Symmetry-Guaranteed Prediction of High-Order Tensor Properties for Crystalline Materials via Irreducible Decomposition

Qiaolin Lu[1] **Qiang Qu**[2] **Hao Jiang**[2] **Aoni Xu**[2] **Fengwang Li**[2]
Bo Han[3] **Tongliang Liu**[2†] **Yi Chang**[4] **Chengqi Zhang**[1]

[1]The Hong Kong Polytechnic University [2]The University of Sydney
[3]Hong Kong Baptist University [4]Jilin University

Abstract

Predicting high-order tensor properties for crystalline materials is crucial for various scientific and engineering applications. Crystal symmetry is one of the primary factors influencing high-order tensor properties, such as elasticity and piezoelectricity, making strict adherence to symmetry constraints essential. However, combining accurate prediction with exact symmetry compliance remains challenging. In this work, we propose a novel method that guarantees exact symmetry compliance by predicting symmetry-constrained irreducible components of high-order tensors. Specifically, we first develop a computational procedure to identify the basis tensors corresponding to symmetry-constrained irreducible components under various symmetry conditions. This symmetry-constrained basis guarantees that the assembled full tensor strictly adheres to the required symmetry constraints. To predict the numerical values for these irreducible components, we then propose a spherical-harmonic convolutional neural network designed to effectively capture essential high-order tensor information. Extensive experiments validate that our method achieves exact symmetry compliance without compromising prediction accuracy, thereby outperforming state-of-the-art approaches. Our code is publicly available at: `github.com/QiaolinLu/IrredNet`.

1 Introduction

Tensorial properties of crystalline materials, such as dielectric permittivity, piezoelectric coefficients, elastic stiffness, and higher-order nonlinear optical susceptibilities, are fundamental for technological advancements in electronics, photonics, energy harvesting, and quantum optics [1, 2, 3]. The accurate prediction of these anisotropic and orientation-dependent properties directly impacts material discovery and design, as their precise characterization is intimately connected to the performance of crystalline materials in practical applications. Achieving reliable predictions requires careful consideration of the intrinsic symmetry properties inherent in crystalline structures.

Indeed, crystal symmetry significantly governs tensor properties by constraining how tensors transform under symmetry operations, thereby determining a material's macroscopic behavior. For instance, materials with cubic symmetry, such as diamond, exhibit uniform deformation under stress, whereas those with hexagonal symmetry, such as graphite, display anisotropic deformation behaviors dependent on crystallographic directions [4, 5, 6]. Moreover, even within a nominal symmetry class, slight deviations or symmetry-breaking distortions can induce profound changes in response; for example, the cubic-to-tetragonal distortion in perovskites such as $BaTiO_3$ gives rise to ferroelectricity

†Corresponding author: Tongliang Liu (`tongliang.liu@sydney.edu.au`).

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

and markedly different elastic and dielectric behaviors [7, 8]. Therefore, explicitly incorporating symmetry constraints into predictions is essential for accurately modeling tensorial behavior; yet strict enforcement, particularly for higher-order tensors such as the elastic tensor, remains computationally and methodologically challenging.

Although first-principles methods, most notably Density Functional Theory (DFT), yield quantitatively reliable and symmetry-consistent predictions across a broad range of properties, their poor scaling and the need for large supercells and dense Brillouin-zone sampling render them prohibitively expensive for large systems, complex configurational spaces, and high-throughput studies [9, 10]. To bypass this bottleneck, machine learning (ML) surrogates are increasingly used to predict tensorial properties at a fraction of the cost; however, the speed gains often come at the expense of fidelity. In practice, ML models provide only approximate emulations of DFT and, crucially, struggle to rigorously preserve the directional dependence and crystallographic symmetry constraints that tensors must satisfy. Many widely adopted frameworks were originally developed with scalar targets in mind—via descriptor-based methods [11, 12, 13, 14] or (rotation-)invariant graph neural networks [15, 16, 17, 18, 19, 20, 21], and when repurposed for tensors they would violate index-level symmetries, equivariance, and crystal point-group constraints, leading to biased or internally inconsistent predictions that become more pronounced for higher-order tensors (e.g., elasticity or piezoelectricity).

To address these shortcomings, recent developments in equivariant neural networks aim to explicitly enforce symmetry transformations on predicted tensors [22, 23, 24, 25, 26, 27, 28]. Despite notable improvements, such symmetry enforcement approaches remain approximate, relying predominantly on crystal-specific rotation or reflection matrices to incorporate symmetry constraints. While effective for lower-order tensors, these methods become increasingly problematic as tensor complexity and order rise, leading to subtle but significant symmetry violations [29, 30, 31]. These discrepancies become critical for properties where strict symmetry compliance is mandatory. For example, elastic tensors (4-th order), which determine direction-dependent mechanical properties such as Young's modulus, require exact symmetry adherence to preserve the expected tensor structure and avoid spurious anisotropic artifacts in downstream calculations [32, 33]. Even minor symmetry violations in high-order tensor predictions can drastically impair the accuracy of subsequent calculations and predictions, undermining their reliability and utility for practical engineering applications.

In this work, we propose an exact, symmetry-guaranteed framework for predicting high-order tensor properties of crystalline materials by learning in the irreducible-component space rather than in the full-tensor space. The core guarantee is that symmetry is satisfied regardless of the prediction accuracy, since even if the learned magnitudes of the irreducible components are imperfect, the reconstructed target tensor obeys the crystallographic constraints by construction. This follows from three facts: (i) any high-order tensor subject to a crystallographic symmetry admits an exact decomposition into a finite set of irreducible components, (ii) a unique reconstruction of the target tensor is obtained via a combination of these components, and (iii) symmetry enforcement is straightforward for the irreducible components via canonical bases or projection, in contrast to the fragile high-order constraints required on the full tensor. Building on this pipeline, we enumerate all the appropriate irreducible components and reconstruct tensors exclusively from them, so any output is symmetry compliant by design. We then use a spherical-harmonic convolutional neural network to predict the magnitudes of irreducible components. Moreover, since the target tensors are high-dimensional and difficult to predict as a whole, decomposing them into irreducible components with straightforward symmetry guarantees has advantages in achieving better accuracy. Specifically, by reconstructing the target tensor using the irreducible components, the search space has been reduced to a small number of degrees of freedom, resulting in accuracy improvement while preserving exact symmetry.

2 Related Work

Datasets. Datasets of tensor-valued material properties (e.g., elasticity and piezoelectric tensors) are critical for modeling complex anisotropic behavior and are increasingly produced by first-principles calculations such as DFT [10]. Prominent public repositories like the Materials Project [34] and JARVIS-DFT [35] now curate extensive collections of these computed tensor properties, facilitating large-scale data-driven materials discovery. These datasets typically include not only high-order tensor properties but also other essential material information such as atomic positions, edge connectivity (bond order), and crystal constants. Additionally, they encompass related downstream properties, such

as Young's modulus, shear modulus, and other mechanical, thermal, or electronic properties, which are vital for evaluating the overall performance and behavior of materials under various conditions.

Equivariant Message Passing Network. Equivariant Message Passing Neural Networks (MPNNs) [24, 36, 37, 38, 39, 22] are a type of Graph Neural Networks (GNNs) designed to efficiently handle Euclidean symmetry operations, such as rotations, reflections, and translations. In most MPNNs, messages are propagated using a set of spherical bases, allowing the model to learn not only vector-based features but also high-order tensor features across different layers. Thus, MPNNs become a powerful tool for learning the representations of crystalline materials.

Methods for Enforcing Symmetry in High-Order Tensor Properties. To enforce symmetry in high-order tensor properties, many models adopt different strategies, such as incorporating space group symmetries [25] or using matrix decomposition [27, 28] techniques. While these methods help in aligning the learned tensor representations with the expected symmetries, most of them rely on approximate symmetry enforcement, rather than ensuring exact symmetry compliance.

3 Methods

This section is devoted to the prediction of high-order tensor properties in crystalline materials and is structured into three complementary components: **theory, computational procedure and model design**. (1) We begin with a rigorous group-theoretic analysis in Section 3.1, demonstrating that any rank-k tensor defined under the rotation group $\mathrm{SO}(3)$, even with external symmetry constraints, admits a finite and exact decomposition into irreducible subspaces. This establishes the mathematical foundation for the subsequent algorithmic and architectural developments. By assembling these basis tensors back into the full tensor, the proposed framework guarantees exact symmetry compliance by construction. (2) Building upon this theoretical result, we develop an efficient computational procedure in Section 3.2 to determine the set of irreducible representations that remain valid under specific symmetry constraints. This enables a systematic identification of relevant subspaces for accurate tensor prediction. (3) Finally, in Section 3.3, we introduce an equivariant neural network with spherical-harmonic convolutions that operates directly in irreducible spaces to predict high-order tensor components. This design ensures compliance with $\mathrm{SO}(3)$-equivariance and enables both physical consistency and predictive accuracy.

3.1 Theory for High-Order Tensor Decomposition

An essential step in predicting high-order tensor properties is to decompose the tensor into its irreducible components and represent it as a direct sum of these fundamental parts under the action of the $\mathrm{SO}(3)$ rotation group. In this section, we demonstrate how such a decomposition can be efficiently performed using group-theoretic principles.

We begin by presenting two propositions to illustrate the logic behind the irreducible decomposition of high-order tensors:

Proposition 1 Let V be a three-dimensional Euclidean vector space. Any rank-k ($k > 1$) tensor $\mathcal{T}_{i_1 \cdots i_k} \in V^{\otimes k}$ admits a unique decomposition into irreducible representations (irreps) of $\mathrm{SO}(3)$ group labeled by integer spins ℓ as follows:

$$V^{\otimes k} \simeq \bigoplus_{\ell \in L_k} \left(\mathcal{S}_k^{(\ell)} \bigoplus \mathcal{A}_k^{(\ell)} \bigoplus \mathcal{M}_k^{(\ell)} \right), \tag{1}$$

where L_k is the set of all distinct integer spins ℓ that appear in the decomposition of $V^{\otimes k}$. $\mathcal{S}_k^{(\ell)}$, $\mathcal{A}_k^{(\ell)}$, and $\mathcal{M}_k^{(\ell)}$ denote the symmetric, the antisymmetric, and the mixed-symmetry components transforming under spin-ℓ irreps, respectively.

Proposition 2 Let $\mathcal{T}$ be a rank-k tensor defined over a three-dimensional Euclidean space. Under decomposition into irreducible representations (irreps) of $\mathrm{SO}(3)$ group: When $\mathcal{T}$ satisfies additional index symmetries (e.g., major/minor symmetry, tracelessness, or antisymmetry), the set of allowed irreps forms a subset of the unconstrained case: $\mathrm{Irrep}(\mathcal{T}_{\text{sym}}) \subseteq \mathrm{Irrep}(\mathcal{T}_{\text{unconstrained}})$.

Together, Propositions 1 and 2 (whose proofs are detailed in Appendix A.1) establish a clear roadmap for exact, symmetry-aware tensor decomposition under $\mathrm{SO}(3)$: **Unconstrained decomposition**. Any rank-k tensor admits a unique decomposition into a finite direct sum of irreducible subspaces, comprising symmetric traceless, antisymmetric, and mixed-symmetry components, where each carries a well-defined spin ℓ and dimension $2\ell + 1$. **Effect of additional symmetries**. Imposing index symmetries (e.g. major/minor index exchange, tracelessness, antisymmetry) simply eliminates those irreducible representations that are incompatible with the constraint. In other words, the set of allowed irreps under symmetry is a subset of the full unconstrained set.

These results imply that once we have computed the full $\mathrm{SO}(3)$-irreducible decomposition of a tensor via Proposition 1, we obtain its symmetry-constrained form by discarding the forbidden irreps identified in Proposition 2. Specifically, by assembling the components based on their basis, we ensure that the resulting full tensor satisfies all symmetry constraints. This process guarantees that the final decomposition is exact and finite, inherently respecting all imposed symmetries, thus providing a solid theoretical foundation for our prediction model based on irreducible components.

3.2 Computational Procedure for Computing Irreducible Decomposition

Algorithm 1 implements a two-step method to extract all symmetry-compatible spin-ℓ components of a rank-k tensor under an arbitrary point-group symmetry G_{sym}:

Unconstrained Decomposition. Based on **Proposition 1**, we introduce the procedure GENERATESO3IRREPS, which enumerates all symmetric components $\ell = k,\ k-2,\ k-4, \ldots$, all antisymmetric components $\ell = 2, 3$ and mixed-symmetric components to record the multiplicity of each ℓ in a temporary map $\mathcal{M}_0[\ell]$ [40, 41, 42]. (see more details in Appendix A, B and J)

Symmetry Compatibility. Starting from the unconstrained decomposition $\mathcal{M}_0$, the function APPLYSYM maintains only the degrees of freedom that are invariant under symmetry constraints G_{sym}. Specifically, for each component (ℓ,m_ℓ) in $\mathcal{M}_0$, we first calculate a class sum over discrete symmetry operations in G_{sym}: $S_\ell = \sum_{C_j \subseteq G_{\mathrm{sym}}} |C_j|\, \chi^{(\ell)}_{O(3)}(\alpha_j, \det_j)$, where $\chi^{(\ell)}_{O(3)}(\alpha, \det) = s(k, \det)\, \chi^{(\ell)}(\alpha)$ is the character of the ℓ-representation evaluated at a representative $(\alpha, \det)$ of each conjugacy class and $s(k, \det) = (\det)^k$ is the parity correction factor that accounts for improper operations [43, 44]. The characters and projectors use the induced rank-k polar-tensor representation; the relation to the spherical-harmonic representation is detailed in Appendix J. The invariant multiplicity is then calculated by $n^{(\ell)}_{\mathrm{triv}} = \max\{0,\ \mathrm{round}(S_\ell/|G_{\mathrm{sym}}|)\}$ [45, 46]. Consequently, we only record the usable coefficients as $\mathcal{M}_{\mathrm{final}}[\ell] = m_\ell\, n^{(\ell)}_{\mathrm{triv}}$ when $n^{(\ell)}_{\mathrm{triv}} > 0$. When an explicit basis is needed, we project the invariant subspace via $\Pi_\ell = |G_{\mathrm{sym}}|^{-1} \sum_{g \in G_{\mathrm{sym}}} \rho^{(\ell)}(g)$. In brief, APPLYSYM counts, for each ℓ, the dimensions that remain unchanged under the symmetry constraint and keeps only those for reconstruction, guaranteeing symmetry by design.

3.3 Model Design

We observe that the irreducible decomposition of high-order tensors aligns naturally with spherical harmonics, as both share the same $\mathrm{SO}(3)$ representation structure. Each irreducible component corresponds to a spherical harmonic with angular quantum number ℓ, making spherical harmonics an ideal basis for modeling rotation-equivariant tensor properties.

Motivated by this insight, we specifically design our crystal tensor prediction model IrredNet (Figure 1) to incorporate spherical harmonic bases, enabling the direct construction of irreducible representations with built-in $\mathrm{SO}(3)$-equivariance. This design not only respects the intrinsic geometric symmetries of tensorial data but also facilitates more physically consistent learning, leading to better generalization and interpretability. It comprises three core components: (1) Crystal Graph Construction Block: building a graph based on the crystal structure; (2) Equivariant Interaction Block: propagating and aggregating features in an $\mathrm{SO}(3)$-equivariant manner via a learnable gating mechanism; (3) Irreducible Decomposition Block: extracting and organizing the tensor's $\mathrm{SO}(3)$ irreducible components for final property estimation.

Algorithm 1 SO(3) Irreducible Decomposition of Rank-k Tensor under Symmetry Constraints G_{sym}

Input: Tensor rank k; Point group symmetry G_{sym}
Output: Final map $\mathcal{M}_{final}$ from irreducible decomposition of k-th order tensor

---

Function GenerateSO3Irreps(k)
$\mathcal{M}_{\text{final}} \leftarrow \{\ell \mapsto 0\}$
$\text{Parts} \leftarrow \text{GeneratePartitions}(k, 3)$
for each $\lambda = (\lambda_1, \lambda_2, \lambda_3) \in \text{Parts}$ **do**
  // Step 1: Enumerate partitions
  $\text{mult}_{\text{GL3}} \leftarrow \text{HookLengthFormula}(\lambda)$
  // Step 2: GL(3) multiplicity
  $p \leftarrow \lambda_1 - \lambda_2$, $q \leftarrow \lambda_2 - \lambda_3$
  // Step 3: Branching SU(3) -> SO(3)
  $\mathcal{M}_{\text{branch}} \leftarrow \text{BranchSU3toSO3}(p, q)$
  // Step 4: Accumulate results
  for each $(\ell, m_\ell) \in \mathcal{M}_{\text{branch}}$ **do**
    $\mathcal{M}_{\text{final}}[\ell] \leftarrow \mathcal{M}_{\text{final}}[\ell] + m_\ell \times \text{mult}_{\text{GL3}}$
  end for
end for
return $\mathcal{M}_{\text{final}}$
End Function

Function CharO3($\ell, \alpha, \det$)
$\alpha \leftarrow \text{mod}(\alpha, 2\pi)$
$s \leftarrow \begin{cases} 1, & \det = +1 \\ (-1)^\ell, & \det = -1 \end{cases}$
if $\alpha = 0$ or $\alpha = 2\pi$ **then**
  return $s \cdot (2\ell + 1)$
else
  return $s \cdot \dfrac{\sin\big((\ell + 0.5)\alpha\big)}{\sin(\alpha/2)}$
end if
End Function

Function IsCompat(ℓ, G_{sym}, k)
$S \leftarrow 0$
for each $(|C_j|, \alpha_j, \det_j) \in G_{\text{sym}}$ **do**
  $\chi_j \leftarrow (\det_j)^{k-\ell}\, \text{CharO3}(\ell, \alpha_j, \det_j)$
  $S \leftarrow S + |C_j| \cdot \chi_j$
end for
$n_\ell \leftarrow \left\lfloor \dfrac{S}{|G_{\text{sym}}|} + 0.5 \right\rfloor$
return n_ℓ
End Function

Function ApplySym($\mathcal{M}_0$, G_{sym}, k)
$\mathcal{M}_{\text{final}} \leftarrow \emptyset$
for each $(\ell, m) \in \mathcal{M}_0$ **do**
  $n_\ell \leftarrow \text{IsCompat}(\ell, G_{\text{sym}}, k)$
  if $n_\ell > 0$ **then**
    $\mathcal{M}_{\text{final}}[\ell] \leftarrow \mathcal{M}_{\text{final}}[\ell] + m \times n_\ell$
  end if
end for
return $\mathcal{M}_{\text{final}}$
End Function

// Main Process:
$\mathcal{M}_0 \leftarrow \text{GenerateSO3Irreps}(k)$
Verify $\sum_\ell (2\ell + 1) \cdot \mathcal{M}_0[\ell] = 3^k$
$\mathcal{M}_{final} \leftarrow \text{ApplySym}(\mathcal{M}_0, G_{sym}, k)$
return $\mathcal{M}_{final}$

// Note: Helper functions in Appendix J.

Crystal Graph Construction Block. A crystal structure is defined by a 3×3 lattice matrix that encodes the three-dimensional periodicity of the unit cell, together with the set of basis atoms occupying that cell. To be specific, the structure can be mathematically represented by **M = (A, P , L)**, where $\mathbf{A} = [\boldsymbol{a}_1, \boldsymbol{a}_2, \cdots, \boldsymbol{a}_n] \in \mathbb{R}^{d_a \times n}$ contains the d_a-dimensional feature vectors for **n** atoms in the unit cell, $\mathbf{P} = [\boldsymbol{p}_1, \boldsymbol{p}_2, \cdots, \boldsymbol{p}_n] \in \mathbb{R}^{3\times n}$, $\boldsymbol{L} = (\boldsymbol{\ell}_1, \boldsymbol{\ell}_2, \boldsymbol{\ell}_3) \in \mathbb{R}^{3\times 3}$ describes the repeating patterns of the unit cell structure in 3D space. To enable networks to handle such infinite crystal structures, we follow the strategy in previous work [47, 30, 24] to build a graph to describe structures and relationships. Each node in the proposed SE(3) invariant crystal graph represents the atom i and all its infinite duplicates in the 3D space with positions $\{\hat{p}_i \mid \hat{p}_i = p_i + k_1\ell_1 + k_2\ell_2 + k_3\ell_3, k_1, k_2, k_3 \in \mathbb{Z}\}$ and atom feature a_i. The edge is built from the source node s to the target node t when the euclidean distance $\|e_{ts}\|_2$ between a duplicate of t and s in a periodic system satisfies $\|e_{ts}\|_2 \leq r$, where $r \in R$ is the cutoff radius determined by the k-th nearest neighbor.

Equivariant Interaction Block. This block is divided into three parts: the first part uses an embedding layer to obtain the basic representation of the graph; the second part employs equivariant spherical-harmonic convolution to capture high-order features; and the final part aggregates hierarchical order information at the crystal graph level. Following previous work [47, 30, 24], we map node type into a one-hot feature representation. We embed each edge length $\|e_{ts}\|$ by expanding it onto a preset family of basis functions, producing a fixed-length vector that captures both the raw distance and its multi-scale variations. Edge directional vectors e_{ts} are embedded by corresponding spherical harmonics $Y(e_{ts})$.

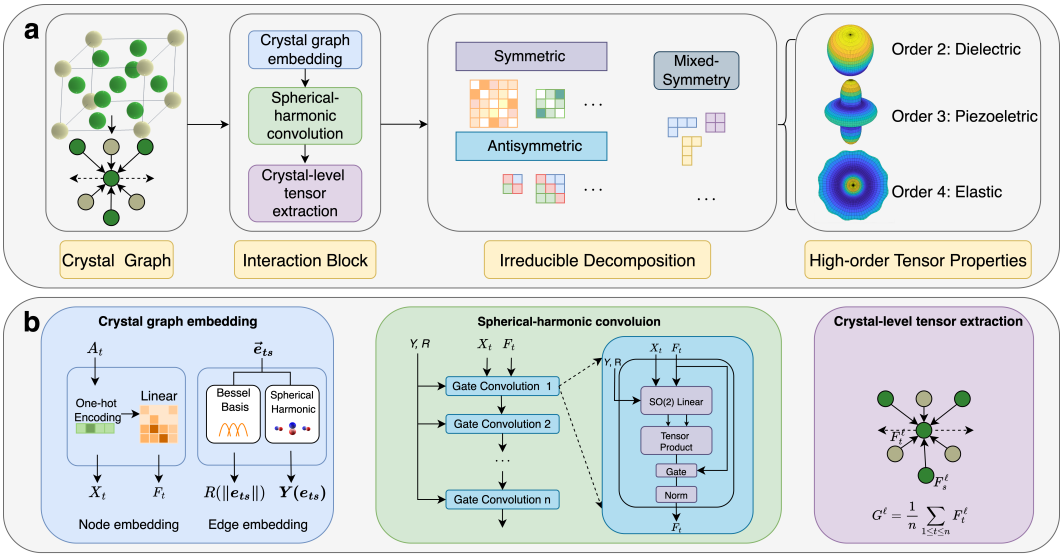


Figure 1: (a) Overview of IrredNet. A periodic crystal structure is first converted into a crystal graph and processed by several Equivariant Interaction Blocks to obtain SO(3)-equivariant node features. These features are then fed into the Irreducible Decomposition Block, where they are expanded on a spherical-harmonic basis, decomposed into irreducible components indexed by angular momentum ℓ, and filtered according to the target tensor symmetries. Different colors indicate groups of spherical harmonics forming bases of irreducible components associated with different parts (dimensionalities) of the high-order tensor property. The remaining symmetry-adapted components are finally combined by the tensor head to reconstruct the physical tensor. (b) Internal structure of a single Equivariant Interaction Block used in (a). It consists of a crystal graph embedding module (initializing node and edge features from atomic species and interatomic geometry), a spherical-harmonic convolution module that updates features in an SO(3)-equivariant manner, and a crystal-level tensor extraction module that aggregates atomic features into crystal-level representations passed to the irreducible decomposition stage in (a).

To efficiently capture high-order features (especially for $\ell > 0$) , we employ a spherical-harmonic convolution module that operates through several equivariant message-passing layers. Specifically, we replace the traditional $\mathrm{SO}(3)$ convolutions with Equivariant Spherical Channel Network [48, 49]. Instead of performing full $\mathrm{SO}(3)$ convolutions, this framework reformulates the required tensor products as $\mathrm{SO}(2)$ linear operations, greatly reducing the computation cost. In a standard $\mathrm{SO}(3)$ convolution, one takes the irreducible representation features $\psi_{m_i}^{\ell_i}$ and spherical-harmonic projections $Y_{m_f}^{\ell_f}(\hat{e}_{ts})$ and computes their tensor product. By applying a rotation matrix D_{ts} such that $D_{ts}\hat{r}_{ts} = \hat{\mathbf{z}}$, the spherical harmonics are evaluated on the canonical axis. In this frame, $Y_{m_f}^{\ell_f}(\hat{\mathbf{z}})$ is nonzero only for $m_f = 0$, while ℓ_f can still take multiple values. Hence, only the $m_f = 0$ terms remain in the tensor product, yielding couplings of the form $C_{(\ell_i,m_i),(\ell_f,0)}^{(\ell_o,m_o)}$. By the Clebsch–Gordan selection rule, these terms are nonzero only when $m_o = m_i$ in the complex spherical-harmonic basis. This converts the original $\mathrm{SO}(3)$ convolution into lower-cost $\mathrm{SO}(2)$ operations on rotated tensors, reducing the scaling from $O(L^6)$ to $O(L^3)$, where L is the maximum degree.

The spherical-harmonic convolution module stacks several gated convolution layers to efficiently propagate high-order feature representations. These gates filter out redundant scalar information, allowing the network to concentrate on more intricate, high-order interactions. Moreover, by applying an iterative training strategy at every layer, we further enhance the module's ability to learn and generalize complex, hierarchical features. Finally, after several equivariant message-passing layers are applied to extract high-order node features, the crystal-level features are aggregated from the node-level embeddings.

Irreducible Decomposition Block. Based on the theory and computational procedure for the irreducible decomposition of high-order tensors under specific symmetry constraints, we can derive

Table 1: Dataset statistics of high-order tensor properties.

Dataset	# Samples	Fnorm	Unit	Order
dielectric	4713	14.7±18.2	Unitless	2
piezoelectric	4998	0.43±3.09	C/m^2	3
elastic	14200	327±249	GPa	4

the basis for each irreducible component of the high-order tensor, as described before. A detailed implementation process for different high-order tensors is provided in Appendix B.

4 Experiments

4.1 Experiment Settings

Crystal High-Order Tensor Dataset. In our study, we thoroughly evaluate our model IrredNet on the JARVIS-DFT dataset [35], which includes crystal property tensors of three different orders. Following the preprocessing protocol established by previous work [25], we apply the same normalization procedure to ensure consistency. The dataset comprises dielectric, piezoelectric, and elastic tensors. We further summarize the physicochemical and geometric characteristics of these tensors in the following Table 1. Additionally, the dielectric property is represented by a 3×3 matrix $\boldsymbol{\epsilon}$ that satisfies the symmetry constraint $\boldsymbol{\epsilon}_{ij} = \boldsymbol{\epsilon}_{ji}$. The piezoelectric tensor $\boldsymbol{e}$ is a third-order tensor of size $3 \times 3 \times 3$, subject to the symmetry constraint $\boldsymbol{e}_{ijk} = \boldsymbol{e}_{ikj}$. The elastic tensor is a fourth-order tensor of dimension $3 \times 3 \times 3 \times 3$, governed by more complex symmetry conditions, including major symmetry ($C_{ijkl} = C_{klij}$) and minor symmetry ($C_{ijkl} = C_{jikl} = C_{ijlk}$). To provide a more compact and interpretable representation, the elastic tensor is often expressed using Voigt notation, which converts the tensor into a symmetric 6×6 matrix form containing 21 independent components. A detailed introduction is provided in the Appendix B.

Baseline Methods. We benchmark several state-of-the-art methods in the field, including MEGNet [15], which is an invariant model commonly used for predicting tensorial data in previous work [50]. We also compare it to ETGNN [51], an equivariant model that excels at capturing tensorial features. Additionally, we evaluate GMTNet [25], another approach that uses space group-related rotational matrices to enforce symmetry and predict high-order tensor properties. We additionally include GoeCTP [52], a canonicalization-based framework for crystal tensor prediction. We independently implement GoeCTP based on the published method and the authors' GitHub repository. All methods are evaluated on the JARVIS-DFT datasets using the metrics defined below. Implementation details are provided in Appendix D.

Evaluation Metrics. To assess the quality of high-order crystal tensor predictions, we adopt several metrics from previous work [25, 53, 2]: (1) **Zero-element accuracy**: This metric evaluates the model's ability to accurately identify zero-valued entries in high-order tensors that arise due to symmetry constraints. (2) **Equality accuracy**: This metric assesses the model's capability to recognize mutually dependent tensor components that are constrained to have identical values due to intrinsic tensor symmetries. (3) **Frobenius norm (Fnorm) distance**: The Frobenius norm distance quantifies the discrepancy between the predicted and ground-truth tensors by computing the square root of the sum of the squared differences of all corresponding elements. It is defined as:$\|y_{\text{pred}} - y_{\text{true}}\|_F = \sqrt{\sum_i (y_{\text{pred},i} - y_{\text{true},i})^2}$. (4) **Error within threshold (EwT)**: EwT is used to evaluate the proportion of high-quality predictions that fall within specified relative error thresholds (e.g., 25%, 10%, 5%). It is defined based on the ratio between the Frobenius norm of the error and that of the ground-truth tensor: $\|y_{\text{pred}} - y_{\text{label}}\|_F \,/\, \|y_{\text{label}}\|_F$, where y_{pred} and y_{label} denote the predic ted and true tensor values, respectively. A prediction is considered high-quality if its EwT value falls below a specified threshold.

Following previous work [25], we re-implement MEGNet, ETGNN, and GMTNet as our baselines, all with their default hyperparameters. Each dataset is randomly split into training, validation, and test subsets in an 80:10:10 ratio. We train our model using the Huber loss, the AdamW optimizer, and a polynomial decay schedule with an initial learning rate of 10^{-5}. For the tensor datasets, we

set the maximum spin quantum number ℓ to 3, 3, and 4 for the dielectric, piezoelectric, and elastic tensors, respectively. Detailed settings are listed in Appendix D and I.

4.2 Symmetry Compliance of High-Order Tensor Predictions

To assess our model's ability to capture crystal symmetry across different crystal systems, we employ two accuracy metrics mentioned before: (1) **Zero-element accuracy** and (2) **Equality accuracy**. These two metrics assess (i) the proportion of tensor entries that are truly zero and correctly predicted as zero, and (ii) the proportion of nonzero entries whose predicted values satisfy the required equivariance constraints. Following the GMTNet setting, we evaluate symmetry compliance using two success-rate metrics: zero-element accuracy with a tolerance of 10^{-5} and equality accuracy with a tolerance of 10^{-4}. Specifically, we assess symmetry preservation on the fourth-order elastic tensor, one of the most challenging cases in Table 2. Our model achieves perfect zero-element and equality accuracies and strictly preserves all symmetry constraints. In contrast, MEGNet and ETGNN show substantially lower zero-element and equality accuracies. Our independent implementation of GoeCTP obtains 0% on both metrics across these crystal systems under the same tolerances. GMTNet, which leverages space group information to enforce symmetry, also achieves 100% zero-element and equality accuracy across all six crystal systems after its full symmetry-enforcement procedure, establishing it as a strong symmetry-preserving baseline. IrredNet matches this compliance through its symmetry-constrained irreducible parameterization.

Table 2: Zero–element and equality accuracy of predicted rank-4 elastic tensors on the JARVIS-DFT test set, broken down by crystal system. For each crystal system, Zero reports the percentage of tensor components that are required to be zero by symmetry and are correctly predicted as zero, and Equality reports the percentage of pairs of components that are required to be equal and are predicted to satisfy this equality. Higher values indicate better symmetry compliance. GoeCTP denotes our independent implementation.

	MEGNet		ETGNN		GoeCTP		GMTNet		IrredNet	
Crystal System	Zero	Equality	Zero	Equality	Zero	Equality	Zero	Equality	Zero	Equality
Cubic	0%	0%	13.9%	23.1%	0%	0%	**100%**	**100%**	**100%**	**100%**
Tetragonal	0%	0%	5.3%	6.3%	0%	0%	**100%**	**100%**	**100%**	**100%**
Hexagonal	0%	0%	2.5%	7.2%	0%	0%	**100%**	**100%**	**100%**	**100%**
Trigonal	0%	0%	1.6%	2.8%	0%	0%	**100%**	**100%**	**100%**	**100%**
Orthorhombic	0%	0%	0%	0%	0%	0%	**100%**	**100%**	**100%**	**100%**
Monoclinic	0%	0%	6.7%	0%	0%	0%	**100%**	**100%**	**100%**	**100%**

4.3 Accuracy of High-Order Tensor Predictions

Beyond symmetry compliance, we benchmark our model's accuracy against several state-of-the-art baselines across three tensor datasets. Table 3 presents the prediction results of various models evaluated using the Fnorm metric and the EwT metrics at the 25%, 10%, and 5% thresholds. Our model delivers three key advantages in high-order tensor prediction: it consistently reduces reconstruction error (Fnorm) relative to GMTNet from 3.50 to 3.23 on dielectric, 0.37 to 0.29 on piezoelectric, and 68.12 to 64.23 on the elastic tensor; it achieves the highest EwT scores at the strict 5% threshold: 29.8% (dielectric), 49.2% (piezoelectric), and 12.1% (elastic), demonstrating superior capture of fine-grained, symmetry-driven structure. For elastic tensors, the EwT score at the 5% threshold increases from 7.7% for GMTNet to 12.1% for IrredNet. Meanwhile, the total runtime remains competitive across all tasks, indicating that the proposed irreducible-parameterization strategy improves physical consistency and accuracy without introducing prohibitive computational overhead. GoeCTP is also competitive on elastic tensors, with an Fnorm of 66.36 and an EwT score of 12.0% at the 5% threshold, while IrredNet achieves 64.23 and 12.1%, respectively.

4.4 Ablation Studies

In this section, we evaluate the importance of symmetry and individual model components. To this end, we conduct ablation experiments on the dielectric dataset using several metrics, with results summarized in Table 4. For instance, on the Fnorm metric, incorporating irreducible decomposition delivers the first notable gain, reducing Fnorm from 4.17 (without decomposition) to 3.61. Notably,

Table 3: Comparison of high-order tensor prediction on the JARVIS-DFT dielectric, piezoelectric (Piezo), and elastic benchmarks. The table reports results for MEGNet, ETGNN, GoeCTP, GMTNet, and the proposed IrredNet. For each method and dataset we show the Fnorm error (↓) of the predicted tensors and EwT metrics (↑), which give the percentage of test samples whose relative Fnorm error is below 25%, 10%, and 5%, respectively. The last column (Total Time (s) (↓)) reports the total wall-clock time in seconds for training and evaluation on the dataset under identical hardware and batch-size settings. GoeCTP results are from our independent reproduction.

Property	Method	Fnorm ↓	EwT 25% ↑	EwT 10% ↑	EwT 5% ↑	Total time (s) ↓
Dielectric	MEGNet	4.16	74.9%	38.9%	19.1%	1005
	ETGNN	3.92	81.3%	41.6%	23.8%	1510
	GoeCTP	4.06	81.5%	47.8%	23.6%	812
	GMTNet	3.50	84.5%	**57.1%**	27.8%	881
	IrredNet	**3.23**	**85.8%**	53.2%	**29.8%**	**798**
Piezoelectric	MEGNet	0.49	46.1%	44.3%	41.1%	381
	ETGNN	0.40	47.1%	46.9%	46.8%	447
	GoeCTP	0.45	0.2%	0.1%	0.0%	376
	GMTNet	0.37	49.1%	46.3%	45.7%	**286**
	IrredNet	**0.29**	**49.3%**	**49.3%**	**49.2%**	288
Elastic	MEGNet	132.21	22.8%	3.0%	0.3%	14666
	ETGNN	90.84	42.0%	12.8%	1.2%	34157
	GoeCTP	66.36	65.7%	26.3%	12.0%	15131
	GMTNet	68.12	65.0%	21.6%	7.7%	$>$ 67000
	IrredNet	**64.23**	**66.5%**	**26.5%**	**12.1%**	**14121**

the zero-element and equality accuracies improve significantly with this addition. However, when neither irreducible decomposition nor symmetry constraints are applied, the model struggles to capture symmetry accurately, leading to poor predictions for these elements. Only when both irreducible decomposition and high-order tensor-specific symmetry constraints are applied can the model fully guarantee symmetry compliance, further reducing the Fnorm to 3.33. This demonstrates the effectiveness of our symmetry-constrained irreducible decomposition approach. Importantly, the comparison between the second and third rows shows that irreducible decomposition alone improves accuracy but does not fully enforce the desired tensor symmetries, whereas the symmetry-constrained basis is necessary to achieve perfect zero-element and equality accuracies. Finally, experiments with the gate module and norm prediction head show that both contribute positively to prediction accuracy, with the best performance achieved when all components are used, resulting in a final Fnorm of 3.23. These results indicate that the performance gain comes from both the exact symmetry-preserving output parameterization and the architectural components that improve feature learning.

Table 4: Ablation studies on the dielectric tensor dataset. Each row toggles components of IrredNet: Irred. Decomp. (irreducible-basis output parameterization), Sym. Cons. (exact symmetry constraints), Gate (gated equivariant convolutions), and Norm (tensor-norm prediction head). We report Fnorm (↓), EwT (↑), and zero/equality accuracy (↑), showing that combining irreducible decomposition with symmetry constraints and gating yields the best overall performance and perfect symmetry compliance.

Irred. Decomp.	Sym. Cons.	Gate	Norm	Fnorm (↓)	EwT 25% (↑)	EwT 10% (↑)	EwT 5% (↑)	Zero (↑)	Equality (↑)
✗	✗	✓	✓	4.17	72.1%	40.1%	14.9%	0%	0%
✓	✗	✓	✓	3.61	81.4%	42.4%	24.1%	84.9%	89.7%
✓	✓	✗	✗	3.44	78.3%	43.3%	26.0%	**100%**	**100%**
✓	✓	✓	✗	3.33	83.9%	51.1%	29.7%	**100%**	**100%**
✓	✓	✓	✓	**3.23**	**85.8%**	**53.2%**	**29.8%**	**100%**	**100%**

4.5 Analysis of High-Order Tensor Predictions

To further evaluate whether the predicted elastic tensors are physically meaningful beyond element-wise tensor errors, we analyze several direction-dependent mechanical properties derived from the

elastic tensor. Specifically, we consider Young's modulus, shear modulus, linear compressibility, and Poisson's ratio, which are widely used quantities for characterizing the anisotropic mechanical response of crystalline materials. These properties are nonlinear functions of the elastic stiffness tensor and its corresponding compliance tensor, and therefore provide a more stringent test of whether a predicted tensor preserves the correct spatial anisotropy and symmetry patterns.

Using HgIBr as a representative case study, we take the DFT reference elastic tensor as the ground truth and use the ELATE tool [32] to derive four directional mechanical response surfaces from the ground-truth tensor, the GMTNet prediction, and the prediction of IrredNet. The resulting directional profiles are visualized in Figure 2. As shown in the figure, GMTNet captures the overall magnitude of some properties but introduces noticeable angular distortions and spurious anisotropic patterns, especially in the shear modulus and Poisson's ratio. In contrast, IrredNet produces directional response surfaces that are more consistent with the DFT reference in both shape and orientation. For example, the predicted Young's modulus preserves the nearly ellipsoidal anisotropy of the ground truth, while the shear modulus and Poisson's ratio more faithfully reproduce the angular locations of high- and low-response regions. This qualitative comparison indicates that IrredNet outperforms GMTNet in reproducing the DFT-derived directional mechanical responses of HgIBr, with closer agreement in profile shape and angular variation, particularly for shear modulus and Poisson's ratio.

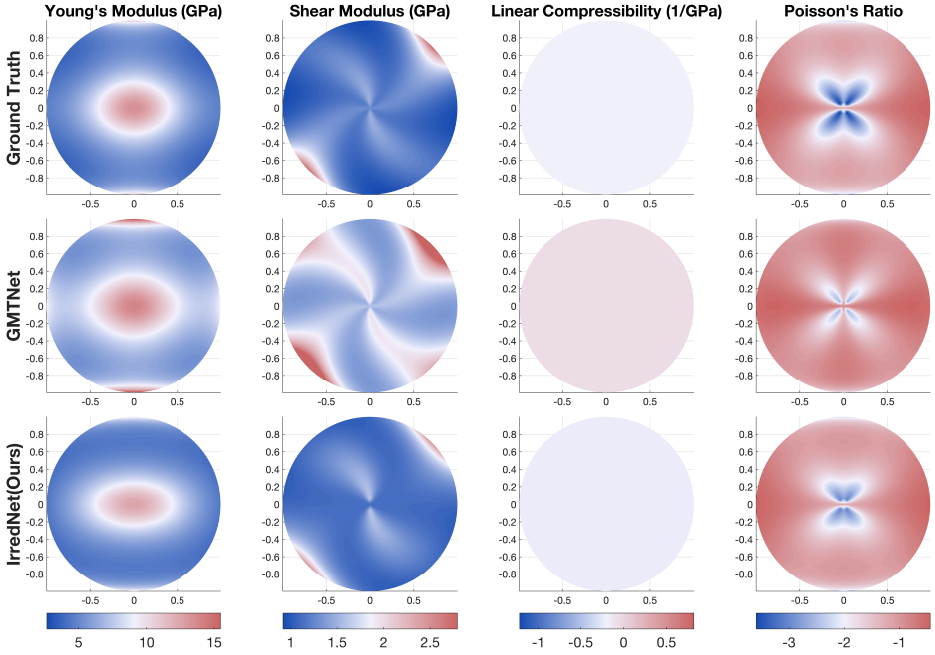


Figure 2: Directional elastic properties for a representative crystal. Columns show polar plots of Young's modulus (GPa), shear modulus (GPa), linear compressibility (1/GPa), and Poisson's ratio as functions of loading direction; rows correspond to ground truth (top), GMTNet (middle), and IrredNet (bottom). IrredNet more closely matches the ground-truth anisotropy and avoids the spurious artifacts visible in GMTNet.

5 Conclusion

We address the problem of enforcing strict symmetry in high-order tensor predictions for crystals. We introduce IrredNet, which guarantees exact compliance by predicting symmetry-constrained irreducible components: we prove that any high-order tensor admits an SO(3) irreducible decomposition and provide an algorithm to enumerate its basis tensors, ensuring symmetry by construction. Experiments confirm exact symmetry preservation and high accuracy on complex tensors and related applications.

Acknowledgments and Disclosure of Funding

Bo Han was supported by the RGC Research Fellow Scheme (No. RFS2627-2S02), the RGC Young Collaborative Research Grant (No. C2005-24Y), the RGC General Research Fund (Nos. 12202026 and 12200725), the NSFC Major Research Plan (No. 92570109), the NSFC General Program (No. 62376235), the HKBU Faculty Niche Research Areas (No. RC-FNRA-IG/25-26/SCI/05), and the HKBU CSD Departmental Incentive Scheme. Tongliang Liu was partially supported by the Australian Research Council through projects DP260102466 and FT220100318. Yi Chang was supported by the Ministry of Science and Technology of China (No. 2023YFF0905400), the National Natural Science Foundation of China (No. U2341229), and the New Cornerstone Science Foundation through the XPLORER PRIZE. We extend our special thanks to Mr. Chenhan Wang from Hyper.AI for generously providing the computing resources essential to this research.

The authors have no competing interests to declare.

References

[1] David Lovett. *Tensor properties of crystals*. CRC Press, 2018.

[2] Ioannis Petousis, Wei Chen, Geoffroy Hautier, Thomas Graf, Thomas D Schladt, Kristin A Persson, and Fritz B Prinz. Benchmarking density functional perturbation theory to enable high-throughput screening of materials for dielectric constant and refractive index. *Physical Review B*, 93(11):115151, 2016.

[3] Jun Xiao, Ying Wang, Hua Wang, Chaitanya Das Pemmaraju, Siming Wang, Peter Muscher, Edbert J Sie, Clara M Nyby, Thomas P Devereaux, Xiaofeng Qian, et al. Berry curvature memory through electrically driven stacking transitions. *Nature Physics*, 16(10):1028–1034, 2020.

[4] Eva G Noya, Carlos Vega, Jonathan PK Doye, and Ard A Louis. The stability of a crystal with diamond structure for patchy particles with tetrahedral symmetry. *The Journal of chemical physics*, 132(23), 2010.

[5] Avni Jain, Jeffrey R Errington, and Thomas M Truskett. Dimensionality and design of isotropic interactions that stabilize honeycomb, square, simple cubic, and diamond lattices. *Physical Review X*, 4(3):031049, 2014.

[6] Dodd Gray, Adam McCaughan, and Bhaskar Mookerji. Crystal structure of graphite, graphene and silicon. *Physics for Solid State Applications*, 6:730, 2009.

[7] Fanhao Jia, Shaowen Xu, Shunbo Hu, Jianguo Chen, Yongchen Wang, Yuan Li, Wei Ren, and Jinrong Cheng. Understanding the intrinsic piezoelectric anisotropy of tetragonal abo3 perovskites through a high-throughput study. *npj Computational Materials*, 11(1):6, 2025.

[8] Sina Hashemizadeh, Alberto Biancoli, and Dragan Damjanovic. Symmetry breaking in hexagonal and cubic polymorphs of batio3. *Journal of Applied Physics*, 119(9), 2016.

[9] Walter Kohn and Lu Jeu Sham. Self-consistent equations including exchange and correlation effects. *Physical review*, 140(4A):A1133, 1965.

[10] Walter Kohn, Axel D Becke, and Robert G Parr. Density functional theory of electronic structure. *The journal of physical chemistry*, 100(31):12974–12980, 1996.

[11] Alexander Dunn, Qi Wang, Alex Ganose, Daniel Dopp, and Anubhav Jain. Benchmarking materials property prediction methods: the matbench test set and automatminer reference algorithm. *npj Computational Materials*, 6(1):138, 2020.

[12] Lukáš Kỳvala, Pablo Montero de Hijes, and Christoph Dellago. Unsupervised identification of crystal defects from atomistic potential descriptors. *npj Computational Materials*, 11(1):50, 2025.

[13] Lauri Himanen, Marc OJ Jäger, Eiaki V Morooka, Filippo Federici Canova, Yashasvi S Ranawat, David Z Gao, Patrick Rinke, and Adam S Foster. Dscribe: Library of descriptors for machine learning in materials science. *Computer Physics Communications*, 247:106949, 2020.

[14] Jingjin He, Junjie Li, Chuanbao Liu, Changxin Wang, Yan Zhang, Cheng Wen, Dezhen Xue, Jiangli Cao, Yanjing Su, Lijie Qiao, et al. Machine learning identified materials descriptors for ferroelectricity. *Acta Materialia*, 209:116815, 2021.

[15] Chi Chen, Weike Ye, Yunxing Zuo, Chen Zheng, and Shyue Ping Ong. Graph networks as a universal machine learning framework for molecules and crystals. *Chemistry of Materials*, 31(9):3564–3572, 2019.

[16] Tian Xie and Jeffrey C Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Physical Review Letters*, 120(14):145301, 2018.

[17] Chi Chen and Shyue Ping Ong. A universal graph deep learning interatomic potential for the periodic table. *Nature Computational Science*, 2(11):718–728, 2022.

[18] Yeonghun Kang, Hyunsoo Park, Berend Smit, and Jihan Kim. A multi-modal pre-training transformer for universal transfer learning in metal–organic frameworks. *Nature Machine Intelligence*, 5(3):309–318, 2023.

[19] Zheyong Fan, Yanzhou Wang, Penghua Ying, Keke Song, Junjie Wang, Yong Wang, Zezhu Zeng, Ke Xu, Eric Lindgren, J Magnus Rahm, et al. Gpumd: A package for constructing accurate machine-learned potentials and performing highly efficient atomistic simulations. *The Journal of Chemical Physics*, 157(11), 2022.

[20] Keqiang Yan, Yi Liu, Yuchao Lin, and Shuiwang Ji. Periodic graph transformers for crystal material property prediction. *Advances in Neural Information Processing Systems*, 35:15066–15080, 2022.

[21] Kamal Choudhary and Brian DeCost. Atomistic line graph neural network for improved materials property predictions. *npj Computational Materials*, 7(1):185, 2021.

[22] Kristof Schütt, Oliver Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial properties and molecular spectra. *International Conference on Machine Learning*, pages 9377–9388, 2021.

[23] Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13:2453, 2022.

[24] Ilyes Batatia, David Peter Kovacs, Gregor Simm, Christoph Ortner, and Gábor Csányi. Mace: Higher order equivariant message passing neural networks for fast and accurate force fields. *Advances in Neural Information Processing Systems*, 35:11423–11436, 2022.

[25] Keqiang Yan, Alexandra Saxton, Xiaofeng Qian, Xiaoning Qian, and Shuiwang Ji. A space group symmetry informed network for o (3) equivariant crystal tensor prediction. *arXiv preprint arXiv:2406.12888*, 2024.

[26] Ivan Grega, Ilyes Batatia, Gábor Csányi, Sri Karlapati, and Vikram S Deshpande. Energy-conserving equivariant gnn for elasticity of lattice architected metamaterials. *arXiv preprint arXiv:2401.16914*, 2024.

[27] Kailai Xu, Daniel Z Huang, and Eric Darve. Learning constitutive relations using symmetric positive definite neural networks. *Journal of Computational Physics*, 428:110072, 2021.

[28] Sikko Van't Sant, Prakash Thakolkaran, Jonàs Martínez, and Siddhant Kumar. Inverse-designed growth-based cellular metamaterials. *Mechanics of Materials*, 182:104668, 2023.

[29] Andrea Grisafi, David M Wilkins, Gábor Csányi, and Michele Ceriotti. Symmetry-adapted machine learning for tensorial properties of atomistic systems. *Physical review letters*, 120(3):036002, 2018.

[30] Mingjian Wen, Matthew K Horton, Jason M Munro, Patrick Huck, and Kristin A Persson. An equivariant graph neural network for the elasticity tensors of all seven crystal systems. *Digital Discovery*, 3(5):869–882, 2024.

[31] Charles F Jekel, Kenneth E Swartz, Daniel A White, Daniel A Tortorelli, and Seth E Watts. Neural network layers for prediction of positive definite elastic stiffness tensors. *arXiv preprint arXiv:2203.13938*, 2022.

[32] Romain Gaillac, Pluton Pullumbi, and François-Xavier Coudert. Elate: an open-source online application for analysis and visualization of elastic tensors. *Journal of Physics: Condensed Matter*, 28(27):275201, 2016.

[33] Jules Thomas Browaeys and Sébastien Chevrot. Decomposition of the elastic tensor and geophysical applications. *Geophysical Journal International*, 159(2):667–678, 2004.

[34] Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, et al. Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL materials*, 1(1), 2013.

[35] Kamal Choudhary, Kevin F Garrity, Andrew CE Reid, Brian DeCost, Adam J Biacchi, Angela R Hight Walker, Zachary Trautt, Jason Hattrick-Simpers, A Gilad Kusne, Andrea Centrone, et al. The joint automated repository for various integrated simulations (jarvis) for data-driven materials design. *npj computational materials*, 6(1):173, 2020.

[36] Brandon Anderson, Truong Son Hy, and Risi Kondor. Cormorant: Covariant molecular neural networks. *Advances in neural information processing systems*, 32, 2019.

[37] Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications*, 13(1):2453, 2022.

[38] Bowen Deng, Peichen Zhong, KyuJung Jun, Janosh Riebesell, Kevin Han, Christopher J Bartel, and Gerbrand Ceder. Chgnet as a pretrained universal neural network potential for charge-informed atomistic modelling. *Nature Machine Intelligence*, 5(9):1031–1041, 2023.

[39] Vıctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks. In *International conference on machine learning*, pages 9323–9332. PMLR, 2021.

[40] Jerome Brachat, Pierre Comon, Bernard Mourrain, and Elias Tsigaridas. Symmetric tensor decomposition. *Linear Algebra and its Applications*, 433(11-12):1851–1872, 2010.

[41] AA Gusev, Vladimir P Gerdt, Sergey I Vinitsky, VL Derbov, A Góźdź, and A Pędrak. Symbolic algorithm for generating irreducible bases of point groups in the space of so (3) group. In *International Workshop on Computer Algebra in Scientific Computing*, pages 166–181. Springer, 2015.

[42] Miguel S Costa and Tobias Hansen. Conformal correlators of mixed-symmetry tensors. *Journal of High Energy Physics*, 2015(2):1–44, 2015.

[43] Igor R Klebanov and Grigory Tarnopolsky. On large n limit of symmetric traceless tensor models. *Journal of High Energy Physics*, 2017(10):1–16, 2017.

[44] Atsushi Higuchi. Symmetric tensor spherical harmonics on the n-sphere and their application to the de sitter group so (n, 1). *Journal of mathematical physics*, 28(7):1553–1566, 1987.

[45] Wu-Ki Tung. *Group theory in physics*, volume 1. World Scientific, 1985.

[46] G Ya Lyubarskii. *The application of group theory in physics*. Elsevier, 2013.

[47] Keqiang Yan, Chengyue Fu, Xiaofeng Qian, Xiaoning Qian, and Shuiwang Ji. Complete and efficient graph transformers for crystal material property prediction. *International Conference on Learning Representations*, 2024.

[48] Saro Passaro and C Lawrence Zitnick. Reducing so (3) convolutions to so (2) for efficient equivariant gnns. In *International conference on machine learning*, pages 27420–27438. PMLR, 2023.

[49] Yi-Lun Liao, Brandon Wood, Abhishek Das, and Tess Smidt. Equiformerv2: Improved equivariant transformer for scaling to higher-degree representations. *arXiv preprint arXiv:2306.12059*, 2023.

[50] Kazuki Morita, Daniel W Davies, Keith T Butler, and Aron Walsh. Modeling the dielectric constants of crystals using machine learning. *The Journal of Chemical Physics*, 153(2), 2020.

[51] Yang Zhong, Hongyu Yu, Xingao Gong, and Hongjun Xiang. Edge-based tensor prediction via graph neural networks. *arXiv preprint arXiv:2201.05770*, 2022.

[52] Haowei Hua, Jingwen Yang, Wanyu Lin, and Pan Zhou. Revisiting the canonicalization for fast and accurate crystal tensor property prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 417–425, 2026.

[53] Haowei Hua, Wanyu Lin, and Jingwen Yang. Fast crystal tensor property prediction: A general o (3)-equivariant framework based on polar decomposition. *arXiv preprint arXiv:2410.02372*, 2024.

[54] Maxim Nazarov. Young's symmetrizers for projective representations of the symmetric group. *advances in mathematics*, 127(2):190–257, 1997.

[55] Andrea Brini and Antonio GB Teolis. Young—capelli symmetrizers in superalgebras. *Proceedings of the National Academy of Sciences*, 86(3):775–778, 1989.

[56] Arieh Iserles, Hans Z Munthe-Kaas, Syvert P Nørsett, and Antonella Zanna. Lie-group methods. *Acta numerica*, 9:215–365, 2000.

[57] Claude Chevalley. *Theory of Lie groups*. Courier Dover Publications, 2018.

[58] Jonathan Brundan and Alexander Kleshchev. Schur–weyl duality for higher levels. *Selecta Mathematica*, 14(1):1–57, 2008.

[59] Stephen Doty and Jun Hu. Schur–weyl duality for orthogonal groups. *Proceedings of the London Mathematical Society*, 98(3):679–713, 2009.

[60] Theodore Chang and Tor Skjelbred. The topological schur lemma and related results. *Annals of Mathematics*, 100(2):307–321, 1974.

[61] Jean-Pierre Serre et al. *Linear representations of finite groups*, volume 42. Springer, 1977.

[62] Yakov Itin and Shulamit Reches. Decomposition of third-order constitutive tensors. *Mathematics and Mechanics of Solids*, 27(2):222–249, 2022.

[63] Yakov Itin. Irreducible matrix resolution of the elasticity tensor for symmetry systems. *arXiv preprint arXiv:1812.03367*, 2018.

[64] Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.

[65] Johannes Gasteiger, Shankari Giri, Johannes T Margraf, and Stephan Günnemann. Fast and uncertainty-aware directional message passing for non-equilibrium molecules. *arXiv preprint arXiv:2011.14115*, 2020.

[66] Stephen C Cowin. The relationship between the elasticity tensor and the fabric tensor. *Mechanics of materials*, 4(2):137–147, 1985.

A Theory for high-order tensor irreducible decomposition under SO(3) group

A.1 Proof for high-order tensor irreducible decomposition under $\mathrm{SO}(3)$ group

Proposition 1 Restated Let V be a three-dimensional Euclidean vector space. Any rank-k ($k > 1$) tensor $\mathcal{T}_{i_1\cdots i_k} \in V^{\otimes k}$ admits a unique decomposition (up to isomorphism) into irreducible representations (irreps) of $\mathrm{SO}(3)$ group labeled by integer spins ℓ as follows:

$$V^{\bigotimes k} \simeq \bigoplus_{\ell\in L_k} \left(\mathcal{S}_k^{(\ell)} \bigoplus \mathcal{A}_k^{(\ell)} \bigoplus \mathcal{M}_k^{(\ell)}\right),$$

where L_k is the set of all distinct integer spins ℓ that appear in the decomposition of $V^{\otimes k}$. $\mathcal{S}_k^{(\ell)}$, $\mathcal{A}_k^{(\ell)}$, and $\mathcal{M}_k^{(\ell)}$ denote the symmetric traceless, the antisymmetric, and the mixed-symmetry components (including symmetric with trace) transforming under spin-ℓ irreps, respectively.

Proof We aim to prove that any rank-k ($k > 1$) tensor $\mathcal{T}_{i_1\cdots i_k} \in V^{\otimes k}$ (where V is a three-dimensional Euclidean vector space) admits a unique decomposition (up to isomorphism) into irreducible representations (irreps) of the $\mathrm{SO}(3)$ group. These irreps are labeled by integer spins ℓ and can be categorized by their permutation symmetry properties.

Specifically, in **Step 1**, we establish the theoretical foundation for decomposing rank-k tensors under the $\mathrm{SO}(3)$ symmetry group by considering the irreducible representations (irreps) associated with $\mathrm{SO}(3)$ and their spin labels. **Step 2** focuses on the role of the symmetric group S_k and Schur-Weyl duality, which enables the decomposition of the tensor space into subspaces defined by both $\mathrm{SO}(3)$ and S_k. **Step 3** addresses the condition of irreducibility under $\mathrm{SO}(3)$ by separating traceless components and trace-derived lower-spin sectors while preserving the required symmetry. In **Step 4**, we classify the irreducible components based on their permutation symmetry, dividing them into fully symmetric traceless, fully antisymmetric, and mixed-symmetry subspaces. Finally, **Step 5** provides the explicit procedure for decomposing the tensor space, using Young symmetrizers [54, 55] and projection operations to obtain a complete irreducible decomposition, including trace-derived lower-spin sectors.

Step 1: Representation-Theoretic Foundations

The group $\mathrm{SO}(3)$ is a compact Lie group [56, 57], and its finite-dimensional irreducible representations (irreps) are characterized by an integer spin (often denoted ℓ), designated as $D(\ell)$ with dimension $2\ell + 1$. The three-dimensional Euclidean vector space V itself is the fundamental representation space for $\mathrm{SO}(3)$, corresponding to spin ℓ=1. The space of rank-k tensors, $V^{\otimes k}$, forms a representation of $\mathrm{SO}(3)$ that transforms under the action of $g \in \mathrm{SO}(3)$. For $k > 1$, this representation is generally reducible, and our goal is to decompose it into a direct sum of $\mathrm{SO}(3)$ irreps $D(\ell)$. By the general theory of compact groups, this decomposition is unique up to isomorphism and the ordering of terms.

Step 2: Role of the Symmetric Group S_k and Schur-Weyl Duality

The symmetric group S_k acts on $V^{\otimes k}$ by permuting the tensor indices. The actions of the $\mathrm{SO}(3)$ group (acting on the values at each index position) and the S_k group (acting on the index positions themselves) commute. This commutativity allows for the use of Schur-Weyl duality [58, 59], which decomposes $V^{\otimes k}$ into subspaces associated with the joint representations of S_k and $\mathrm{SO}(3)$. Specifically,

$$V^{\otimes k} \simeq \bigoplus_{\lambda} \mathcal{R}_\lambda \otimes \mathcal{P}_\lambda,$$

where λ are Young diagrams for partitions of k (with at most dim(V)=3 rows), $\mathcal{P}_\lambda$ is the irrep of S_k, and $\mathcal{R}_\lambda$ carries a representation of $\mathrm{SO}(3)$. Importantly, $\mathcal{R}_\lambda$ itself is generally reducible under $\mathrm{SO}(3)$ and requires further decomposition.

Step 3: Irreducibility under $\mathrm{SO}(3)$ and Tracelessness

An $\mathrm{SO}(3)$ spin-ℓ irrep can be realized by symmetric traceless tensors of rank ℓ, but its embedding in $V^{\otimes k}$ need not be traceless. For example, $\mathrm{span}\{\delta_{ij}\}$ is a spin-0 irrep with nonzero trace. Contraction with δ_{ij} is $\mathrm{SO}(3)$-equivariant [43]; therefore, a complete decomposition retains the trace-derived lower-spin sectors alongside the traceless components.

Step 4: Classification of Irreducible Components based on Permutation Symmetry

The irreducible subspaces $D^{(\ell)}$ obtained from $V^{\otimes k}$ are classified by the permutation symmetry (Young diagram λ) of the tensors that carry them, together with their trace-derived lower-spin sectors:

- **Fully Symmetric** ($\mathcal{S}_k^{(\ell)}$): These components arise from tensors of Young diagram $\lambda = [k]$ (a single row, i.e. fully symmetric tensors). By convention $\mathcal{S}_k^{(\ell)}$ denotes the highest–spin component, namely $D^{(k)}$ (spin $\ell = k$), a fully symmetric traceless tensor of rank k.
- **Fully Antisymmetric** ($\mathcal{A}_k^{(\ell)}$): These components arise from tensors of Young diagram $\lambda = [1^k]$ (a single column, i.e. fully antisymmetric tensors). In $\dim(V) = 3$ they are nonzero only for $k \leq 3$, yielding for example $D^{(1)}$ (axial vector) when $k = 2$ and $D^{(0)}$ (pseudoscalar) when $k = 3$.
- **Mixed Symmetry** ($\mathcal{M}_k^{(\ell)}$): This category includes all other $\mathrm{SO}(3)$ irreps $D^{(\ell)}$ found in $V^{\otimes k}$. They come from Young diagrams that are neither a single row without trace nor a single column (e.g. $\lambda = [k-1, 1]$ for $k \geq 3$), after resolving traceless and trace-derived sectors. Consistent with the proposition's convention on traces of symmetric tensors, $\mathcal{M}_k^{(\ell)}$ also contains the lower–spin irreps $D^{(k-2j)}$ ($j > 0$) originating from the trace of the $[k]$ diagram once the leading $\mathcal{S}_k^{(k)}$ component is removed.

Step 5: Explicit Decomposition, Projections, and Uniqueness

The decomposition of the $V^{\otimes k}$ space into subspaces of specific permutation symmetry (corresponding to Young diagrams λ) can be achieved using Young symmetrizers $\mathbb{P}_\lambda$, which are linear combinations of permutation operators from S_k; examples include the full symmetrizer $\mathbb{P}_S$ and antisymmetrizer $\mathbb{P}_A$. Within each such symmetry-defined subspace $\mathcal{R}_\lambda$, further projections separate traceless components and trace-derived lower-spin sectors, retaining both in the full decomposition.

The resulting components supply bases for the SO(3) irreps $D^{(\ell)}$. The specific spin value ℓ for each irreducible subspace can be confirmed by the eigenvalue equation for the total angular momentum operator squared, $J^2\,\mathcal{T}^{(\ell)} = \ell(\ell+1)\,\mathcal{T}^{(\ell)}$. The decomposition of $V^{\otimes k}$ into a direct sum of SO(3) irreps $D^{(\ell)}$ is unique up to isomorphism. The categorization into $\mathcal{S}_k^{(\ell)}$, $\mathcal{A}_k^{(\ell)}$, and $\mathcal{M}_k^{(\ell)}$ (adhering to the specific definitions where $\mathcal{S}_k^{(\ell)}$ is primarily the $\ell = k$ part from symmetric tensors) provides a complete partitioning of these irreps based on their permutation-symmetry origins and trace-derived sectors.

Proposition 2 Restated Let $\mathcal{T}$ be a rank-k tensor defined over a three-dimensional Euclidean space. Under decomposition into irreducible representations (irreps) of $\mathrm{SO}(3)$ group: When $\mathcal{T}$ satisfies additional index symmetries (e.g., major/minor symmetry, tracelessness, or antisymmetry), the set of allowed irreps forms a subset of the unconstrained case: $\mathrm{Irrep}(\mathcal{T}_{\text{sym}}) \subseteq \mathrm{Irrep}(\mathcal{T}_{\text{unconstrained}})$.

Proof We prove the Proposition 2 based on the Schur's Lemma [60, 61].

Schur's lemma Let (ρ, V) and (ρ', W) be two real irreducible representations of $G = \mathrm{SO}(3)$. If a linear map $\phi : V \to W$ satisfies:

$$\phi \circ \rho(g) = \rho'(g) \circ \phi \quad (\forall g \in G),$$

Then if V and W are inequivalent, then $\phi = 0$. If $V \simeq W$, then after identifying the two representation spaces, ϕ is a scalar multiple of the identity ($\phi = \lambda I$).

Let $\mathcal{T}$ be a rank-k tensor in 3D Euclidean space. We rigorously prove that imposing symmetry constraints reduces the set of admissible irreducible representations (irreps) of $\mathrm{SO}(3)$ in its decomposition by Schur's lemma.

For a tensor space without symmetry constraints $\mathcal{T} = V^{\otimes k}$, its $\mathrm{SO}(3)$-irreducible decomposition is:

$$\mathcal{T} \simeq \bigoplus_{\ell} m_\ell \mathcal{H}^{(\ell)}$$

where $\mathcal{H}^{(\ell)}$ is the spin-ℓ irrep. Here, $m_\ell \mathcal{H}^{(\ell)} \simeq \mathbb{R}^{m_\ell} \otimes \mathcal{H}^{(\ell)}$, with $\mathrm{SO}(3)$ acting trivially on the multiplicity space. An $\mathrm{SO}(3)$-invariant index constraint $\mathcal{S}$ induces a projection operator $\Pi_{\mathcal{S}} : \mathcal{T} \to \mathcal{T}_{\mathcal{S}}$ satisfying:

$$\Pi_{\mathcal{S}} \circ \rho(g) = \rho(g) \circ \Pi_{\mathcal{S}} \quad (\forall g \in \mathrm{SO}(3))$$

i.e., $\Pi_{\mathcal{S}}$ is an $\mathrm{SO}(3)$-equivariant map. The projection $\Pi_{\mathcal{S}}$ acts on the isotypic components as follows:

$$\Pi_{\mathcal{S}} = \bigoplus_{\ell} \left(A_\ell \otimes I_{2\ell+1} \right), \qquad A_\ell^2 = A_\ell.$$

By Schur's lemma, $A_\ell \in \mathbb{R}^{m_\ell \times m_\ell}$ acts on the multiplicity space, while $I_{2\ell+1}$ acts on $\mathcal{H}^{(\ell)}$. The constrained multiplicity is $m_\ell^{\mathcal{S}} = \mathrm{rank}(A_\ell) \leq m_\ell$, and $\mathcal{T}_{\mathcal{S}} \simeq \bigoplus_\ell m_\ell^{\mathcal{S}} \mathcal{H}^{(\ell)}$. Thus, a constraint may reduce the multiplicity of a spin without removing that spin entirely. Therefore, the constraint introduces no new spin types and may reduce their multiplicities: $\mathcal{I} = \{\ell \mid m_\ell^{\mathcal{S}} > 0\} \subseteq \{\ell \mid m_\ell > 0\}$.

B Examples of high-order tensor irreducible decomposition under $\mathrm{SO}(3)$ group

In this section, we present three illustrative examples demonstrating the irreducible decomposition of tensors of order 2 to 4 under the $\mathrm{SO}(3)$ group, based on **Proposition 1** and **Proposition 2**. We have validated the correctness of these examples through a series of references [62, 63, 33, 64, 40].

B.1 Irreducible decomposition of 2nd-order tensors

1. Unconstrained Case Any rank-2 tensor T_{ij} in 3D space can be decomposed into symmetric and antisymmetric parts, which further reduce into irreducible representations (irreps) of the rotation group $\mathrm{SO}(3)$.

$$T_{ij} \simeq \underbrace{1 \times \mathbf{1}}_{\ell=0} \oplus \underbrace{1 \times \mathbf{3}}_{\ell=1} \oplus \underbrace{1 \times \mathbf{5}}_{\ell=2}.$$

(1) Symmetric Traceless Component ($\mathcal{S}_2^{(\ell)}$)

$\mathcal{S}_2^{(2)}$: Traceless part (**5** independent modes)

$$\mathcal{S}_2^{(2)} : \quad T_{(ij)} - \frac{1}{3}\delta_{ij} T_{kk}, \quad \text{where } T_{(ij)} = \frac{1}{2}(T_{ij} + T_{ji}).$$

(2) Antisymmetric Component ($\mathcal{A}_2^{(\ell)}$):

($\mathcal{A}_2^{(1)}$): The skew-symmetric part maps to a pseudovector via the Levi-Civita symbol (**3** independent modes):

$$\mathcal{A}_2^{(1)} : \quad T_{[ij]} = \frac{1}{2}(T_{ij} - T_{ji}), \quad \text{equivalent to } \mathbf{v}_k = \epsilon_{ijk} T_{[ij]}.$$

(3) Symmetric Trace Components ($\mathcal{S}_2^{(\ell)}$) (can also be formulated as $\mathcal{M}_2^{(\ell)}$ based on **Proposition 1**)

The scalar (isotropic) part formed by the trace (**1** independent mode):

$$\mathcal{S}_2^{(0)} : \quad \frac{1}{3}\delta_{ij} T_{kk}, \quad \text{where } T_{kk} = \delta_{ij} T_{ij}.$$

Dimension Verification: $1 \times 1 + 1 \times 3 + 1 \times 5 = 9 = 3^2$ (matches rank-2 tensor in 3D).

2. Symmetry constraints In the dielectric tensor, due to the symmetry $T_{ij} = T_{ji}$, its antisymmetric part automatically disappears.

$$T_{ij} \simeq \underbrace{1 \times \mathbf{1}}_{\ell=0} \oplus \underbrace{1 \times \mathbf{5}}_{\ell=2}.$$

B.2 Irreducible decomposition of 3rd-order tensors

1. Unconstrained Case T_{ijk} decomposes into irreducible representations (irreps) of SO(3) labeled by angular momentum L:

$$T_{ijk} \simeq \underbrace{1 \times \mathbf{1}}_{\ell=0} \oplus \underbrace{3 \times \mathbf{3}}_{\ell=1} \oplus \underbrace{2 \times \mathbf{5}}_{\ell=2} \oplus \underbrace{1 \times \mathbf{7}}_{\ell=3},$$

(1) Symmetric Traceless Components ($\mathcal{S}_3^{(\ell=3)}$)

$\ell = 3$ (Octupole): Fully symmetric traceless part (with **7** independent modes):

$$\mathcal{S}_{ijk}^{(3)}: \quad T_{(ijk)} - \frac{1}{5}\left(\delta_{ij}\langle T\rangle_k + \delta_{ik}\langle T\rangle_j + \delta_{jk}\langle T\rangle_i\right), \qquad \langle T\rangle_k \equiv T_{(ppk)}.$$

(2) Antisymmetric Components ($\mathcal{A}_3^{(\ell)}$)

$\mathcal{A}_3^{(\ell=0)}$ with only **1** independent mode using Levi-Civita contraction:

$$\mathcal{A}_{ijk}^{(\ell=0)}: \quad \frac{1}{6}\epsilon_{ijk}\left(\epsilon_{pqr}T_{pqr}\right).$$

(3) Trace and mixed-symmetry components ($\ell = 1, 2$)

Type-I (Symmetric in 2 indices):

$$H_{ijk} = T_{(ij)k} - \frac{1}{3}\delta_{ij}a_k, \qquad a_k = T_{ppk}.$$

Decompose $T_{(ij)k}$ into:

- Trace: $a_k = T_{ppk}$ gives one $\ell = 1$ component (**3** independent modes).
- Pair-traceless remainder: H_{ijk} decomposes as $\mathbf{3} \oplus \mathbf{5} \oplus \mathbf{7}$ ($\ell = 1, 2, 3$). Its $\ell = 3$ component is the fully symmetric traceless part in (1); the other two components contribute to the remaining $\ell = 1, 2$ sectors.

Type-II (Antisymmetric in 2 indices):

$$T_{[ij]k} - \frac{1}{6}\epsilon_{ijk}\big(\epsilon_{pqr}T_{[pq]r}\big).$$

Define $B_{mk} = \frac{1}{2}\epsilon_{mij}T_{[ij]k}$. Its trace is the $\ell = 0$ component already given in (2); after removing that component, the remainder splits into:

- Vector: $v_a = \epsilon_{amk}B_{[mk]}$ gives the $\ell = 1$ component (**3** independent modes).
- Symmetric traceless part: $Q_{mk} = B_{(mk)} - \frac{1}{3}\delta_{mk}B_{pp}$ gives the $\ell = 2$ component (**5** independent modes).

After removing the $\ell = 3$ and $\ell = 0$ components in (1) and (2), the remaining trace and mixed-symmetry components are:

$$\mathcal{M}_3 = \underbrace{3 \times \mathbf{3}}_{\ell=1} \oplus \underbrace{2 \times \mathbf{5}}_{\ell=2}.$$

Dimension Verification: $1 \times 1 + 3 \times 3 + 2 \times 5 + 1 \times 7 = 27 = 3^3$. (matches rank-3 tensor in 3D)

2. Symmetry Constraints When T_{ijk} is symmetric in its last two indices ($T_{ijk} = T_{ikj}$ in piezoelectric tensor ϵ_{ijk}), the decomposition simplifies dramatically:

$$T_{ijk} \simeq \underbrace{2 \times \mathbf{3}}_{\ell=1} \oplus \underbrace{1 \times \mathbf{5}}_{\ell=2} \oplus \underbrace{1 \times \mathbf{7}}_{\ell=3},$$

reducing the independent components from 27 to 18.

B.3 Irreducible decomposition of 4th-order tensors

1. Unconstrained Case For a rank-4 tensor $C_{ijkl} \in V^{\otimes 4}$, the decomposition is explicitly constructed as follows:

$$T_{ijkl} \simeq \underbrace{3 \times \mathbf{1}}_{\ell=0} \oplus \underbrace{6 \times \mathbf{3}}_{\ell=1} \oplus \underbrace{6 \times \mathbf{5}}_{\ell=2} \oplus \underbrace{3 \times \mathbf{7}}_{\ell=3} \oplus \underbrace{1 \times \mathbf{9}}_{\ell=4},$$

(1) Symmetric Components ($\mathcal{S}_4^{(\ell)}$)

$\ell = 4$ (Hexadecapole): Fully symmetric and traceless part:

$$\mathcal{S}_4^{(4)}: \quad C_{(ijkl)} - \text{lower-order traces}.$$

This part is constructed by using the symmetrization operator $\mathcal{P}_4$ and subtracting all possible traces (e.g., δ_{ij} terms)

(2) Antisymmetric Components ($\mathcal{A}_4^{(\ell)}$)

In 3D, all $\mathcal{A}_4^{(\ell)}$ vanish because there is no nonzero fully antisymmetric rank-4 tensor (the Levi-Civita symbol ϵ_{ijk} only supports 3 indices).

(3) Lower-spin components. Let Π_ℓ denote the orthogonal projector onto the spin-ℓ isotypic component of $V^{\otimes 4}$ under $\mathrm{SO}(3)$.

$\ell = 0$ (Scalar): Independent modes are fully contracted scalars (can also be formulated as $\mathcal{S}_4^{(0)}$:

$$\mathcal{M}_4^{(0)}: \quad C_{iijj},\ C_{ijij},\ C_{ijji} \quad (\text{3 independent scalars}).$$

$\ell = 1$ (Vector): Six independent vector copies are obtained by tracing each possible pair of tensor indices and contracting the remaining pair with the Levi-Civita symbol:

$$\begin{aligned} V_i^{(1)} &= \epsilon_{ijk} C_{ppjk}, \quad V_i^{(2)} = \epsilon_{ijk} C_{pjpk}, \\ V_i^{(3)} &= \epsilon_{ijk} C_{pjkp}, \quad V_i^{(4)} = \epsilon_{ijk} C_{jppk}, \\ V_i^{(5)} &= \epsilon_{ijk} C_{jpkp}, \quad V_i^{(6)} = \epsilon_{ijk} C_{jkpp}. \end{aligned}$$

$\ell = 2$ (Quadrupole): The spin-2 component comprises six copies of the five-dimensional symmetric traceless representation:

$$\mathcal{M}_4^{(2)} = \Pi_2 C \qquad (\text{the spin-0, 1, 3, and 4 components are removed}).$$

$\ell = 3$ (Octupole): The spin-3 component comprises three copies of the seven-dimensional symmetric traceless representation:

$$\mathcal{M}_4^{(3)} = \Pi_3 C \qquad (\text{the spin-0, 1, 2, and 4 components are removed}).$$

Dimension Verification: $3 \times 1 + 6 \times 3 + 6 \times 5 + 3 \times 7 + 1 \times 9 = 81 = 3^4$. matches rank-4 tensor in 3D

2. Symmetry Constraints When applied major symmetry ($C_{ijkl} = C_{klij}$) and minor symmetry ($C_{ijkl} = C_{jikl} = C_{ijlk}$) on 4th-order tensor, the decomposition components are reduced:

$$T_{ijkl} \simeq \underbrace{2 \times \mathbf{1}}_{\ell=0} \oplus \underbrace{2 \times \mathbf{5}}_{\ell=2} \oplus \underbrace{1 \times \mathbf{9}}_{\ell=4},$$

C Details of the physical principles in the properties of high-order tensors

C.1 Dielectric

Dielectric tensor is calculated by $\mathbf{D} = \epsilon\mathbf{E}$, which can be written as:

$$\begin{pmatrix} D_x \\ D_y \\ D_z \end{pmatrix} = \begin{pmatrix} \varepsilon_{xx} & \varepsilon_{xy} & \varepsilon_{xz} \\ \varepsilon_{yx} & \varepsilon_{yy} & \varepsilon_{yz} \\ \varepsilon_{zx} & \varepsilon_{zy} & \varepsilon_{zz} \end{pmatrix} \begin{pmatrix} E_x \\ E_y \\ E_z \end{pmatrix}$$

where $\mathbf{E} \in \mathbb{R}^3$ is an externally applied electric field and $\mathbf{D}$ is the resultant electric displacement field $\mathbf{D} \in \mathbb{R}^3$

Due to the symmetry $\varepsilon_{ij} = \varepsilon_{ji}$, i.e., $\varepsilon_{xy} = \varepsilon_{yx}$, $\varepsilon_{xz} = \varepsilon_{zx}$, etc., the antisymmetric components are reduced to zero, leaving six independent modes, as discussed in Appendix B.1.

C.2 Piezoelectric

Piezoelectric Tensor is calculated by

$$e_{ijk} = \left(\frac{\partial D_i}{\partial \varepsilon_{jk}}\right)_E = -\left(\frac{\partial \sigma_{jk}}{\partial E_i}\right)_{\boldsymbol{\epsilon}}$$

where $\mathbf{D} \in \mathbb{R}^3, \mathbf{E} \in \mathbb{R}^3$ are the electric displacement tensor and the electric field, $\boldsymbol{\sigma} \in \mathbb{R}^{3\times3}, \boldsymbol{\epsilon} \in \mathbb{R}^{3\times3}$ are the second-order elastic stress tensor and strain tensor.

Due to the symmetry of the stress tensor $\boldsymbol{\sigma}_{jk} = \boldsymbol{\sigma}_{kj}$, the piezoelectric tensor satisfies the symmetry relation:

$$\mathbf{e}^{ijk} = \mathbf{e}^{ikj}$$

In general, such a pair-symmetric tensor has 18 independent components which correspond to the dimensions of Appendix B.2. Formally, the piezoelectric tensor can also be formulated by the Voigt notation as:

$$\mathbf{e} = \begin{pmatrix} e_{11} & e_{12} & e_{13} & e_{14} & e_{15} & e_{16} \\ e_{21} & e_{22} & e_{23} & e_{24} & e_{25} & e_{26} \\ e_{31} & e_{32} & e_{33} & e_{34} & e_{35} & e_{36} \end{pmatrix}$$

C.3 Elastic

The elastic tensor (or elastic stiffness tensor) C_{ijkl} is a fourth-rank tensor that quantifies the linear relationship between strain ϵ_{kl} and stress σ_{ij} in a solid material under small deformations (Hooke's law). It fully characterizes the stiffness (resistance to deformation) of an elastic material. According to the principle of elastic tensor, here are two definitions:

(1) **Stress-Strain Relationship** In the linear elasticity regime, the constitutive equation (generalized Hooke's law) is:

$$\sigma_{ij} = C_{ijkl}\epsilon_{kl}$$

where σ_{ij} = Cauchy stress tensor (force per unit area, symmetric: $\sigma_{ij} = \sigma_{ji}$) ϵ_{kl} = infinitesimal strain tensor (symmetric: $\epsilon_{kl} = \epsilon_{lk}$) C_{ijkl} = elastic stiffness tensor, with 81 components (reducible to 21 independent components due to symmetry).

(2) **Strain Energy Density** The elastic tensor is derived from the strain energy density Ψ (energy per unit volume stored due to deformation), which is quadratic in strain for a linear elastic material:

$$\Psi = \frac{1}{2}C_{ijkl}\epsilon_{ij}\epsilon_{kl}$$

Stress is obtained from the thermodynamic derivative:

$$\sigma_{ij} = \frac{\partial \Psi}{\partial \epsilon_{ij}} = C_{ijkl}\epsilon_{kl}$$

The elastic tensor can also be defined as the second derivative of energy with respect to strain:

$$C_{ijkl} = \frac{\partial^2 \Psi}{\partial \epsilon_{ij} \partial \epsilon_{kl}}$$

Due to symmetry ($\sigma_{ij} = \sigma_{ji}, \epsilon_{kl} = \epsilon_{lk}$), the tensor is compressed into a 6×6 stiffness matrix $\mathbf{C}_{\alpha\beta}$:

$$\begin{bmatrix} \sigma_1 \\ \sigma_2 \\ \sigma_3 \\ \sigma_4 \\ \sigma_5 \\ \sigma_6 \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & C_{14} & C_{15} & C_{16} \\ C_{12} & C_{22} & C_{23} & C_{24} & C_{25} & C_{26} \\ C_{13} & C_{23} & C_{33} & C_{34} & C_{35} & C_{36} \\ C_{14} & C_{24} & C_{34} & C_{44} & C_{45} & C_{46} \\ C_{15} & C_{25} & C_{35} & C_{45} & C_{55} & C_{56} \\ C_{16} & C_{26} & C_{36} & C_{46} & C_{56} & C_{66} \end{bmatrix} \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \\ \epsilon_3 \\ \epsilon_4 \\ \epsilon_5 \\ \epsilon_6 \end{bmatrix}$$

where stress and strain are in Voigt Notation: (1) $\sigma_1 = \sigma_{11}$, $\sigma_4 = \sigma_{23}$ (shear components) (2) $\epsilon_1 = \epsilon_{11}$, $\epsilon_4 = 2\epsilon_{23}$ (engineering shear strain). Considering the whole symmetry of elastic tensor ($C_{ijkl} = C_{klij} = C_{jikl} = C_{ijlk}$), 36 components can be further reduced into 21 components, which correspond to B.3:

$$C^{ijkl} = \begin{bmatrix} C^{1111} & C^{1122} & C^{1133} & C^{1123} & C^{1131} & C^{1112} \\ * & C^{2222} & C^{2233} & C^{2223} & C^{2231} & C^{2212} \\ * & * & C^{3333} & C^{3323} & C^{3331} & C^{3312} \\ * & * & * & C^{2323} & C^{2331} & C^{2312} \\ * & * & * & * & C^{3131} & C^{3112} \\ * & * & * & * & * & C^{1212} \end{bmatrix}$$

$$= \begin{bmatrix} C^{11} & C^{12} & C^{13} & C^{14} & C^{15} & C^{16} \\ * & C^{22} & C^{23} & C^{24} & C^{25} & C^{26} \\ * & * & C^{33} & C^{34} & C^{35} & C^{36} \\ * & * & * & C^{44} & C^{45} & C^{46} \\ * & * & * & * & C^{55} & C^{56} \\ * & * & * & * & * & C^{66} \end{bmatrix}$$

D Parameter setting

D.1 Details of baselines

MEGNET Implementations of MEGNet [15]. Following the original paper, we use three layers of MEGNet message passing with the same feature dimensions as mentioned in the paper. We train MEGNet for 200 epochs using Huber loss with a learning rate of 0.001 and AdamW optimizer with 10^{-5} weight decay. The same polynomial learning rate decay scheduler as our model implementation is used. To predict a tensor matrix of shape 3×3, we change the original output dimension from one to nine.

ETGNN Following ETGNN [51], we use four DimeNet++ [65] layers with hidden dimension 128 and ELU activation function to serve as the invariant message-passing network. Since the code of ETGNN is not publicly accessible, we implement DimeNet++ layers following the code provided by GMTNet. We train ETGNN for 200 epochs using Huber loss with a learning rate of 0.001 and AdamW optimizer with 10^{-5} weight decay. The same polynomial learning rate decay scheduler as our model implementation is used.

GMTNet Following the original paper [25], we use 2 layers of node invariant feature updating and 3 layers of equivariant message passing for all tasks including dielectric, piezoelectric, and elastic tensors. Learning rate of 0.001, epoch number of 200, and batch size of 64 are used for dielectric, piezoelectric, and elastic tensors. The radius of construction crystal graphs is determined by the 16-th nearest neighbor.

GoeCTP We independently implement GoeCTP [52], using the published method and the authors' GitHub repository as implementation references. For each crystal, we first apply polar decomposition to the lattice matrix to obtain a canonical lattice and the associated orthogonal transformation. We

then construct the crystal graph in the canonical frame and use the prediction network to estimate tensor components in that frame. Finally, we map the predicted tensor back to the input coordinate frame by applying the orthogonal transformation to each tensor index. Our implementation uses the same property-specific datasets and evaluation metrics as the other baselines. Training uses 200 epochs, a batch size of 64, the AdamW optimizer with a weight decay of 10^{-5}, Huber loss, and polynomial learning-rate decay. The reported GoeCTP results are obtained from this independent implementation.

D.2 Details of IrredNet

Our model is implemented based on the MACE [24] and GMTNet repositories. In the graph construction block, we set the radius to 5 and use 20 basis functions of the Bessel function for encoding edge length. The angular momentum number ℓ = 3 (for dielectric), ℓ = 3 (for piezoelectric), and ℓ = 4 (for elastic) are used to encode the edge directions. Additionally, we employ four spherical harmonic convolution layers, along with gates and batch normalization, to learn the node features. Learning rate of 0.00001, epoch number of 200, and batch size of 64 are used for dielectric, piezoelectric, and elastic tensors. All experiments were conducted on a computing platform with four NVIDIA RTX 3090 GPUs, running Ubuntu 22.04, an Intel Xeon Gold 5120 CPU at 2.20 GHz, and 512 GB of system memory. Our code will be released when the paper is publicly available.

E Limitations and Broader Impacts

This work focuses on symmetry-guaranteed prediction of high-order tensor properties for crystalline materials. Our experiments cover three representative tensor properties, including dielectric, piezoelectric, and elastic tensors, spanning second-, third-, and fourth-order cases. While these tasks provide a broad evaluation of the proposed framework, future work could further test the method on additional tensor-valued properties, larger materials datasets, and experimental measurements. The symmetry guarantee depends on the specified crystal symmetry and tensor index constraints; when these inputs are correctly identified, the reconstructed tensor satisfies the required symmetry constraints by construction.

The proposed method may help accelerate computational materials discovery by improving the physical consistency of machine-learning predictions and reducing the cost of screening candidate materials for applications such as mechanics, electronics, photonics, and energy-related devices. The work is primarily intended as a scientific modeling and screening tool. As with other data-driven materials prediction methods, predictions should be interpreted as a complement to, rather than a replacement for, first-principles calculations or experimental characterization in downstream applications.

We do not anticipate direct privacy, fairness, or human-subject risks, since the method uses crystalline materials data rather than personal or behavioral data. Potential societal impacts are mainly indirect and depend on how predicted materials are further validated and deployed in real-world engineering settings.

F Analysis of high-order tensor properties

F.1 Analysis of Elastic Tensor

The elastic stiffness tensor C_{ijkl} and the compliance tensor S_{ijkl} are two fundamental fourth-rank tensors in linear elasticity theory. They are inverse tensors of each other and together describe the stress-strain response of a material. In Voigt notation, $\mathbf{C}$ and $\mathbf{S}$ become 6×6 matrices, with $\mathbf{S} = \mathbf{C}^{-1}$. Formally, it's more common to use compliance tensor S_{ijkl} for calculating relevant properties [33, 66].

Young's modulus $E(n)$ Young's modulus describes the material's stiffness in the direction defined by the unit vector n. It is the inverse of the compliance tensor contracted along the direction of the applied stress.

$$E(\mathbf{n}) = \Big(S_{ijkl}\, n_i n_j n_k n_l \Big)^{-1}.$$

Poisson's ratio $\nu(n, m)$ Poisson's ratio characterizes the material's deformation in the lateral directions when stretched or compressed along the direction defined by the unit vector n. It is the ratio of lateral strain to axial strain.

$$\nu(\mathbf{n}, \mathbf{m}) = -\frac{S_{ijkl} n_i n_j m_k m_l}{S_{pqrs} n_p n_q n_r n_s}$$

Linear compressibility $\beta(n)$ Linear compressibility measures the material's response to uniform compression, defined as the negative of the inverse of the bulk modulus P, using the compliance tensor to calculate the volumetric strain.

$$\beta(\mathbf{n}) = -\frac{1}{P}\,\varepsilon^{(\mathbf{n})} = M_{ij}\,n_i\,n_j = S_{ijkl}\,\delta_{kl}\,n_i\,n_j = S_{ijkk}\,n_i\,n_j.$$

Shear modulus $G(n, m)$ The shear modulus $G(n, m)$ describes the material's resistance to shear deformation. It is calculated from the compliance tensor in the direction of the applied shear stress and strain.

$$G(\mathbf{n}, \mathbf{m}) = \frac{\tau}{\gamma} = \frac{1}{4\,S_{ijkl}\,n_i\,m_j\,n_k\,m_l}\,.$$

G Complexity analysis in algorithm 1

Complexity of basis generation. Algorithm 1 consists of two stages: (i) the unconstrained SO(3) irreducible decomposition via GENERATESO3IRREPS(k) and (ii) the symmetry-compatibility filtering APPLYSYM(M_0, G_{sym}, k), with an optional projector step to obtain explicit invariant bases. The function GENERATESO3IRREPS(k) enumerates all partitions of the tensor rank k with at most 3 parts and applies the SU(3)→SO(3) branching rules (Appendix J). The number of such partitions grows as

$$p_3(k) = \Theta(k^2),$$

and each branching call has quadratic cost in k, leading to an overall runtime

$$\mathcal{O}(k^4)$$

and $\mathcal{O}(|L_k|) = \mathcal{O}(k)$ memory for storing the map $M_0[\ell]$, where L_k is the set of distinct spins. The symmetry step APPLYSYM(M_0, G_{sym}, k) loops over all $\ell \in L_k$ and, for each spin, evaluates the O(3) character on all conjugacy classes of the point group G_{sym}. Its runtime is therefore

$$\mathcal{O}(|L_k| \cdot C_G) = \mathcal{O}(k \cdot C_G),$$

where C_G is the number of conjugacy classes (a small constant for crystallographic point groups), with again $\mathcal{O}(k)$ memory. When an explicit invariant basis is required, we form the projectors

$$\Pi_\ell = \frac{1}{|G_{\text{sym}}|} \sum_{g \in G_{\text{sym}}} \rho^{(\ell)}(g)$$

on each SO(3) block, which costs $\mathcal{O}(|G_{\text{sym}}| d_\ell^2 + d_\ell^3)$ per spin, where $d_\ell = 2\ell + 1$. Summed over all $\ell \le k$, this yields at most

$$\mathcal{O}(|G_{\text{sym}}| k^3 + k^4) \text{ time and } \mathcal{O}(k^3) \text{ memory.}$$

For the tensor orders considered in this work ($k \le 4$), the entire basis generation is a one-off precomputation that takes negligible time and memory compared to training or inference of the neural network.

Case Study (elastic tensor). Algorithm 1 describes the generic rank-k decomposition and point-group filtering. For the elastic tensor, we additionally restrict the rank-4 basis to the major/minor-symmetric subspace before calling APPLYSYM:

$$\mathcal{C} = \operatorname{im} \mathcal{P}_{\text{el}}, \qquad \mathcal{P}_{\text{el}} = \frac{1}{8}(I + P_{(13)(24)})(I + P_{(12)})(I + P_{(34)}),$$

where P_σ permutes tensor indices according to σ. This reduces the 81-dimensional space of rank-4 tensors to the 21-dimensional elasticity space, which decomposes under SO(3) as

$$\mathcal{C} \cong 2 \times \ell = 0 \;\oplus\; 2 \times \ell = 2 \;\oplus\; 1 \times \ell = 4,$$

i.e.

$$M_0^{\text{elastic}}[0] = 2, \quad M_0^{\text{elastic}}[2] = 2, \quad M_0^{\text{elastic}}[4] = 1.$$

The corresponding dimensions satisfy

$$2 \times (2 \cdot 0 + 1) \; + \; 2 \times (2 \cdot 2 + 1) \; + \; 1 \times (2 \cdot 4 + 1) = 2 + 10 + 9 = 21,$$

matching the 21 independent components of the elastic tensor. Given a crystal point group G_{sym}, the procedure APPLYSYM then computes, for each spin ℓ, the multiplicity n_ℓ of the trivial representation in the restriction of the $(2\ell+1)$-dimensional SO(3) irrep to G_{sym} via a class-sum character test. For a cubic crystal with $G_{\text{sym}} = O_h$, this yields

$$n_0 = 1, \qquad n_2 = 0, \qquad n_4 = 1,$$

and hence

$$M_{\text{final}}^{\text{elastic}}[0] = 2 \times n_0 = 2, \quad M_{\text{final}}^{\text{elastic}}[2] = 2 \times n_2 = 0, \quad M_{\text{final}}^{\text{elastic}}[4] = 1 \times n_4 = 1.$$

The total number of invariant coefficients is therefore

$$2 + 0 + 1 = 3,$$

in one-to-one correspondence with the three independent elastic constants (C_{11}, C_{12}, C_{44}) of a cubic crystal. For each spin ℓ, we finally construct explicit invariant bases $\{\mathcal{B}_a^{(\ell)}\}_{a=1}^{M_{\text{final}}^{\text{elastic}}[\ell]}$ via the projector

$$\Pi_\ell = \frac{1}{|G_{\text{sym}}|} \sum_{g \in G_{\text{sym}}} \rho^{(\ell)}(g).$$

The network predicts only the scalar coefficients of these irreducible basis tensors, so any reconstructed elastic tensor is guaranteed to satisfy both the index symmetries and the crystal point-group symmetry by construction.

H Guarantees for Symmetry and Accuracy

Building on the unconstrained decomposition and symmetry compatibility step in Section 3.1, we fix a canonical irreducible basis for the symmetry constraint G and learn only the associated basis coefficients. Reconstruction is restricted to the symmetry-compatible subspace S_G. Operationally, the group projector P_G implements the constraint by zeroing symmetry-forbidden coefficients and enforcing equality among symmetry-related ones, yielding an exactly constrained, compact parameterization. The number of free degrees of freedom drops from the ambient dimension D to the symmetry dimension d_G; estimation noise may perturb coefficient magnitudes but cannot introduce symmetry-forbidden patterns, so symmetry is guaranteed by construction.

Below we make this precise. Theorem 1 characterizes P_G via character projectors and explains why the reconstruction is exactly symmetry-compliant. Proposition 3 shows that projecting cannot increase the per-sample Frobenius error when the ground truth lies in S_G. Theorem 2 gives the expected risk identity $\mathbb{E}\|P_G(Y) - T^\star\|_F^2 = \operatorname{tr}(P_G\Sigma)$ and reduces to $\sigma^2 d_G$ under isotropic noise, while Corollary 1 shows that stronger symmetry implies no larger risk and typically a smaller one. Together these results justify the claim from the introduction that predicting symmetry-allowed irreducible components preserves symmetry exactly and improves accuracy in expectation.

Setting Let $(\mathcal{V}, \langle\cdot,\cdot\rangle)$ be the real inner-product space of order-k tensors with Frobenius norm $\|\cdot\|_F$ and dimension $D = \dim \mathcal{V}$. Let G be the crystallographic point group acting orthogonally via $\rho : G \to O(\mathcal{V})$. By complete reducibility, $\mathcal{V} = \bigoplus_\ell \mathcal{V}^{(\ell)}$, where $\mathcal{V}^{(\ell)}$ is the isotypic component for irrep ℓ with multiplicity m_ℓ and irrep dimension r_ℓ, so $\dim \mathcal{V}^{(\ell)} = m_\ell r_\ell$. Given the set of *allowed* irreps $\mathcal{L}_G$, define the symmetry subspace

$$S_G = \bigoplus_{\ell \in \mathcal{L}_G} \mathcal{V}^{(\ell)}, \qquad d_G = \dim S_G,$$

and let $P_G : \mathcal{V} \to S_G$ be the orthogonal projector.

Theorem 1 (Character projectors and exact symmetry)**.** *Define*

$$\Pi_\ell = \frac{r_\ell}{|G|} \sum_{g \in G} \overline{\chi_\ell(g)}\, \rho(g), \qquad P_G = \sum_{\ell \in \mathcal{L}_G} \Pi_\ell.$$

Then Π_ℓ projects onto $\mathcal{V}^{(\ell)}$ and P_G projects onto S_G. Consequently P_G zeros all forbidden components and ties symmetry-equivalent ones, yielding an exactly symmetry-compliant reconstruction.

Theorem 2 (Risk identity and isotropic case)**.** *Let* $\mathbb{E}[\varepsilon] = 0$ *and* $\Sigma = \mathbb{E}[\varepsilon\varepsilon^\top] \succeq 0$.

Then

$$\mathbb{E}\|P_G(Y) - T^\star\|_F^2 = \mathrm{tr}(P_G\Sigma), \qquad \mathbb{E}\|Y - T^\star\|_F^2 = \mathrm{tr}(\Sigma).$$

If the noise is isotropic, $\Sigma = \sigma^2 I$, then

$$\mathbb{E}\|P_G(Y) - T^\star\|_F^2 = \sigma^2 d_G, \qquad \mathbb{E}\|Y - T^\star\|_F^2 = \sigma^2 D.$$

Proposition 3 Inequality of point projection. If $T^\star \in S_G$ and $Y = T^\star + \varepsilon$, then

$$\|P_G(Y) - T^\star\|_F \le \|Y - T^\star\|_F,$$

with equality if and only if $\varepsilon \in S_G$.

Corollary 1 Monotonicity under stricter symmetry. If H is stricter than G so that $S_H \subset S_G$ and $P_H \preceq P_G$, then for any $\Sigma \succeq 0$,

$$\mathrm{tr}(P_H\Sigma) \le \mathrm{tr}(P_G\Sigma) \le \mathrm{tr}(\Sigma).$$

In particular, for $\Sigma = \sigma^2 I$, $\sigma^2 d_H \le \sigma^2 d_G \le \sigma^2 D$.

Finite-sample implication. For linear heads mapping features into S_G with Frobenius norm bounded by B, the empirical Rademacher complexity scales as

$$\mathfrak{R}_n \;\le\; \frac{B}{\sqrt{n}}\sqrt{d_G},$$

so restricting to S_G reduces capacity from a $\sqrt{D}$ to a $\sqrt{d_G}$ factor, improving sample efficiency.

In conclusion, predicting irreducible components confines outputs to S_G. This enforces symmetry by construction and reduces expected error from $\mathrm{tr}(\Sigma)$ to $\mathrm{tr}(P_G\Sigma)$, which equals $\sigma^2 d_G$ under isotropic noise. As symmetry becomes stricter and d_G shrinks, accuracy improves in expectation.

I Sensitivity Analysis for $\ell_{\max}$ and Channel Allocation

Choice of $\ell_{\max}$. Our architecture is implemented using the SO(2) reformulation of SO(3) convolutions (eSCN), while the feature space is still organized as SO(3) irreps up to degree $\ell_{\max}$. Thus $\ell_{\max} = \ell$ plays the usual role of the maximum angular momentum and controls the angular resolution of the model.

From representation theory, the *target tensor* itself only needs irreps up to a certain order. Concretely, symmetric rank-2 tensors: $\ell \le 2$, pair-symmetric rank-3 tensors: $\ell \le 3$, elastic rank-4 tensors: $\ell \le 4$. In principle, this means that a model predicting only a symmetric rank-2 property could already be implemented with $\ell_{\max} = 2$.

In our experiments, we choose $\ell_{\max}$ slightly more conservatively for two reasons.

First, for dielectric and piezoelectric prediction, we set $\ell_{\max} = 3$. For piezoelectric tensors (pair-symmetric rank-3), $\ell_{\max} = 3$ matches the maximal physically allowed ℓ. For dielectric tensors (symmetric rank-2), $\ell_{\max} = 3$ is one order higher than strictly required ($\ell \le 2$ would be sufficient), but it allows us to use a unified architecture across the second- and third-order tasks and provides slightly richer intermediate angular features. In all cases, the final tensor head only uses the irreps that are consistent with the target tensor's index symmetries.

Second, for elastic tensors, we set $\ell_{\max} = 4$, which fully covers the irreps needed for rank-4 elasticity while avoiding the cubic growth in cost with $\ell_{\max}$ inherent to eSCN-like convolutions.

These settings are also consistent with the hyperparameter choices reported in GMTNet. In Table 5, we further provide an ablation over $\ell_{\max}$, showing that smaller values (e.g., $\ell_{\max} = 2$ for dielectric, $\ell_{\max} = 2, 3$ for piezoelectric) do not lead to significant gains and that our default choice offers a good trade-off between accuracy and computational cost.

Table 5: $\ell_{\max}$ sensitivity experiments on the JARVIS-DFT dataset.

Dataset	$\ell_{\max}$	Fnorm (↓)	EwT 25% (↑)	EwT 10% (↑)	EwT 5% (↑)	Total Time (s) (↓)
Dielectric	2	3.31	**85.9%**	52.1%	**29.9%**	**711**
Dielectric	3	**3.23**	85.8%	**53.2%**	29.8%	798
Piezoelectric	3	**0.29**	49.3%	49.3%	**49.2%**	**288**
Piezoelectric	4	0.30	49.1%	**48.9%**	48.5%	314
Piezoelectric	5	0.33	49.2%	**48.9%**	48.6%	454
Elastic	4	64.23	**66.5%**	26.5%	**12.1%**	**14121**
Elastic	5	64.61	65.7%	26.1%	**12.1%**	16237
Elastic	6	**63.88**	**66.5%**	**26.9%**	12.0%	18711

Degree-wise channels. Channel counts are then tuned under a fixed compute budget: we allocate more channels to low and mid degrees (ℓ=0, 1, 2), which carry most of the physical signal, and fewer channels to higher degrees. This strategy provides sufficient capacity to model complex high-order anisotropy while maintaining stable training and reasonable runtime. For each degree ℓ and parity (even/odd), we assign channels in a monotonically decreasing manner with ℓ, roughly following a geometric decay (e.g., 32, 16, 8, 4, 2 for ℓ = 0,1,2,3,4). Low and mid degrees (ℓ = 0,1,2) are given significantly more channels because they carry most of the physical signal (scalar, vector, and low-order anisotropy), while higher degrees ($\ell \geq 3$) are assigned fewer channels and mainly serve to capture fine-grained anisotropy. Even and odd irreps at the same ℓ share the same width, reflecting that we do not impose any prior bias between parities. This degree-wise allocation gives a good trade-off between expressivity of high-order components and computational efficiency.

To assess the sensitivity of this design, we perform a channel-allocation ablation under a comparable total width (Table 6). Besides our default monotonically decreasing allocation (e.g., $32 \times 0e/0o, 16 \times 1e/1o, 8 \times 2e/2o, 4 \times 3e/3o, 2 \times 4e/4o$), we consider two alternatives: (i) an approximately flat allocation with 12 channels per parity at each degree (e.g., $12 \times 0e/0o, \ldots, 12 \times 4e/4o$), where the value 12 is chosen so that the overall number of channels is similar to the default setting, and (ii) a high-degree-heavy allocation that reverses the pattern and allocates most channels to $\ell = 3, 4$. Across all three tensor prediction tasks, our default scheme achieves the best or near-best Fnorm and EwT, while the high-degree-heavy configuration clearly degrades accuracy (e.g., for elastic tensors, Fnorm increases from 64.23 to 69.77 and EwT 25% drops from 66.5% to 61.5%). The flat allocation is only slightly worse than the default, indicating that the model is reasonably robust to moderate changes in the exact per-degree counts. Overall, these results support our design principle: allocating more capacity to low and mid degrees ($\ell = 0, 1, 2$) is important for performance, whereas over-emphasizing high degrees is not beneficial and mainly increases computational cost.

Table 6: Channel-sensitivity experiments on the JARVIS-DFT dataset.

Dataset	hidden irreps	Fnorm (↓)	EwT 25% (↑)	EwT 10% (↑)	EwT 5% (↑)
Dielectric	32x0e+32x0o+16x1e+16x1o+8x2e+8x2o+4x3e+4x3o	**3.23**	**85.8%**	**53.2%**	**29.8%**
Dielectric	12x0e+12x0o+12x1e+12x1o+12x2e+12x2o+12x3e+12x3o	3.43	83.8%	51.2%	26.7%
Dielectric	4x0e+4x0o+8x1e+8x1o+16x2e+16x2o+32x3e+32x3o	3.86	78.3%	40.9%	20.4%
Piezoelectric	32x0e+32x0o+16x1e+16x1o+8x2e+8x2o+4x3e+4x3o	**0.29**	**49.3%**	**49.3%**	**49.2%**
Piezoelectric	12x0e+12x0o+12x1e+12x1o+12x2e+12x2o+12x3e+12x3o	0.33	48.9%	48.7%	48.4%
Piezoelectric	4x0e+4x0o+8x1e+8x1o+16x2e+16x2o+32x3e+32x3o	0.41	47.3%	47.2%	46.8%
Elastic	32x0e+32x0o+16x1e+16x1o+8x2e+8x2o+4x3e+4x3o+2x4e+2x4o	**64.23**	66.5%	**26.5%**	**12.1%**
Elastic	12x0e+12x0o+12x1e+12x1o+12x2e+12x2o+12x3e+12x3o+12x4e+12x4o	64.41	**66.8%**	23.4%	11.5%
Elastic	2x0e+2x0o+4x1e+4x1o+8x2e+8x2o+16x3e+16x3o+32x4e+32x4o	69.77	61.5%	20.9%	7.5%

J Supplementary pseudo code analysis

J.1 Process of GENERATESO3IRREPS

The process begins with the construction of an unconstrained decomposition map, $\mathcal{M}_0$. The theoretical basis for this step is rooted in the representation theory of the general linear group $\mathrm{GL}(3)$ and the rotation group $\mathrm{SO}(3)$. This relies on key group-theoretical tools: first, Schur-Weyl Duality [40] is

employed to establish a correspondence between the tensor's permutation symmetries and the $\mathrm{GL}(3)$ irreducible representations indexed by Young Diagrams. Subsequently, Branching Rules are applied to decompose each $\mathrm{GL}(3)$ representation into a direct sum of $\mathrm{SO}(3)$ representations upon restriction to the $\mathrm{SO}(3)$ subgroup [41, 42]. This systematic approach, by iterating through all valid Young diagrams, automatically accounts for the contributions from all of the tensor's symmetry components, including the totally symmetric (corresponding to single-row diagrams), totally antisymmetric (single-column diagrams), and mixed symmetries (all other diagram shapes).

Parity in symmetry compatibility. CHARO3 evaluates the character of the standard spherical-harmonic representation, whose inversion parity is $(-1)^\ell$. For the rank-k polar tensors considered here, inversion instead acts as $(-1)^k$ on every irreducible component. Thus, for $g \in O(3)$, the target-tensor character is

$$\chi_{k,\ell}(g) = (\det g)^{k-\ell}\chi_{\mathrm{SH},\ell}(g).$$

Here, the rotation angle supplied to CHARO3 is that of $R = (\det g)g \in SO(3)$. For improper operations, the correction gives $(-1)^{k-\ell}(-1)^\ell = (-1)^k$; for proper rotations, the factor is one. The same correction is applied to the representation used in the invariant projector:

$$\rho_{k,\ell}(g) = (\det g)^{k-\ell}\rho_{\mathrm{SH},\ell}(g).$$

For the fixed target rank k, $\chi^{(\ell)}_{O(3)}$ and $\rho^{(\ell)}$ in the main text denote $\chi_{k,\ell}$ and $\rho_{k,\ell}$, respectively.

J.2 Complexity analysis of irreducible decomposition

We also analyze the computational complexity of irreducible decomposition of high-order tensor. The algorithm consists of two main parts: 1. $\mathcal{M}_0 \leftarrow$ GENERATEIRREPSk : Generates the unconstrained decomposition, which is the most computationally expensive part. 2. $\mathcal{M}_{final} \leftarrow$ APPLYSYMMETRY$\mathcal{M}_0, G_{sym}, k$.

Step 1: GENERATEPARTITIONS(K, MAX_ROWS=3) The number of integer partitions of k into at most 3 parts, $P_3(k)$, is on the order of $O(k^2)$. The time and space required to generate these partitions are both $O(k^2)$.

Main Loop This loop iterates $P_3(k) \approx O(k^2)$ times. The complexity of the loop body is as follows:

- **Step 2** (HOOKLENGTHFORMULA): Time complexity is $O(k)$.
- **Step 3** (BRANCHSU3TOSO3): The complexity is approximately $O(p \times q)$, which in the worst case is $O(k^2)$.
- **Step 4** (Accumulation): This loop runs $O(k^2)$ times, resulting in a complexity of $O(k^2)$.

Summary for GENERATESO3IRREPS The complexity is summarized as follows:

- **Time Complexity:** The main loop runs $O(k^2)$ times, and the most expensive step within it is $O(k^2)$. The total time complexity is $O(k^2) \times O(k^2) = O(k^4)$.
- **Space Complexity:** Dominated by storing the $O(k^2)$ partitions, the space complexity is $O(k^2)$.

2. Complexity Analysis of APPLYSYMMETRY($\mathcal{M}_0$, G_{SYM}, k)

This function filters the unconstrained decomposition map $\mathcal{M}_0$.

Loop Iterations The size of $\mathcal{M}_0$ is at most $k+1$ (since $\ell_{max} = k$), so the main loop runs $O(k)$ times.

Function Call The ISCOMPATIBLE function's complexity depends on the number of conjugacy classes, $|C_{G_{\mathrm{sym}}}|$, which is a constant with respect to k. Thus, its complexity is $O(|C_{G_{\mathrm{sym}}}|)$.

Summary for APPLYSYMMETRY The complexity is summarized as follows:

- **Time Complexity:** $O(k) \times O(|C_{G_{\mathrm{sym}}}|) = O(k \cdot |C_{G_{\mathrm{sym}}}|)$.
- **Space Complexity:** A new map of size at most $O(k)$ is created, so the space complexity is $O(k)$.

Final Conclusion

By combining the analysis of all parts, we can determine the overall complexity of the algorithm:

- **Total Time Complexity:** $T(k) = O(k^4) + O(k \cdot |C_{G_{\text{sym}}}|)$. Since k^4 is the dominant term, the final time complexity is $\boldsymbol{O(k^4)}$.
- **Total Space Complexity:** $S(k) = O(k^2) + O(k)$. Since k^2 is the dominant term, the final space complexity is $\boldsymbol{O(k^2)}$.

The computational cost of the algorithm is primarily driven by the GENERATESO3IRREPS function. In practice, during training this computation has minimal impact on overall cost; for each tensor it only needs to be performed once upfront.

Algorithm 2 Supplementary Pseudo Code: HOOKLENGTHFORMULA

Input: An integer partition $\lambda = [\lambda_1, \lambda_2, \dots]$
Output: The dimension of the S_k irreducible representation for λ

Function HookLengthFormula(λ)
$k \leftarrow \sum_{i=1}^{\text{length}(\lambda)} \lambda_i$
if $k = 0$ **then**
    return 1
end if
$k_{\text{factorial}} \leftarrow k!$
hook_product $\leftarrow 1$
for $i \leftarrow 1$ to length(λ) **do**
    // Iterate rows
    for $j \leftarrow 1$ to λ_i **do**
        // Iterate columns
        boxes_right $\leftarrow \lambda_i - j$
        boxes_below $\leftarrow 0$
        for $m \leftarrow i + 1$ to length(λ) **do**
            if $\lambda_m \geq j$ **then**
                boxes_below $\leftarrow$ boxes_below $+ 1$
            end if
        end for
        hook_length $\leftarrow$ boxes_right $+$ boxes_below $+ 1$
        hook_product $\leftarrow$ hook_product $\times$ hook_length
    end for
end for
dimension $\leftarrow k_{\text{factorial}}$/hook_product
return round(dimension)
End Function

Algorithm 3 Supplementary Pseudo Code: BRANCHRULESU3TOSO3 and GENERATEPARTITIONS

Part 1: Branching Rules U(3) to SO(3)
Input: Integers $p = \lambda_1 - \lambda_2$, $q = \lambda_2 - \lambda_3$
Output: A map M_{branch} from SO(3) irrep L to multiplicity

Function BranchRulesU3toSO3(p, q)
$M_{\text{branch}} \leftarrow$ new Map()
$\lambda \leftarrow p$, $\mu \leftarrow q$
for $K \leftarrow \mu$ down to 0 step -2 **do**
    if $K = 0$ **then**
        for $L \leftarrow \lambda$ down to 0 step -2 **do**
            $M_{\text{branch}}[L] \leftarrow (M_{\text{branch}}[L] \text{ or } 0) + 1$
        end for
    else
        for $L \leftarrow K$ to $K + \lambda$ **do**
            $M_{\text{branch}}[L] \leftarrow (M_{\text{branch}}[L] \text{ or } 0) + 1$
        end for
    end if
end for
return M_{branch}
End Function

Part 2: Partition Generation
Input: Integer k, max rows max_rows
Output: List of partitions

Function GeneratePartitions(k, max_rows)
all_parts $\leftarrow$ new List()
cur_part $\leftarrow$ new List()
if $k > 0$ **then**
    ReFindParts(k, max_rows, k, cur_part, all_parts)
end if
return all_parts
End Function

Procedure ReFindParts(k, max, start, cur, all)
if $k = 0$ **then**
    Add copy of cur to all
    return
end if
if $k < 0$ **or** max $= 0$ **then**
    return
end if
for $i \leftarrow \min(k, \text{start})$ down to 1 **do**
    Add i to cur
    ReFindParts($k - i$, max $- 1$, i, cur, all)
    Remove last element from cur *// Backtrack*
end for
End Procedure

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The main claims in the abstract and introduction are aligned with the proposed contribution.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper includes a combined limitations and broader impacts section discussing the empirical scope of the evaluation, the dependence on correctly specified symmetry constraints, and possible extensions to additional tensor properties and material settings.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We provide the assumptions and proof in the paper.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The paper describes the datasets, preprocessing protocol, baselines, evaluation metrics, data split, optimization details, and main hyperparameters needed to reproduce the core experimental results. Additional implementation details and sensitivity settings are provided in the appendix.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in

some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: Our code is publicly available at: `github.com/QiaolinLu/IrredNet`.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: We specify all the details in the appendix.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: We calculated the mean results and don't report error bars.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The paper reports wall-clock training/evaluation time for the main experiments and provides hardware details including GPU type, CPU, memory, and operating system. These details are sufficient to estimate the compute needed to reproduce the reported experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: We follow and respect NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: The paper includes a combined limitations and broader impacts section discussing potential benefits for materials discovery and computational screening, as well as indirect risks from unvalidated downstream use in engineering applications.

    Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: No such risks.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: We respect the license to use other assets.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: We don't release new assets.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing, user studies, or research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve human subjects or crowdsourcing experiments. Therefore, IRB approval or equivalent review is not applicable.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: The core method development does not use LLMs as an important, original, or non-standard component.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.