GNNs in materials discovery pipelines and modern deep learning frameworks [27], the guidance is equally concrete [21]. Do not rely on invariant GNNs for single-crystal anisotropy calculations without directional augmentation or equivariant fallback. When mechanical properties govern device performance—thermal mismatch in microelectronics, toughness in alloys, or piezoelectric response in sensors [28]—use equivariant architectures [8, 9] or hybrid directional encodings. Inverse design loops [11] that optimize for target elastic tensors will otherwise converge to spurious isotropic solutions.

Collectively, these implications shift materials GNN design from a one-size-fits-all invariance default toward symmetry-aware, property-specific architectures. The field can retain the remarkable scalar-prediction successes of CGCNN [1], MEGNet [3], SchNet [4], and broader machine learning/deep learning applications [29] while upgrading tensor pathways to equivariant or augmented designs.

Conclusion

Rotational invariance, once celebrated as the key to scalable materials property prediction, is revealed here as a liability when the target property is a tensor of rank two or higher. The same architectural choices—distance-only edges, global pooling, and symmetry averaging—that enable accurate scalar forecasts in CGCNN, MEGNet, and SchNet systematically erase the directional information required by elastic constants. This paper has articulated four mechanisms of information loss, developed a typology of four inevitable failure modes (diagonal indistinguishability, off-diagonal blindness, shear underspecification, and symmetry-induced degeneracy), supplied four detection principles, and outlined five mitigation principles centered on equivariant architectures [4, 5], directional encoding, and stress multi-task learning.

The consequences extend far beyond elasticity. Any directional or tensorial quantity—piezoelectric tensors, thermal conductivity, dielectric permittivity—will exhibit analogous failures in invariant GNNs. Zunger and Montoya *et al.* have called for physically consistent autonomous discovery; that consistency is impossible while invariance remains the default for tensor tasks.

The community is therefore urged to treat rotational invariance as a deliberate, property-specific choice rather than an architectural default. For scalar properties it remains powerful; for anisotropic elastic constants and other tensors it must be replaced or augmented. By adopting the detection principles and mitigation strategies presented here, future materials GNNs can deliver the directional fidelity that engineering design demands. The era of silent tensorial collapse in invariant predictions must end; the path forward is equivariant, directional, and rigorously symmetry-aware.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 05 May 2021 Revised: 16 Aug 2021 Accepted: 01 Nov 2021
Published online: 18 January 2022

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

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.

Gilmer J, Schoenholz SS, Riley PF, Vinyals O, Dahl GE. Neural message passing for quantum chemistry. In: Proceedings of the 34th International Conference on Machine Learning; 2017. p. 1263-72.

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.

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.

Li Z, Liu F, Yang W, Peng S, Zhou J. A survey of convolutional neural networks: Analysis, applications, and prospects. IEEE Trans Neural Netw Learn Syst. 2022;33(12):6999-7019.

Wu J. Introduction to convolutional neural networks. Nanjing: National Key Laboratory for Novel Software Technology, Nanjing University; 2017.

Cao Y, Romero J, Olson JP, Degroote M, Johnson PD, Kieferová M, et al. Quantum chemistry in the age of quantum computing. Chem Rev. 2019;119(19):10856-915.

Thomas N, Smidt T, Kearnes S, Yang L, Li L, Kohlhoff K, et al. Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds. arXiv preprint arXiv:1802.08219. 2018.

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.

Zhang Y, Tang Y, Liu Y, Zhang D, Zheng J. Design and engineering of amyloid aggregation-prone fragments and their antimicrobial conjugates with multi-target functionality. Adv Funct Mater. 2021;31(32):2102978.

Zunger A. Inverse design in search of materials with target functionalities. Nat Rev Chem. 2018;2(4):0121.

Stach E, DeCost B, Kusne AG, Hattrick-Simpers J, Brown KA, Reyes KG, et al. Autonomous experimentation systems for materials development: A community perspective. Matter. 2021;4(9):2702-26.

Szymanski NJ, Zeng Y, Huo H, Bartel CJ, Kim H, Ceder G. Toward autonomous design and synthesis of novel inorganic materials. Mater Horiz. 2021;8(8):2169-98.

Uchino K. The development of piezoelectric materials and the new perspective. In: Advanced piezoelectric materials. Woodhead Publishing; 2017. p. 1-92.

Wu Z, Pan S, Chen F, Long G, Zhang C, Yu PS. A comprehensive survey on graph neural networks. IEEE Trans Neural Netw Learn Syst. 2021;32(1):4-24.

Calin O. Deep learning architectures: A mathematical approach. Cham: Springer; 2020.

Choudhary K, DeCost B. Atomistic line graph neural network for improved materials property predictions. npj Comput Mater. 2021;7(1):185.

Gasteiger J, Becker F, Günnemann S. GemNet: Universal directional graph neural networks for molecules. Adv Neural Inf Process Syst. 2021;34:6790-802.

Jin D, Yu Z, Huo C, Wang R, Wang X, He D, et al. Universal graph convolutional networks. Adv Neural Inf Process Syst. 2021;34:10654-64.

Jackson R, Zhang W, Pearson J. TSNet: Predicting transition state structures with tensor field networks and transfer learning. Chem Sci. 2021;12(29):10022-40.

de Haan P, Cohen TS, Welling M. Natural graph networks. Adv Neural Inf Process Syst. 2020;33:3636-46.

Sanyal S, Balachandran J, Yadati N, Kumar A, Rajagopalan P, Sanyal S, et al. MT-CGCNN: Integrating crystal graph convolutional neural network with multitask learning for material property prediction. arXiv preprint arXiv:1811.05660. 2018.

Park CW, Wolverton C. Developing an improved crystal graph convolutional neural network framework for accelerated materials discovery. Phys Rev Mater. 2020;4(6):063801.

Xu K, Hu W, Leskovec J, Jegelka S. How powerful are graph neural networks? In: International Conference on Learning Representations; 2019.

Barceló P, Kostylev EV, Monet M, Pérez J, Reutter J, Silva JP. The logical expressiveness of graph neural networks. In: International Conference on Learning Representations; 2020.

Ying R, Bourgeois D, You J, Zitnik M, Leskovec J. GNNExplainer: Generating explanations for graph neural networks. Adv Neural Inf Process Syst. 2019;32.

Nguyen G, Dlugolinsky S, Bobák M, Tran V, Lopez Garcia A, Heredia I, et al. Machine learning and deep learning frameworks and libraries for large-scale data mining: A survey. Artif Intell Rev. 2019;52(1):77-124.

Behera A. Piezoelectric materials. In: Advanced materials: An introduction to modern materials science. Cham: Springer International Publishing; 2021. p. 43-76.

Shinde PP, Shah S. A review of machine learning and deep learning applications. In: 2018 Fourth International Conference on Computing Communication Control and Automation (ICCUBEA); 2018. p. 1-6. IEEE.