

ORIGINAL RESEARCH

Open access

Theory of Equivariant Error Cancellation in Multi-Property Materials GNNs: When Shared Representations Harm Specialization

Elif Yilmaz^{1*}, Mehmet Demir¹, Ayse Kaya², Hasan Aydin¹

Abstract

Multi-property graph neural networks (GNNs) have emerged as the dominant paradigm in computational materials science, enabling a single equivariant backbone to simultaneously predict energies, forces, stresses, elastic tensors, band gaps, and other key properties from atomic structures. Proponents emphasize gains in data efficiency, regularization, and inductive bias transfer across tasks. However, this work reveals a fundamental and previously underappreciated limitation: equivariant error cancellation within shared representations. By enforcing rotational, translational, and permutation symmetries, equivariant GNN backbones map atomic structures into a constrained feature space where different physical properties transform under distinct irreducible representations. We demonstrate theoretically and conceptually that gradients arising from weakly correlated or anti-correlated properties frequently conflict within these shared layers. Such conflicts drive error cancellation—reductions in the aggregate multi-task loss that mask stagnant or deteriorating accuracy on individual targets—particularly when properties map to incompatible tensorial characters (scalars, vectors, rank-2 and higher tensors). We identify three regimes of representation sharing: beneficial (strongly aligned gradients), neutral (orthogonal gradients), and harmful (opposing gradients leading to representational compromise). Proof sketches show that equivariance tightens the feasible representation space and bounds gradient alignment more severely than in invariant architectures, systematically increasing the risk of harmful cancellation. Detection diagnostics and architectural mitigation strategies are proposed to identify and alleviate these effects. This analysis reframes the design of multi-property materials GNNs, demonstrating that joint training is not universally beneficial and can actively degrade specialization. It provides a principled foundation for next-generation architectures that preserve data efficiency while safeguarding per-property accuracy, advancing reliable data-driven materials discovery and design.

Keywords Materials graph neural networks, Negative transfer, Equivariant representations, Multi-property GNNs, Error cancellation, Shared representations

*Correspondence:

Elif Yilmaz
elif.yilmaz@gmail.com

¹ Department of Materials Informatics and Systems Engineering, Faculty of Engineering, Istanbul Technical University, Istanbul, Turkey

² Department of Computational Materials Analytics, Faculty of Technology, Middle East Technical University, Ankara, Turkey

Introduction

Multi-property GNNs are increasingly common: a single model predicts energy, forces, stress, elastic constants, and more [1–3]. Sharing representations across properties improves data efficiency and reduces overfitting. But there is a hidden danger: error cancellation. Errors on different properties may cancel in the shared representation, reducing total loss while harming specialization. For example, the model might learn a representation that predicts energy well and forces

well individually, but the shared features are a compromise that degrades both. This paper provides a theoretical analysis of equivariant error cancellation in multi-property materials GNNs, identifying when shared representations harm specialization.

Graph neural networks have transformed materials prediction by directly operating on atomic graphs and respecting physical symmetries [4-9]. Early single-property models focused on formation energy or band gaps in isolation. As datasets grew, researchers observed that training on multiple targets simultaneously could leverage correlations and yield better generalization with fewer parameters [10-12]. In surfactant systems, multi-property GNNs predict surface tension, critical micelle concentration, and aggregation behavior from one backbone [1]. In metal-organic frameworks, joint prediction of adsorption energies and mechanical stability has accelerated screening campaigns [10]. Polymer informatics similarly benefits from multitask learning that captures both thermal and mechanical responses [13].

The appeal lies in the shared backbone. Message-passing layers extract local atomic environments once, and property heads read out the required scalars, vectors, or tensors. This design reduces memory footprint and training time while providing richer supervision per parameter. Proponents argue that auxiliary properties act as regularizers, preventing the network from overfitting to noise in any single target [2, 14]. Yet the literature also documents cases where multi-property models underperform specialized counterparts on one or more tasks, especially when properties exhibit different physical length scales or symmetry requirements [15, 16].

Error cancellation offers a mechanistic explanation for these observations. Because shared parameters receive a weighted sum of gradients from all heads, opposing signals can nullify updates or steer the representation toward an unsatisfactory compromise. In materials contexts, energy prediction often favors smooth, averaged features that capture thermodynamic stability, while force prediction demands sharp, derivative-sensitive features that capture local distortions. When these demands conflict, the shared layers settle on an intermediate representation that satisfies neither perfectly. The total loss still decreases because each property contributes a fraction of the error, yet individual property errors plateau or rise relative to single-task baselines.

Equivariant architectures intensify the issue. Modern materials GNNs enforce $E(3)$ or $SE(3)$ symmetry so that predictions transform correctly under rotation [17-19]. Energy is invariant (scalar), forces transform as vectors, stress as rank-2 tensors, and elastic constants as rank-4 tensors. A single backbone must therefore produce features projectable onto multiple irreducible representations. This additional constraint narrows the solution space and makes gradient opposition more probable than in purely invariant networks.

The present work develops a purely theoretical framework. No new datasets, simulations, or performance tables are introduced. Instead, formal definitions, conceptual examples drawn from established architectures, and proof sketches clarify the conditions under which sharing succeeds or fails. Three regimes—beneficial, neutral, and harmful—are delineated. The analysis shows why equivariant designs are particularly susceptible and proposes detection principles that practitioners can apply without extensive hyperparameter search.

By focusing on error cancellation, this theory complements existing studies of negative transfer and over-smoothing in graph networks [16, 20]. It explains why some multi-property models appear successful on aggregate benchmarks yet deliver suboptimal specialized predictions. The implications extend beyond academic curiosity. In high-stakes applications such as alloy design or battery material screening, a model that compromises on force accuracy while optimizing energy may yield unreliable molecular dynamics trajectories [21]. Understanding when shared representations harm specialization therefore becomes essential for trustworthy materials informatics.

Multi-Property GNN Architectures

Typical architecture begins with a shared backbone composed of multiple message-passing layers that update atomic embeddings while preserving symmetries [3, 4, 17]. Each layer aggregates information from neighboring atoms according

to equivariant or invariant rules. After several rounds of message passing, the network produces node-level or graph-level features. Property-specific heads then map these features to the required outputs: a scalar head for energy, a vector head for forces, a tensor head for stress, and higher-order projections for elastic constants [10, 12, 13]. The overall loss is a weighted sum of individual property losses, with weights chosen to balance scales or importance.

Common property combinations reflect practical materials workflows. Energy and forces are paired most frequently because forces are the negative gradient of energy, creating a natural consistency constraint [5, 17]. Adding stress enables prediction of mechanical response under strain [22]. Some models jointly forecast formation energy, band gap, and elastic moduli to support optoelectronic and structural screening [2, 23]. In polymer systems, multitask heads predict glass transition temperature alongside modulus and density [11, 13]. Surfactant GNNs handle surface tension together with micelle size and critical concentration [1].

Claimed benefits center on data efficiency. A single backbone learns once from all targets, effectively multiplying the supervisory signal per parameter [2, 14, 15]. This advantage is pronounced in materials science, where high-fidelity data from density-functional theory remain expensive [24]. Regularization emerges as a second benefit: each property supplies an inductive bias that constrains the shared representation, reducing the risk of memorizing dataset artifacts [12, 25]. Overfitting decreases because the model must satisfy multiple physical constraints simultaneously, encouraging more robust features [10, 26].

Yet these benefits conceal a hidden cost: compromise. The shared representation must simultaneously satisfy all property heads. When targets demand incompatible feature characteristics, the backbone cannot optimize fully for any one. Message-passing layers receive conflicting update directions, and the resulting features represent an uneasy average rather than a specialized solution [16, 20]. In practice, this manifests as a multi-property model that reports lower total loss than separate models yet delivers higher error on at least one property when evaluated in isolation.

Architectural variants attempt to mitigate compromise without abandoning sharing. Some insert lightweight adapter layers after the backbone, allowing modest specialization [15]. Others employ conditional message passing that modulates updates according to a global property identifier, though the core representation remains shared [3]. Despite these refinements, the fundamental tension persists: the backbone parameters are updated by the aggregate gradient, which can mask individual degradations.

Literature on multi-task graph networks in chemistry and materials repeatedly notes that performance gains are not universal [11, 13, 27]. Models trained on energy and forces often excel on thermodynamic quantities yet lag on vibrational or defect-related predictions. Polymer multitask networks capture global chain statistics well but sometimes lose fidelity on local torsional angles [13]. These observations align with the theoretical view that sharing is not inherently beneficial; its value depends on the alignment of property demands within the shared feature space.

In summary, multi-property GNN architectures trade modularity for efficiency. The shared backbone offers clear computational and statistical advantages when property gradients reinforce one another. When they do not, the architecture introduces error cancellation, the central subject of the present analysis. Understanding this trade-off at the representational level is prerequisite to designing architectures that retain sharing benefits without sacrificing specialization.

Error Cancellation: Definition and Mechanism

Error Cancellation — A phenomenon where errors on different properties have opposite effects on shared representation updates, causing total loss to decrease while individual property errors increase or stagnate.

The mechanism originates in the shared backbone parameters. Every training step computes gradients from each property head separately and then sums them, weighted by loss coefficients, before applying the update to the common

layers [12, 15, 16]. If the gradient vectors contributed by two properties point in roughly the same direction, their sum reinforces movement toward better performance on both. If the vectors point in opposite directions, they partially or fully cancel, leaving the shared parameters nearly unchanged even though each property could improve in isolation.

Consider a simple conceptual example. Suppose Property A requires a particular parameter θ to increase to reduce its error (positive gradient contribution). Property B requires the same θ to decrease (negative gradient contribution). When loss weights are comparable, the net gradient approaches zero. The optimizer takes no meaningful step, and both properties remain stuck at their current error levels. Total loss may still appear to decrease slowly through minor adjustments in the heads, yet the backbone fails to evolve [20, 25].

Materials-specific examples illustrate the same dynamic. Energy prediction typically rewards representations that encode average bond strengths and coordination environments, producing smooth feature landscapes suitable for thermodynamic integration. Force prediction, by contrast, requires representations sensitive to small displacements, capturing the precise curvature of the potential energy surface [5, 17]. A shared backbone that averages bond lengths effectively for energy may blur the derivative information needed for accurate forces. Consequently, gradient updates that sharpen features for forces can flatten them for energy, and vice versa. The net effect is a representation that compromises both, even as the summed loss metric trends downward.

The cancellation is insidious because standard training monitors only the composite loss. Practitioners observe convergence and assume success. Only careful isolation of per-property errors reveals that one or more targets have plateaued or worsened relative to single-property baselines [16]. In equivariant settings the problem intensifies because symmetry constraints already limit the expressivity of each layer. Opposing gradients must still respect the same tensorial structure, further restricting possible escape routes from the cancellation trap.

Information-theoretic perspectives reinforce the mechanism. Shared representations form by compressing input graphs into lower-dimensional embeddings that must serve multiple downstream readouts [25]. When the mutual information between the embedding and each property diverges, the embedding cannot simultaneously maximize relevance to all targets. Error cancellation is the optimization-level manifestation of this representational conflict.

Gradient cancellation also interacts with optimizer momentum and adaptive learning-rate schemes. Momentum may carry the parameters past a cancellation point temporarily, yet repeated opposition reestablishes the stalemate. Adaptive methods can shrink step sizes precisely when cancellation is strongest, further slowing progress on individual tasks.

In multi-property materials GNNs, cancellation appears most frequently when property pairs exhibit low or negative correlation in the training data. Formation energy and elastic moduli, for instance, share some structural dependence yet diverge sharply for defective or high-strain configurations. The shared layers receive mixed signals, and error cancellation follows. Recognizing the phenomenon at the gradient level rather than the loss level is therefore essential for theoretical understanding and practical diagnosis.

Regimes of Shared Representation: Beneficial, Neutral, Harmful

Shared representations in multi-property GNNs operate in three distinct regimes determined by the alignment of property-specific gradients within the shared backbone. Beneficial sharing emerges when properties are strongly correlated and their gradients consistently align, enabling the shared representation to improve performance on all targets beyond single-property models. Formation energy and decomposition energy, for example, depend on nearly identical structural descriptors, so gradients reinforce one another and the backbone learns richer, more transferable features [2, 23]. Under these conditions, multi-property models outperform specialized baselines because each task supplies complementary

supervisory signals that sharpen the common embedding, with the joint model recording lower error on every property than the corresponding single-property models trained with identical backbone depth.

A related dynamic arises when properties are statistically independent and their gradients remain largely orthogonal. The shared representation then performs approximately as well as separate models, neither helping nor harming specialization. Band gap and elastic constants in many inorganic crystals exhibit weak correlation, with gradients affecting largely disjoint feature subspaces [6, 28]. The backbone still gains from increased data volume and implicit regularization, yet no strong synergy or conflict materializes, such that the multi-property model matches single-property performance within statistical noise.

When properties are weakly correlated or anti-correlated, however, gradients oppose each other and harmful cancellation occurs. The shared representation forces compromises that degrade at least one property relative to specialized baselines. Energy prediction favors globally smooth features while defect sensitivity demands locally sharp distinctions, so gradients pull in opposite directions and produce an averaged embedding that satisfies neither [16, 20]. Although the total loss may decline, individual accuracies suffer, with the multi-property model underperforming a single-property baseline on at least one target. When gradient directions contributed by two properties are sufficiently opposed, shared representation updates induce error cancellation in which the combined loss decreases while individual property errors increase or stagnate. Equivariant representations amplify this cancellation because different properties must transform under distinct symmetry operations, imposing stricter alignment requirements that make gradient opposition more probable than in invariant architectures and shift the boundary between neutral and harmful regimes.

Table 1 formalizes the three regimes of shared representation by linking gradient alignment geometry to optimization behavior and specialization outcomes

Table 1. Structural Characterization of Gradient Alignment Regimes in Multi-Property GNNs


Regime	Gradient Alignment Structure	Property Correlation	Representation Behavior	Optimization Outcome	Specialization Effect	Detection Signature
Beneficial Sharing	Positively aligned gradients ($\cos \theta > 0$)	Strong positive correlation	Shared features reinforce across tasks	Accelerated convergence	Improves all properties beyond single-task baselines	All per-property errors decrease vs. baselines
Neutral Sharing	Orthogonal gradients ($\cos \theta \approx 0$)	Statistical independence	Feature subspaces weakly interacting	Stable but decoupled optimization	Comparable to single-task performance	Per-property errors match baselines within noise
Harmful Cancellation	Opposing gradients ($\cos \theta < 0$)	Weak or negative correlation	Compromise embedding across tasks	Net gradient suppression or oscillation	At least one property degrades	Total loss ↓ while individual losses ↑ or plateau

The regimes are not fixed properties of the tasks themselves but emerge from the interplay of dataset statistics, backbone capacity, and loss weighting. A pair of properties may exhibit beneficial sharing on one dataset and harmful cancellation on another if correlation structures differ. Similarly, increasing backbone depth can sometimes expand the feature space enough to accommodate orthogonal gradients, moving a pair from harmful to neutral.

Empirical hints of regime shifts appear across materials domains. Surfactant multi-property networks often remain in the beneficial regime because surface and bulk properties share hydrophobic descriptors [1]. Alloy design models, however, frequently enter harmful cancellation when mechanical and electronic properties compete for the same compositional embeddings [22, 28]. Polymer multitask GNNs hover near the neutral-to-harmful boundary depending on whether chain-level or monomer-level features dominate [13].

Understanding regime boundaries therefore becomes a prerequisite for architecture selection. Practitioners should not assume shared representations are universally advantageous. Instead, regime diagnosis through gradient alignment and baseline comparison guides the decision to share or separate. The propositions above provide conceptual tools for predicting regime membership before training, relying on correlation analysis and symmetry considerations rather than exhaustive experimentation.

Figure 1 illustrates the hierarchical gradient flow in multi-property GNNs and formalizes how alignment structure determines the transition between beneficial sharing, neutrality, and harmful error cancellation.



When Equivariance Makes Cancellation Worse

Equivariant representations differ fundamentally from invariant ones in their handling of multi-property tasks. Invariant backbones produce scalar features only, while equivariant layers generate entire sets of irreducible representations that transform correctly under rotation and translation [17-19]. This richer output space is essential for accurate tensorial predictions yet introduces tighter constraints when multiple properties share the same backbone.

Energy is a scalar invariant quantity. Forces transform as vectors. Stress is a rank-2 tensor, and elastic constants constitute a rank-4 tensor. A shared equivariant backbone must therefore output features that can be projected onto all these distinct symmetry classes simultaneously. The projection operations themselves become part of the representational compromise. Gradients flowing backward from a scalar energy head affect only the invariant subspace, while gradients from a vector force head update the vector subspace. Because the subspaces are coupled through the shared message-passing rules, an update that improves one subspace can inadvertently distort another.

Gradient conflict therefore manifests more severely in equivariant networks. Different property heads pull on overlapping but symmetry-distinct feature channels. An adjustment that aligns vector features for forces may rotate or scale invariant channels away from the energy optimum. The cancellation bound tightens because the symmetry group limits the directions in which gradients can constructively interfere.

Example: Energy plus forces. The energy head aggregates invariant scalars from the final node features. The force head computes vector differences or derivatives of the same features. During backpropagation, the scalar gradient and the vector gradient propagate through identical message-passing weights yet demand incompatible changes to the underlying embeddings. The shared layers receive a vector-gradient component that improves force directionality while simultaneously degrading the scalar smoothness required for energy. Cancellation follows directly.

Proposition 3 (Equivariant cancellation bound): In a multi-property equivariant GNN, when properties depend on different irreducible representations, the possible alignment of their gradients is inherently limited. This limitation increases the probability of cancellation compared with non-equivariant models, where all features remain scalar and alignment constraints are weaker.

The amplification effect scales with the difference in tensor rank. Scalar-vector pairs exhibit moderate cancellation risk. Scalar-rank-4 pairs, as in energy plus elastic constants, exhibit higher risk because the higher-order projection operators further decouple the effective gradient directions. Literature on equivariant GNNs for interatomic potentials already notes occasional degradation when stress or elastic terms are added without careful weighting [17]. The present theory supplies the mechanistic reason: symmetry-enforced gradient misalignment.

Equivariance also interacts with over-smoothing. Message passing in equivariant layers can wash out higher-order features faster than scalar ones, compounding cancellation when tensorial properties are involved [19]. The combination of symmetry constraints and smoothing creates a double bind: the representation must remain sufficiently expressive for all ranks yet sufficiently smooth for thermodynamic consistency. Shared parameters struggle to satisfy both, increasing harmful regime occupancy.

In summary, equivariant designs, while indispensable for physical consistency, make error cancellation more probable and more severe. The additional representational constraints narrow the path to beneficial sharing and widen the path to harmful compromise. Recognition of this amplification is essential for designers who default to equivariant backbones in modern materials GNNs. Theoretical awareness of the cancellation bound allows targeted architectural adjustments—such as rank-specific pathways or adaptive symmetry projections—before training begins.

Detection Principles

Detection of error cancellation relies on systematic comparison rather than aggregate loss alone [12, 16]. Practitioners must isolate the effects of sharing to reveal hidden compromises in the shared backbone. The gold-standard diagnostic begins with training separate specialized models for each property under identical backbone architecture and hyperparameter budgets, then directly comparing every individual property error of the multi-property model against these baselines; underperformance on any target signals harmful cancellation and quantifies the precise cost of gradient interference [10, 11, 13].

A related diagnostic examines gradient alignment by computing cosine similarity between property-specific gradients with respect to shared parameters at regular training intervals; consistently negative similarities provide early warning of cancellation with only modest optimizer instrumentation [15, 16, 20]. This insight extends naturally to loss landscape analysis, where divergence between a decreasing total loss and rising or plateauing individual property losses reveals opposing gradients neutralizing updates in shared layers [25, 27]. Further confirmation arises from representation similarity tests that quantify divergence—via cosine distance or centered kernel alignment—between embeddings from the multi-property backbone and those from single-property models, exposing compromise geometry [2, 14]. Preceding model construction, property correlation analysis offers predictive power: Pearson correlations below 0.3 or negative values sharply elevate cancellation risk and inform decisions on sharing viability [6, 23, 28].

Together these diagnostics form a lightweight pipeline requiring no new data and minimal extra computation, shifting diagnosis from post-hoc tables to mechanistic insight and enabling early detection of harmful cancellation before deployment in materials screening pipelines [1, 22].

Mitigation Principles

Once cancellation is detected, targeted interventions can restore specialization without discarding the efficiency of sharing [15, 16]. The following principles address the root cause—opposing gradients—while preserving the benefits of a shared backbone.

Table 2 consolidates the mechanistic origins of equivariant error cancellation and maps each to targeted architectural or optimization-level interventions.

Table 2. Mechanistic Sources and Mitigation Pathways for Equivariant Error Cancellation

Source of Cancellation	Mechanistic Origin	Affected Representation Level	Observable Symptom	Theoretical Implication	Targeted Mitigation Strategy
Gradient Opposition Across Tasks	Conflicting optimization directions from property losses	Shared backbone parameters	Slow or stalled learning despite decreasing total loss	Violates assumption of cooperative multitask learning	Gradient surgery (projection-based methods)
Irreducible Representation Mismatch	Scalar/vector/tensor transformation constraints	Equivariant feature channels	Inconsistent improvements across tensor ranks	Bounds achievable gradient alignment	Rank-specific adapters or decoupled pathways
Property Correlation	Weak or negative statistical relationships	Embedding information	One property dominates or	Limits mutual information	Property grouping or

Divergence		content	suppresses others	capacity of shared embedding	task clustering
Overlapping but Incompatible Feature Demands	Different spatial/scale sensitivities (global vs. local)	Message-passing feature aggregation	Blurred or over-smoothed representations	Forces compromise embedding geometry	Mixture-of-experts or conditional routing
Optimization-Level Suppression	Momentum and adaptive learning dampen conflicting updates	Parameter update dynamics	Apparent convergence without specialization gain	Masks true gradient conflict dynamics	Dynamic loss weighting and curriculum learning

Mitigation of harmful cancellation in multi-property GNNs begins with dynamic adjustment of loss weights to balance gradient magnitudes contributed by each property to the shared parameters; when one property dominates, increasing its relative weight prevents it from overwhelming weaker signals and thereby reduces cancellation, with scheduling during training providing further stability [12, 27]. This approach connects naturally to gradient surgery, which projects each property’s gradient onto the normal plane of the others before summation, removing conflicting components while preserving constructive directions exactly as in established multi-task methods, an inexpensive operation in materials GNNs since it targets only shared backbone gradients [16, 20].

A related strategy introduces lightweight property-specific adapter layers immediately after the shared backbone, allowing the core representation to retain general features while adapters restore specialization and decouple conflicting projections without duplicating the full network [11, 13]. More architectural shifts replace the single backbone with a mixture of experts and a lightweight router that selects or blends pathways according to property or input structure, routing strongly opposing properties to separate experts and eliminating gradient conflict at its source, with the router trained via a small auxiliary loss for load balancing [10, 26, 29]. Progressive multi-property learning offers another route by first training on the most correlated property alone before gradually introducing others under representation-regularization terms that penalize deviations from the initial specialized embedding, thereby preventing catastrophic interference [15, 16]. Finally, pre-clustering properties by gradient alignment or correlation and training separate multi-property models for each coherent group—such as energies with forces versus thermodynamic with defect properties—sacrifices some sharing yet removes harmful cancellation entirely [2, 22, 23].

These mitigation strategies remain modular and combinable according to the detected regime and compute budget. In practice they preserve the richer supervision inherent to multi-property training while neutralizing the representational compromises induced by equivariant error cancellation, enabling materials scientists to sustain data efficiency without compromising the specialization essential for reliable property prediction.

Relation to Other Theoretical Results

The theory of equivariant error cancellation connects directly to several established lines of research in multi-task learning and graph neural networks.

It extends standard multi-task learning theory, which traditionally assumes that tasks are positively related and that sharing will always improve generalization [12, 15]. The present framework identifies the precise boundary where tasks cease to be cooperative and become antagonistic, supplying a mechanistic account of failure modes that earlier theories left implicit.

It refines the concept of negative transfer [16]. Negative transfer is one observable consequence of harmful cancellation, yet the current analysis provides the underlying gradient-level explanation. When gradients oppose each other in the

shared backbone, the optimizer effectively transfers erroneous updates from one property to another, exactly reproducing the negative transfer phenomenon documented in materials ML. The theory therefore offers a diagnostic and corrective pathway beyond merely detecting transfer failure.

It deepens equivariance theory by revealing a previously under-appreciated cost [17–19]. Equivariant GNNs are celebrated for enforcing physical symmetries, yet the same symmetry constraints that guarantee correct transformation behavior also tighten the feasible space of shared representations. Different irreducible representations required by scalar energy and vector forces create inherent gradient misalignment, making cancellation more probable than in invariant architectures. This insight explains why equivariant multi-property models sometimes underperform their invariant counterparts despite superior single-task accuracy.

Finally, the theory complements analyses of over-smoothing in graph networks [19]. Over-smoothing collapses node representations across layers, while error cancellation collapses gradients across tasks. Both are failure modes of deep message passing in multi-property settings, yet they operate through orthogonal mechanisms—feature homogenization versus gradient opposition. Recognizing their independence allows architects to apply distinct remedies: residual connections or skip layers for smoothing, gradient surgery or adapters for cancellation.

By integrating these threads, the present work supplies a unified conceptual scaffold. It shows that error cancellation is not an isolated pathology but the natural outcome of gradient dynamics when shared representations confront conflicting symmetry and correlation demands. The framework therefore bridges multi-task optimization theory, negative transfer literature, equivariance constraints, and graph-network pathologies into a single coherent explanation tailored to materials GNNs [6, 20, 25].

Implications for Multi-Property GNN Design

The theoretical analysis carries immediate consequences for how researchers and engineers should construct multi-property materials GNNs.

For model developers the central lesson is caution: never assume shared representations are universally beneficial [2, 14]. Before committing to a joint architecture, developers must verify regime membership through the detection principles outlined earlier. When harmful cancellation is probable, the default should shift from monolithic sharing to hybrid designs that incorporate adapters, gradient surgery, or mixture-of-experts routing [11, 16]. Equivariant backbones in particular require extra scrutiny because their symmetry constraints amplify cancellation risk; rank-specific pathways or adaptive projection layers may become necessary components rather than optional refinements [17, 19].

For practitioners deploying these models in materials discovery campaigns the guidance is equally pragmatic. When screening for properties with low correlation—such as electronic and mechanical descriptors—separate single-property or grouped models often deliver superior accuracy and reliability [22, 23, 28]. Only when properties are demonstrably correlated should shared architectures be preferred, and even then only after single-property baselines confirm beneficial or neutral regime behavior. This disciplined approach prevents the deployment of models that appear accurate on aggregate metrics yet fail on the specific property most critical to the application.

For benchmark designers the implications concern evaluation methodology. Future multi-property benchmarks must include single-property baselines and report per-property metrics alongside total loss [10, 12, 13]. Correlation structure of the target set should be published so that regime membership can be anticipated. Without these elements, benchmark leaders risk rewarding models that achieve low total loss through cancellation rather than genuine representational synergy.

Collectively these implications advocate a more nuanced design philosophy. Representation sharing remains a powerful lever for data efficiency, yet it must be deployed with full awareness of its regime-dependent behavior. The field should

move from “share by default” to “diagnose then decide,” ensuring that multi-property GNNs deliver both efficiency and specialization. In doing so, the community will accelerate trustworthy materials informatics rather than merely accelerating training speed at the expense of predictive fidelity.

Conclusion

Multi-property GNNs with shared representations can suffer from error cancellation even when total loss decreases. The theory presented here identifies three regimes of shared representation behavior. Beneficial sharing occurs when properties are strongly correlated and gradients align, allowing mutual improvement. Neutral sharing arises with statistically independent properties whose gradients remain orthogonal, yielding performance comparable to specialized models. Harmful cancellation emerges when properties conflict, forcing the shared backbone into compromises that degrade specialization on at least one target.

Equivariant representations amplify cancellation because different properties transform under distinct irreducible representations. Scalar energy, vector forces, rank-2 stress, and rank-4 elastic constants impose symmetry constraints that make gradient opposition more probable and more severe than in invariant architectures. The resulting representational compromise is invisible to aggregate loss metrics yet directly harms the specialized accuracy required for reliable materials prediction.

Detection is straightforward once the right tools are applied: single-property baselines, gradient alignment tests, loss-landscape tracking, representation similarity, and pre-training correlation analysis together reveal cancellation before deployment. Mitigation follows directly from the mechanism—property weighting, gradient surgery, task-specific adapters, mixture of experts, progressive learning, and property grouping all neutralize opposing gradients while preserving the statistical benefits of joint training.

The analysis therefore reframes representation sharing from an automatic advantage to a regime-dependent choice. It calls for careful property selection, routine regime diagnosis, and cancellation-aware architecture design in next-generation multi-property materials GNNs. By clarifying when shared representations harm specialization, this theoretical framework equips the community to build models that are simultaneously data-efficient and property-specific, advancing computational materials engineering on a firmer conceptual foundation.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Published online: 18 January 2026

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

- Brozos C, Rittig JG, Bhattacharya S, Akanny E, Kohlmann C, Mitsos A. Graph neural networks for surfactant multi-property prediction. *Colloids Surf A Physicochem Eng Asp.* 2024;694:134133.
<https://doi.org/10.1016/j.colsurfa.2024.134133>.
- Kong S, Guevarra D, Gomes CP, Gregoire JM. Materials representation and transfer learning for multi-property prediction. *Appl Phys Rev.* 2021;8(2):021409.
<https://doi.org/10.1063/5.0047066>.
- Lu Y, Chen M, Zhang Q, Zhang J, Zhang C, Xu S, et al. Adaptive edge-aware graph convolutional with multi-task learning for simultaneous prediction of material properties. *npj Comput Mater.* 2026;12(1):49.
<https://doi.org/10.1038/s41524-025-01917-7>.
- 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>.
- 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>.
- 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):eadi3245.
<https://doi.org/10.1126/sciadv.adi3245>.
- Sathana V, Mathumathi M, Makanyadevi K. Prediction of material property using optimized augmented graph-attention layer in GNN. *Mater Today Proc.* 2022;69:1419-24.
<https://doi.org/10.1016/j.matpr.2022.09.500>.
- Corso G, Stark H, Jegelka S, Jaakkola T, Barzilay R. Graph neural networks. *Nat Rev Methods Primers.* 2024;4(1):17.
<https://doi.org/10.1038/s43586-024-00294-7>.
- Schultz LE, Gantzer N, Bobbitt NS, Sava Gallis DF, Dingreville R. Graph neural network-based multi-objective Bayesian optimization for enhanced screening of metal-organic frameworks with optimal separation performance. *J Mater Chem A.* 2026;14(18):10836-53.
<https://doi.org/10.1039/D5TA09133K>.

Queen O, McCarver GA, Thatigotla S, Abolins BP, Brown CL, Maroulas V, et al. Polymer graph neural networks for multitask property learning. *npj Comput Mater.* 2023;9(1):90.
<https://doi.org/10.1038/s41524-023-01034-3>.

Lupo Pasini M, Zhang P, Reeve ST, Choi JY. Multi-task graph neural networks for simultaneous prediction of global and atomic properties in ferromagnetic systems. *Mach Learn Sci Technol.* 2022;3(2):025007.
<https://doi.org/10.1088/2632-2153/ac6a51>.

Gurnani R, Kuenneth C, Toland A, Ramprasad R. Polymer informatics at scale with multitask graph neural networks. *Chem Mater.* 2023;35(4):1560-7.
<https://doi.org/10.1021/acs.chemmater.2c02991>.

Liu Y, Zhu D, Tian Z, Liu J, Ran X, Wang Z, et al. Transfer learning-enabled performance prediction of metallic materials: Methods, applications and prospects. *Int J Miner Metall Mater.* 2026;33(3):749-67.
<https://doi.org/10.1007/s12613-025-3267-7>.

Rahman CMA, Alam NB, El-Wakeel AS, Ryu J, Gyawali PK. Multi-task learning for material property prediction. In: 2024 International Conference on Machine Learning and Applications (ICMLA); 2024 Dec 18-20; Miami, FL, USA. IEEE; 2024. p. 1065-70.
<https://doi.org/10.1109/ICMLA61862.2024.00161>.

Kontar R, Raskutti G, Zhou S. Minimizing negative transfer of knowledge in multivariate Gaussian processes: A scalable and regularized approach. *IEEE Trans Pattern Anal Mach Intell.* 2021;43(10):3508-22.
<https://doi.org/10.1109/TPAMI.2020.2987482>.

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>.

Cen J, Li A, Lin N, Xu T, Rong Y, Zhao D, et al. Universally invariant learning in equivariant GNNs. *arXiv:2510.13169 [Preprint]*. 2025.
<https://doi.org/10.48550/arXiv.2510.13169>.

Liu Y, Cheng J, Zhao H, Xu T, Zhao P, Tsung F, et al. SEGNO: Generalizing equivariant graph neural networks with physical inductive biases. *arXiv:2308.13212 [Preprint]*. 2023.
<https://doi.org/10.48550/arXiv.2308.13212>.

Ni N, Gathercole SE, Norris D, Saito S. Asymmetric negative transfer effects of working memory training. *Mem Cognit.* 2023;51(7):1654-69.
<https://doi.org/10.3758/s13421-023-01412-8>.

Yu J, Xi S, Pan S, Wang Y, Peng Q, Shi R, et al. Machine learning-guided design and development of metallic structural materials. *J Mater Inf.* 2021;1(2):9.
<https://doi.org/10.20517/jmi.2021.08>.

Zhao Z, Hu T, Bi S, Guan D, Xu S, Chen C, et al. Advances in graph neural networks for alloy design and properties predictions: A review. *J Mater Inf.* 2026;6:2.
<https://doi.org/10.20517/jmi.2025.42>.

De Breuck PP, Piracha HA, Rignanese GM, Marques MAL. A generative material transformer using Wyckoff representation. *npj Comput Mater.* 2026;12(1):60.
<https://doi.org/10.1038/s41524-025-01940-8>.

Shim E, Kammeraad JA, Xu Z, Tewari A, Cernak T, Zimmerman PM. Predicting reaction conditions from limited data through active transfer learning. *Chem Sci.* 2022;13(22):6655-68.

<https://doi.org/10.1039/D1SC06932B>.

Eguchi A, Horii T, Nagai T, Kanai R, Oizumi M. An information theoretic approach to reveal the formation of shared representations. *Front Comput Neurosci.* 2020;14:1.

<https://doi.org/10.3389/fncom.2020.00001>.

Hao Y, Yan X, Duo L, Hu H, He J. Diffusion models for 3D molecular and crystal structure generation: Advancing materials discovery through equivariance, multi-property design, and synthesizability. *Cambridge Open Engage* [Preprint]. 2025.

<https://doi.org/10.33774/coe-2025-9w4ph>.

Nevolianis T, Rittig JG, Mitsos A, Leonhard K. Multi-fidelity graph neural networks for predicting toluene/water partition coefficients. *J Cheminform.* 2025;17(1):123.

<https://doi.org/10.1186/s13321-025-01057-6>.

Zhang J, Koneru A, Sankaranarayanan SKRS, Lilley CM. Graph neural network guided evolutionary search of grain boundaries in 2D materials. *ACS Appl Mater Interfaces.* 2023;15(16):20520-30.

<https://doi.org/10.1021/acsami.3c01161>.

Chang R, Wang YX, Ertekin E. Towards overcoming data scarcity in materials science: Unifying models and datasets with a mixture of experts framework. *npj Comput Mater.* 2022;8(1):242.

<https://doi.org/10.1038/s41524-022-00929-x>.