EACO: PREDICTING $\mathrm{O}(3)$-EQUIVARIANT QUANTUM OPERATORS VIA EDGE-ALIGNED $\mathrm{O}(2)$ CONVOLUTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Predicting quantum operator matrices from atomic structures with group symmetries enables large-scale electronic-structure calculations and high-throughput materials screening. Focusing on this task, current edge-aligned $\mathrm{SO}(2)$ convolutions guarantee global $\mathrm{SO}(3)$ equivariance. Still, under reflections or spatial inversion, they cannot force particular operator components to vanish or protect electronic degeneracies. To overcome this issue, we propose Edge-Aligned $\mathrm{O}(2)$ Convolution (**EACO**), a message-passing layer that achieves local $\mathrm{O}(2)$ equivariance with strict locality and high computational efficiency. In particular, EACO assigns each $\mathrm{O}(3)$ irreducible representation (irrep) of angular degree ℓ and inversion parity p the sign $\tau_{\ell,p} = p(-1)^{\ell}$, indicating how it transforms under a reflection in the edge-aligned local frame and thus distinguishing the irreps with the same angular degree but opposite inversion parities. We prove that local $\mathrm{O}(2)$ equivariance is necessary and sufficient for global $\mathrm{O}(3)$ equivariance within this construction, so the model built on EACO preserves the transformation laws of the full $\mathrm{O}(3)$ group. Compared with SLEM, EACO reduces Hamiltonian MAE by up to 22.1% across five materials, with 17.2%–29.6% fewer parameters. It also reduces equivariance errors under reflections and spatial inversion by three to four orders of magnitude. On graphene, EACO recovers the horizontal-mirror (σ_h) selection rule for flexural phonons, suppressing forbidden electron–phonon coupling while accurately reproducing allowed coupling.

1 INTRODUCTION

In density functional theory (DFT), quantum operators, such as the Hamiltonian and density matrices, encode the electronic structures of molecules and materials (Hohenberg & Kohn, 1964; Kohn & Sham, 1965; Zhouyin et al., 2025). Predicting these matrices from atomic structures can bypass costly self-consistent field iterations (Payne et al., 1992; Kudin et al., 2002), enabling large-scale electronic-structure simulations and high-throughput materials screening (Unke et al., 2021; Gong et al., 2023; Zhouyin et al., 2025; Zhong et al., 2026; Schwade et al., 2026; Yin et al., 2026). Since atomic structures and operator matrices transform under the same coordinate systems, a quantum operator predictor should consider the group symmetries of the atomic structures, and accordingly, preserve geometric equivariance.

To build a physically consistent predictor, current equivariant message-passing networks represent atomic features as irreducible representations (irreps) and update them through tensor products with spherical harmonics (Unke et al., 2021; Gong et al., 2023; Yu et al., 2023; Zhong et al., 2023; Qian et al., 2026). Among them, models based on full Clebsch–Gordan tensor products have an angular cost of $\mathcal{O}(\ell_{\max}^6)$, where $\ell_{\max}$ is the maximum angular degree. Recently, edge-aligned convolutions express irreps in a local frame aligned with each edge and replace tensor products with $\mathrm{SO}(2)$ convolutions (Passaro & Zitnick, 2023; Liao et al., 2024; Zhouyin et al., 2025; Yu et al., 2026). Although these models reduce the computational complexity from $\mathcal{O}(\ell_{\max}^6)$ to $\mathcal{O}(\ell_{\max}^3)$, they only impose local $\mathrm{SO}(2)$ constraints on rotations while leaving reflections and spatial inversion unconstrained. As a result, these models merely guarantee global $\mathrm{SO}(3)$ equivariance, which cannot distinguish $\mathrm{O}(3)$ irreps of the same degree but opposite inversion parities.

Preserving $\mathrm{O}(3)$ equivariance is an explicit design goal in recent Hamiltonian models (Yin et al., 2026). In practice, reflections and spatial inversion matter for quantum operator prediction, as

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

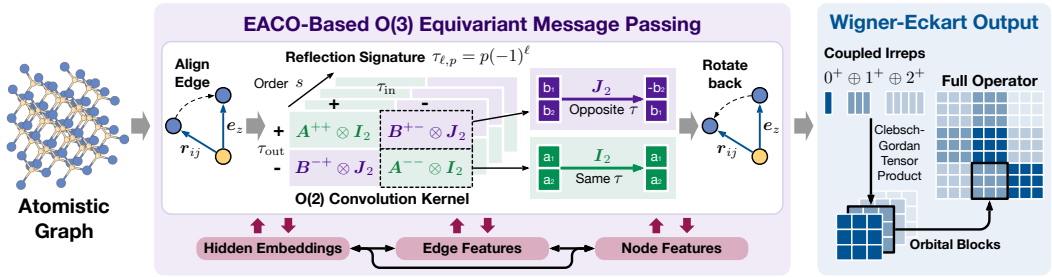

Figure 1: The modeling principle and pipeline of EACO. In particular, EACO augments an edge-aligned $\mathrm{SO}(2)$ convolution with the local reflection constraint required for global $\mathrm{O}(3)$ equivariance.

they can force Hamiltonian components to vanish or protect electronic degeneracies. For example, horizontal mirror symmetry forces the first-order p_z–p_z Hamiltonian response to out-of-plane displacements to vanish in graphene (Mariani & von Oppen, 2008; Alidoosti et al., 2022), inversion symmetry forbids the conventional homogeneous shift-current response (Yang et al., 2026; Shi et al., 2021), and mirror symmetry protects surface Dirac cones in SnTe (Hsieh et al., 2012).[1] **Therefore, it is critical to achieve $\mathrm{O}(3)$-equivariant quantum operator prediction while maintaining the complexity as low as $\mathcal{O}(\ell_{\max}^3)$.**

To overcome the above limitation, we propose Edge-Aligned $\mathrm{O}(2)$ Convolution (**EACO**), an efficient and strictly local message-passing layer for $\mathrm{O}(3)$-equivariant quantum operator prediction. As illustrated in Figure 1, EACO constructs an edge-aligned axis for each local frame, so that the rotation about the axis and the reflection through the planes containing the axis form an $\mathrm{O}(2)$ stabilizer. For each $\mathrm{O}(3)$ irreducible representation (irrep) of angular degree ℓ and inversion parity p, EACO uses $\tau_{\ell,p} = p(-1)^\ell$ to indicate how the irrep transforms under a local reflection and thus distinguish it from the irrep with the same angular degree but opposite inversion parities. The resulting convolution kernel is $\mathrm{O}(2)$-equivariant, which reduces the number of trainable parameters and preserves strict locality and $\mathcal{O}(\ell_{\max}^3)$ computational complexity. Moreover, we prove that global $\mathrm{O}(3)$ equivariance holds if and only if the local convolution kernel is $\mathrm{O}(2)$-equivariant within the class of edge-aligned local-frame maps, which guarantees the global $\mathrm{O}(3)$ equivariance of EACO in theory.

We evaluate EACO on Hamiltonian prediction for five materials (Si, GaN, HfO_2, graphene, and MoS_2). Compared with SLEM, EACO reduces Hamiltonian MAE by up to 22.1% across five materials, with 17.2%–29.6% fewer parameters. It also reduces equivariance errors under reflections and spatial inversion by **three to four orders of magnitude**. The graphene experiment further demonstrates its physical effectiveness: EACO preserves mirror equivariance and suppresses forbidden electron–phonon coupling for the out-of-plane acoustic (ZA) and optical (ZO) modes, while reproducing the allowed coupling with relative errors of 0.100% and 0.113%, respectively.

2 Preliminaries and Related Work

Given an atom set $\mathcal{V}$, in a localized atomic-orbital basis (Soler et al., 2002), we denote $\phi_{in\ell m}$ as an orbital centered at atom $i \in \mathcal{V}$, where n, ℓ, and $m \in \{-\ell, ..., \ell\}$ denote its radial index, angular degree, and magnetic quantum number, respectively. A one-electron quantum operator O, such as the Kohn–Sham Hamiltonian (Kohn & Sham, 1965), is represented by atom-pair orbital blocks

$$\boldsymbol{O}^{ij}_{n_1\ell_1,n_2\ell_2} = \left[\langle \phi_{in_1\ell_1m_1} | O | \phi_{jn_2\ell_2m_2} \rangle\right]_{m_1,m_2} = [O^{ij}_{n_1\ell_1m_1,n_2\ell_2m_2}]_{m_1,m_2}. \tag{1}$$

The block in Equation (1) measures the coupling between two orbitals, and the complete operator matrix O is assembled from all such orbital blocks $\{\boldsymbol{O}^{ij}_{n_1\ell_1,n_2\ell_2}\}$.

An Equivariant Model Framework of Quantum Operators. Early works like SchNOrb (Schütt et al., 2019) and DeepH (Li et al., 2022) learn to predict the quantum operator or its orbital blocks directly, which, however, cannot guarantee $\mathrm{O}(3)$ equivariance in general. Later equivariant Hamiltonian models use coupled irreps and predict orbital blocks through a *Wigner–Eckart output layer* (Gong et al., 2023; Yu et al., 2023; Zhong et al., 2023; Yu et al., 2026; Zhouyin et al., 2025). In particular, denote an $\mathrm{O}(3)$ irrep as $V_{\ell,p}$, which essentially is a representation space having dimension

[1] Please refer to the work in (Qian et al., 2026) and Appendix C for a broader discussion and more evidence.

$2\ell+1$ and determined by angular degree ℓ and inversion parity $p \in \{\pm 1\}$. The *Clebsch–Gordan (CG) tensor product* $V_{\ell_1,p_1} \otimes V_{\ell_2,p_2}$ decomposes into irreps V_{ℓ_3,p_3} with $|\ell_1-\ell_2| \le \ell_3 \le \ell_1+\ell_2$ and $p_3 = p_1 p_2$. For each orbital block $\boldsymbol{O}^{ij}_{n_1\ell_1,n_2\ell_2}$ in Equation (1), the Wigner–Eckart output layer predicts each of its elements as

$$\widehat{O}^{ij}_{n_1\ell_1 m_1,n_2\ell_2 m_2} = \sum\nolimits_{\ell_3=|\ell_1-\ell_2|}^{\ell_1+\ell_2} \sum\nolimits_{m_3=-\ell_3}^{\ell_3} c^{\ell_3,m_3}_{\ell_1,m_1,\ell_2,m_2} \hat{o}^{ij}_{\ell_3,p_3,m_3}. \tag{2}$$

Here, $\{c^{\ell_3,m_3}_{\ell_1,m_1,\ell_2,m_2}\}_{\ell_3,m_3}$ are fixed Clebsch–Gordan coefficients, which specify how the orbital indices transform, and $\hat{\boldsymbol{o}}^{ij}_{\ell_3,p_3} = [\hat{o}^{ij}_{\ell_3,p_3,m_3}] \in V_{\ell_3,p_3}$ is the *coupled component* parameterized by neural networks. Because Equation (2) is a fixed equivariant change of basis, predicting $\mathrm{O}(3)$-equivariant coupled components yields an orbital block satisfying $\mathrm{O}(3)$ equivariance. As shown in Appendices A.1–A.3, recent studies improve this framework through self-consistency (Zhang et al., 2024; Wang et al., 2024), adaptive sparsity (Luo et al., 2025), nonlinear equivariant representations (Yin et al., 2025), and flow matching (Kim et al., 2025) and extend it to various application scenarios (Tang et al., 2024; Qian et al., 2026; Li et al., 2025b; Kaniselvan et al., 2026; Yin et al., 2026; Zhong et al., 2026; Schwade et al., 2026; Kim et al., 2026).

Equivariant Message Passing. In the above framework, the coupled components (i.e., the $\hat{o}$'s in Equation (2)) are often parameterized by various message-passing neural networks (MPNNs). In general, MPNNs represent an atomic structure as a graph and update each atomic feature by aggregating messages from its neighbors (Gilmer et al., 2017; Schütt et al., 2018). Directional MPNNs incorporate angular and spherical geometric information (Gasteiger et al., 2020a;b; 2021; Liu et al., 2022). Equivariant MPNNs organize geometric features into irreps and preserve their transformation laws (Thomas et al., 2018; Fuchs et al., 2020; Schütt et al., 2021; Batzner et al., 2022; Brandstetter et al., 2022; Geiger & Smidt, 2022; Batatia et al., 2022). Recent equivariant Transformers combine these representations with attention mechanisms (Liao & Smidt, 2023; Liao et al., 2024; Li et al., 2025a; Liao et al., 2026; Huang et al., 2026). For the high-degree features required by quantum operators, conventional full CG tensor products have an angular cost of $\mathcal{O}(\ell^6_{\max})$.

Edge-aligned Methods. To improve computational efficiency of equivariant message passing, edge-aligned methods rotate each edge direction to the reference axis, e.g., $\boldsymbol{e}_z = (0,0,1)^\top$, and perform the angular operation in the corresponding local frame (Passaro & Zitnick, 2023; Liao et al., 2024). After alignment, rotations about $\boldsymbol{e}_z$ form the $\mathrm{SO}(2)$ group. Restricting the irrep $V_{\ell,p}$ to this group reduces the computational complexity to $\mathcal{O}(\ell^3_{\max})$. Following this strategy, recent models use axis-aligned sparsification or edge-local $\mathrm{SO}(2)$ convolutions to improve efficiency in atomistic modeling (Huang et al., 2026; Li et al., 2026). Typically, QHNetV2 (Yu et al., 2026) performs feature updates in $\mathrm{SO}(2)$ local frames, and SLEM (Zhouyin et al., 2025) combines $\mathrm{SO}(2)$ convolutions with strictly localized message passing. Details are provided in Appendices A.2–A.5.

3 Proposed Model

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote an atomistic graph with atom set $\mathcal{V}$ and directed edge set $\mathcal{E}$. Each node $i \in \mathcal{V}$ carries atomic type z_i and Cartesian position $\boldsymbol{r}_i$. Each directed edge $i \to j$ carries the relative vector $\boldsymbol{r}_{ij} = \boldsymbol{r}_j - \boldsymbol{r}_i$, its length $r_{ij} = \|\boldsymbol{r}_{ij}\|_2$, and its unit direction $\hat{\boldsymbol{r}}_{ij} = \boldsymbol{r}_{ij}/r_{ij}$. Let $\mathcal{N}(i)$ denote the neighbors of atom i within the cutoff radius. Denote $\boldsymbol{x}_{ij,\ell_{\text{in}}}$ as an input feature for edge $i \to j$ of angular degree ℓ_{in}, which may collect node, edge, or hidden features. The angular part of rotation-equivariant message passing has the following general form (Passaro & Zitnick, 2023):

$$\boldsymbol{m}_{i,\ell_{\text{out}}} = \sum\nolimits_{j\in\mathcal{N}(i)} \sum\nolimits_{\ell_{\text{in}},\ell_f} w^{\ell_{\text{out}}}_{\ell_{\text{in}},\ell_f}(\boldsymbol{z}_{ij}) [\underbrace{\boldsymbol{x}_{ij,\ell_{\text{in}}} \otimes \boldsymbol{Y}^{(\ell_f)}(\hat{\boldsymbol{r}}_{ij})}_{\text{CG tensor product}}]_{\ell_{\text{out}}}, \tag{3}$$

where ℓ_f is the filter degree, and $\boldsymbol{m}_{i,\ell_{\text{out}}}$ is the aggregated message of output degree ℓ_{out}. The angular degrees satisfy $|\ell_{\text{in}} - \ell_f| \le \ell_{\text{out}} \le \ell_{\text{in}} + \ell_f$, and channel indices are suppressed. The learnable function w takes rotation-invariant pair embeddings $\boldsymbol{z}_{ij}$, which encode r_{ij}, z_i, and z_j and may include learned scalar features. The fixed real spherical harmonics $\boldsymbol{Y}^{(\ell_f)}(\hat{\boldsymbol{r}}_{ij})$ encode the edge direction as a vector with $2\ell_f+1$ components that transforms under the degree-ℓ_f irrep of $\mathrm{SO}(3)$.

In practice, the general form in Equation (3) can be implemented by various equivariant MPNNs (EquiMPNNs), e.g., TFN (Thomas et al., 2018), NequIP (Batzner et al., 2022), EquiformerV2 (Liao

et al., 2024), and recent SLEM (Zhouyin et al., 2025). Typically, given a graph $\mathcal{G}$, an EquiMPNN uses T message-passing layers to update equivariant node features $\{\boldsymbol{n}_i\}_{i\in\mathcal{V}}$, edge features $\{\boldsymbol{e}_{ij}\}_{i\to j\in\mathcal{E}}$, hidden features $\{\boldsymbol{h}_{ij}\}_{i\to j\in\mathcal{E}}$, and invariant pair embeddings $\{\boldsymbol{z}_{ij}\}_{i\to j\in\mathcal{E}}$. For $t=1,\dots,T$,

$$\{\boldsymbol{n}_i^{(t)}\}_{i\in\mathcal{V}},\{\boldsymbol{e}_{ij}^{(t)},\boldsymbol{h}_{ij}^{(t)},\boldsymbol{z}_{ij}^{(t)}\}_{i\to j\in\mathcal{E}}=\mathrm{EquiMPNN}(\{\boldsymbol{n}_i^{(t-1)}\}_{i\in\mathcal{V}},\{\boldsymbol{e}_{ij}^{(t-1)},\boldsymbol{h}_{ij}^{(t-1)},\boldsymbol{z}_{ij}^{(t-1)}\}_{i\to j\in\mathcal{E}}). \quad (4)$$

Finally, we obtain the coupled components by $\hat{\boldsymbol{o}}^{ii}=\mathrm{Readout}(\boldsymbol{n}_i^{(T)})$ and $\hat{\boldsymbol{o}}^{ij}=\mathrm{Readout}(\boldsymbol{e}_{ij}^{(T)})$ and estimate orbital blocks using Equation (2), denoted as $\{\widehat{\boldsymbol{O}}^{ij}_{n_1\ell_1,n_2\ell_2}[\mathcal{G}]\}$.

Pursuing $\mathrm{O}(3)$-Equivariance. In this study, we require the predicted blocks to be equivariant under the group $\mathrm{O}(3)$ (containing rotations, reflections, and spatial inversion), i.e.,

$$\widehat{\boldsymbol{O}}^{ij}_{n_1\ell_1,n_2\ell_2}[\boldsymbol{Q}\mathcal{G}]=\boldsymbol{P}_{\ell_1,p_1}(\boldsymbol{Q})\widehat{\boldsymbol{O}}^{ij}_{n_1\ell_1,n_2\ell_2}[\mathcal{G}]\boldsymbol{P}_{\ell_2,p_2}(\boldsymbol{Q})^\top,\quad \text{for } \boldsymbol{Q}\in\mathrm{O}(3), \quad (5)$$

where $\boldsymbol{Q}\mathcal{G}$ is the graph transformed by replacing each $\boldsymbol{r}_i$ with $\boldsymbol{Q}\boldsymbol{r}_i$, $p_k=(-1)^{\ell_k}$ for $k\in\{1,2\}$, and $\boldsymbol{P}_{\ell,p}(\boldsymbol{Q})$ is the representation matrix of the transformation $\boldsymbol{Q}$, acting on the $\mathrm{O}(3)$ irrep $V_{\ell,p}$.

Equation (5) requires prediction after transforming the structure to agree with transforming the predicted matrix in the orbital basis. To achieve this aim with a mild computational cost, we propose the following edge-aligned $\mathrm{O}(2)$ convolution.

3.1 Edge-Aligned $\mathrm{O}(2)$ Convolution

Edge-Aligned Convolution Kernel. Let $\boldsymbol{u}=\hat{\boldsymbol{r}}_{ij}\in\mathbb{S}^2$ be the unit direction of edge $i\to j$, where $\mathbb{S}^2$ is the unit sphere in $\mathbb{R}^3$. We would like to design a convolution for the local frame associated with each $\boldsymbol{u}$. Let V_{in} and V_{out} denote the input and output spaces of the convolution, and $\boldsymbol{P}_{\mathrm{in}}(\boldsymbol{Q})$ and $\boldsymbol{P}_{\mathrm{out}}(\boldsymbol{Q})$ be their representation matrices for $\boldsymbol{Q}\in\mathrm{O}(3)$. For $a\in\{\mathrm{in},\mathrm{out}\}$, V_a is a direct sum of irreps up to angular degree $\ell_{\max}$, and $\boldsymbol{P}_a(\boldsymbol{Q})$ is a block diagonal matrix constructed by $\boldsymbol{P}_{\ell,p}(\boldsymbol{Q})$:

$$V_a=\bigoplus_{\ell=0}^{\ell_{\max}}\bigoplus_{p\in\{\pm1\}}\bigoplus_{c=1}^{C^a_{\ell,p}}V_{\ell,p},\quad \boldsymbol{P}_a(\boldsymbol{Q})=\mathrm{diag}(\{\mathrm{diag}(\underbrace{\boldsymbol{P}_{\ell,p}(\boldsymbol{Q}),\dots,\boldsymbol{P}_{\ell,p}(\boldsymbol{Q})}_{\text{Copy } C^a_{\ell,p} \text{ times}})\}_{\ell,p}). \quad (6)$$

Here, $C^a_{\ell,p}$ is the number of channels of type (ℓ,p) in V_a. Each channel c contributes one copy of $V_{\ell,p}$, and $C^a_{\ell,p}=0$ for unused types.

Let $\boldsymbol{e}_z=(0,0,1)^\top$ be the fixed reference direction. We choose a rotation $\boldsymbol{F}_{\boldsymbol{u}}\in\mathrm{SO}(3)$ satisfying $\boldsymbol{F}_{\boldsymbol{u}}\boldsymbol{e}_z=\boldsymbol{u}$. For a convolution kernel $\boldsymbol{K}_{\boldsymbol{e}_z}:V_{\mathrm{in}}\to V_{\mathrm{out}}$ defined along the reference direction, EACO constructs the direction-dependent kernel $\boldsymbol{K}_{\boldsymbol{u}}:V_{\mathrm{in}}\to V_{\mathrm{out}}$ as

$$\boldsymbol{K}_{\boldsymbol{u}}=\boldsymbol{P}_{\mathrm{out}}(\boldsymbol{F}_{\boldsymbol{u}})\boldsymbol{K}_{\boldsymbol{e}_z}\boldsymbol{P}_{\mathrm{in}}(\boldsymbol{F}_{\boldsymbol{u}})^{-1}. \quad (7)$$

Here, $\boldsymbol{P}_{\mathrm{in}}(\boldsymbol{F}_{\boldsymbol{u}})^{-1}$ transforms the input into the edge-aligned frame, $\boldsymbol{K}_{\boldsymbol{e}_z}$ applies the kernel at $\boldsymbol{e}_z$, and $\boldsymbol{P}_{\mathrm{out}}(\boldsymbol{F}_{\boldsymbol{u}})$ transforms the output back to the global frame.

Let $\mathcal{S}_{\boldsymbol{e}_z}=\{\boldsymbol{Q}\in\mathrm{O}(3)\mid\boldsymbol{Q}\boldsymbol{e}_z=\boldsymbol{e}_z\}\cong\mathrm{O}(2)$ denote the *stabilizer* of the reference direction. For any $\boldsymbol{Q}\in\mathrm{O}(3)$, the residual transformation between the frames aligned with $\boldsymbol{u}$ and $\boldsymbol{Q}\boldsymbol{u}$ is

$$\boldsymbol{g}(\boldsymbol{Q},\boldsymbol{u})=\boldsymbol{F}_{\boldsymbol{Q}\boldsymbol{u}}^{-1}\boldsymbol{Q}\boldsymbol{F}_{\boldsymbol{u}}\in\mathcal{S}_{\boldsymbol{e}_z}\cong\mathrm{O}(2),\quad \text{with }\ \det\boldsymbol{g}(\boldsymbol{Q},\boldsymbol{u})=\det\boldsymbol{Q}. \quad (8)$$

Therefore, a proper rotation ($\det\boldsymbol{Q}=1$) induces a local $\mathrm{SO}(2)$ rotation, whereas a transformation with $\det\boldsymbol{Q}=-1$ induces a reflection of the local $\mathrm{O}(2)$ stabilizer. This observation gives the following equivalence between the local kernel constraint and the global transformation law.

Theorem 1 (Local and global equivariance). *Let $\boldsymbol{I}_d$ be the $d\times d$ identity matrix. Assume $\boldsymbol{F}_{\boldsymbol{e}_z}=\boldsymbol{I}_3$ and define $\boldsymbol{K}_{\boldsymbol{u}}$ by Equation (7). Then*

$$\boldsymbol{K}_{\boldsymbol{Q}\boldsymbol{u}}\boldsymbol{P}_{\mathrm{in}}(\boldsymbol{Q})=\boldsymbol{P}_{\mathrm{out}}(\boldsymbol{Q})\boldsymbol{K}_{\boldsymbol{u}}\quad \textit{for every } \boldsymbol{Q}\in\mathrm{O}(3),\ \boldsymbol{u}\in\mathbb{S}^2, \quad (9)$$

if and only if

$$\boldsymbol{P}_{\mathrm{out}}(\boldsymbol{g})\boldsymbol{K}_{\boldsymbol{e}_z}=\boldsymbol{K}_{\boldsymbol{e}_z}\boldsymbol{P}_{\mathrm{in}}(\boldsymbol{g})\quad \textit{for every } \boldsymbol{g}\in\mathcal{S}_{\boldsymbol{e}_z}\cong\mathrm{O}(2). \quad (10)$$

Indeed, substituting $\boldsymbol{Q}\boldsymbol{F}_{\boldsymbol{u}}=\boldsymbol{F}_{\boldsymbol{Q}\boldsymbol{u}}\boldsymbol{g}(\boldsymbol{Q},\boldsymbol{u})$ into Equation (7) reduces Equation (9) to Equation (10). Conversely, setting $\boldsymbol{u}=\boldsymbol{e}_z$ and $\boldsymbol{Q}=\boldsymbol{g}\in\mathcal{S}_{\boldsymbol{e}_z}$ recovers the local constraint. The complete proof, including independence from the choice of alignment frame, is provided in Appendix B.2.

Local Kernel Parameterization. The key point of EACO is parametrizing the local kernel $\boldsymbol{K}_{\boldsymbol{e}_z}$ in an efficient way and making it satisfy Theorem 1. Given an $\mathrm{O}(3)$ irrep $V_{\ell,p}$, we denote W_s, $0 \le s \le \ell$ as its s-th subspace associated with magnetic orders $m = \pm s$. We choose $\boldsymbol{Q}_{\boldsymbol{e}_z} = \mathrm{diag}(1, -1, 1)$, the reflection through the xz-plane, as a generating reflection in $\mathcal{S}_{\boldsymbol{e}_z}$. Its action on W_s is

$$\boldsymbol{P}_{\ell,p}(\boldsymbol{Q}_{\boldsymbol{e}_z})\big|_{W_s} = \tau_{\ell,p}\boldsymbol{\Sigma}_s, \quad \text{with } \tau_{\ell,p} = p(-1)^\ell,\ \boldsymbol{\Sigma}_s = \begin{cases} 1 & s = 0, \\ \mathrm{diag}(1,-1) & s > 0. \end{cases} \tag{11}$$

The sign $\tau_{\ell,p}$ retains the parity information that identifies reflections. As a result, a polar irrep $V_{1,-1}$ and an axial irrep $V_{1,1}$ have the same $\mathrm{SO}(2)$ rotation action, but their local reflection actions are $\boldsymbol{\Sigma}_1$ and $-\boldsymbol{\Sigma}_1$, respectively, making them distinguishable.

For $s > 0$, let $C_{\mathrm{in}}(s)$ and $C_{\mathrm{out}}(s)$ count the input and output channels containing magnetic orders $\pm s$. An arbitrary real $\mathrm{SO}(2)$-equivariant block has the form $\boldsymbol{A}_s \otimes \boldsymbol{I}_2 + \boldsymbol{B}_s \otimes \boldsymbol{J}_2$. Here, $\boldsymbol{A}_s, \boldsymbol{B}_s \in \mathbb{R}^{C_{\mathrm{out}}(s) \times C_{\mathrm{in}}(s)}$ are learnable channel-mixing matrices, $\boldsymbol{J}_2 = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$, and $\otimes$ denotes the Kronecker product in this expression. For every rotation $\boldsymbol{R}$ about $\boldsymbol{e}_z$, both $\boldsymbol{I}_2$ and $\boldsymbol{J}_2$ commute with $\boldsymbol{P}_{\ell,p}(\boldsymbol{R})|_{W_s}$.

Let $\boldsymbol{\tau}_s^{\mathrm{in}} = [\tau_{s,a}^{\mathrm{in}}] \in \{\pm 1\}^{C_{\mathrm{in}}(s)}$ and $\boldsymbol{\tau}_s^{\mathrm{out}} = [\tau_{s,b}^{\mathrm{out}}] \in \{\pm 1\}^{C_{\mathrm{out}}(s)}$ collect the reflection signs of all input and output channels at order s. For input channel a and output channel b, we define two binary mask matrices $\boldsymbol{M}_s^+$ and $\boldsymbol{M}_s^-$, whose elements are

$$[\boldsymbol{M}_s^+]_{ba} = \frac{1}{2}\left(1 + \tau_{s,b}^{\mathrm{out}}\tau_{s,a}^{\mathrm{in}}\right) \quad \text{and} \quad [\boldsymbol{M}_s^-]_{ba} = \frac{1}{2}\left(1 - \tau_{s,b}^{\mathrm{out}}\tau_{s,a}^{\mathrm{in}}\right). \tag{12}$$

The fixed mask $\boldsymbol{M}_s^+$ selects channel pairs with the same reflection sign, while $\boldsymbol{M}_s^-$ selects pairs with opposite reflection signs. The full reflection constraint then gives the EACO block

$$\boldsymbol{K}_{\boldsymbol{e}_z}(s) = (\boldsymbol{A}_s \odot \boldsymbol{M}_s^+) \otimes \boldsymbol{I}_2 + (\boldsymbol{B}_s \odot \boldsymbol{M}_s^-) \otimes \boldsymbol{J}_2, \qquad 0 < s \le \ell_{\max}. \tag{13}$$

Here, $\odot$ denotes the Hadamard product.

Proposition 1 (Characterization of the local kernel)**.** *For $s > 0$, Equation (13) gives all linear intertwiners[2] between the restricted input and output $\mathrm{O}(2)$ representations. Channels with the same reflection sign couple through $\boldsymbol{I}_2$, while channels with opposite reflection signs couple through $\boldsymbol{J}_2$. At $s = 0$, a weight is allowed only between channels with the same reflection sign, and a bias is allowed only for output channels with reflection sign $+1$.*

Given an input feature $\boldsymbol{x} \in V_{\mathrm{in}}$ and a unit edge direction, EACO computes the output $\boldsymbol{y}$ by the procedure in Algorithm 1. It uses the same "align–convolve–rotate" pipeline as an edge-aligned $\mathrm{SO}(2)$ convolution (Passaro & Zitnick, 2023; Zhouyin et al., 2025), with binary masks determined by the reflection signs in Equation (12). Thus, each input–output channel pair at $s > 0$ retains exactly one of the two $\mathrm{SO}(2)$-equivariant coefficients, removing the component that violates reflection equivariance. Appendices B.4 and B.5 provide the derivation and its relation to parity-allowed Clebsch–Gordan couplings.

Algorithm 1 Edge-Aligned $\mathrm{O}(2)$ Convolution

Require: $\boldsymbol{x}$, $\boldsymbol{u}$, V_{in}, V_{out}, weights and biases
1: Choose $\boldsymbol{F}_{\boldsymbol{u}}\boldsymbol{e}_z = \boldsymbol{u}$, and $\bar{\boldsymbol{x}} \leftarrow \boldsymbol{P}_{\mathrm{in}}(\boldsymbol{F}_{\boldsymbol{u}})^{-1}\boldsymbol{x}$.
2: **For** $s = 0, \dots, \ell_{\max}$ **do**
3: Get channels $\bar{\boldsymbol{x}}_s$ with magnetic orders $\pm s$.
4: **If** $s = 0$
5: $\bar{\boldsymbol{y}}_0 \leftarrow$ equal-sign coupling, bias for $\tau = +1$.
6: **else**
7: Compute $\boldsymbol{M}_s^{\pm}$ by Equation (12).
8: $\bar{\boldsymbol{y}}_s \leftarrow \boldsymbol{K}_{\boldsymbol{e}_z}(s)\bar{\boldsymbol{x}}_s$ by Equation (13).
Ensure: $\boldsymbol{y} \leftarrow \boldsymbol{P}_{\mathrm{out}}(\boldsymbol{F}_{\boldsymbol{u}})(\bigoplus_{s=0}^{\ell_{\max}} \bar{\boldsymbol{y}}_s)$

Advantages over Existing Equivariant Angular Operations. Table 1 summarizes the resulting differences from existing equivariant angular operations. EACO extends the stabilizer constraint from $\mathrm{SO}(2)$ to $\mathrm{O}(2)$ while retaining the angular cost and strict locality of SLEM (Zhouyin et al., 2025). It has advantages in the following aspects.

- **Equivariance and Strictness of Locality.** Beyond $\mathrm{SO}(2)$ kernels in (Passaro & Zitnick, 2023; Zhouyin et al., 2025), EACO enforces the full local $\mathrm{O}(2)$ constraint and guarantees global $\mathrm{O}(3)$ equivariance by Theorem 1. It also retains SLEM's fixed node neighborhoods and combines the two endpoint neighborhoods for each edge. Its receptive fields therefore remain independent of network depth, unlike those of DeepH-E3 (Gong et al., 2023) and MACE-H (Qian et al., 2026).

[2] A linear map that commutes with the group actions is called an intertwiner and is equivariant.

Table 1: **Comparison of equivariant angular operations.** The last column counts independent coefficients per input–output channel pair at magnetic orders $\pm s$ with $s > 0$.

Method	Stabilizer	Equivariance	Locality	Angular cost	#Coeffs.
Full $\mathrm{O}(3)$ Tensor Product	–	$\mathrm{O}(3)$	model-dependent	$\mathcal{O}(\ell_{\max}^6)$	–
DeepH-E3 (Gong et al., 2023)	–	$\mathrm{O}(3)$	depth-dependent	$\mathcal{O}(\ell_{\max}^6)$	–
MACE-H (Qian et al., 2026)	–	$\mathrm{O}(3)$	depth-dependent	$\mathcal{O}(\ell_{\max}^6)$	–
eSCN (Passaro & Zitnick, 2023)	$\mathrm{SO}(2)$	$\mathrm{SO}(3)$	depth-dependent	$\mathcal{O}(\ell_{\max}^3)$	2
SLEM (Zhouyin et al., 2025)	$\mathrm{SO}(2)$	$\mathrm{SO}(3)$	strict	$\mathcal{O}(\ell_{\max}^3)$	2
EACO	$\mathrm{O}(2)$	$\mathrm{O}(3)$	strict	$\mathcal{O}(\ell_{\max}^3)$	**1**

Table 2: **Numerical comparisons of various methods on five materials.** Bold marks the lowest Hamiltonian MAE for each material. Light gray highlights violations of $\mathrm{O}(3)$ equivariance.

Material	Method	Model Size (M)↓	Hamiltonian MAE (meV)↓ All	On-site	Off-site	Equivariance error↓ $\Delta_{\mathrm{imp}}H$	$\epsilon_{H,\mathrm{imp}}^{\max}$
Si	DeepH-E3	0.628	0.143	0.650	0.139	5.16×10^{-5}	4.41×10^{-6}
	MACE-H	1.184	0.114	0.588	0.110	3.21×10^{-5}	7.31×10^{-7}
	SLEM	0.689	0.069	0.366	0.067	3.65×10^{-2}	2.97×10^{-4}
	EACO	0.567	**0.054**	**0.250**	**0.052**	9.13×10^{-6}	1.96×10^{-7}
MoS_2	DeepH-E3	1.340	0.571	1.286	0.558	7.45×10^{-5}	3.16×10^{-7}
	MACE-H	2.794	**0.408**	0.947	**0.399**	5.97×10^{-5}	7.14×10^{-7}
	SLEM	0.467	0.432	**0.287**	0.435	3.68×10^{-1}	1.03×10^{-3}
	EACO	0.384	0.430	0.330	0.432	6.57×10^{-5}	9.63×10^{-7}
Graphene	DeepH-E3	0.628	0.360	0.721	0.352	8.12×10^{-5}	1.03×10^{-6}
	MACE-H	2.350	0.258	0.838	0.245	5.17×10^{-5}	1.11×10^{-6}
	SLEM	0.433	0.227	**0.628**	0.218	1.67×10^{-1}	6.05×10^{-4}
	EACO	0.358	**0.196**	0.665	**0.186**	5.55×10^{-5}	1.82×10^{-7}
GaN	DeepH-E3	1.903	1.244	30.77	1.161	6.13×10^{-5}	6.11×10^{-6}
	MACE-H	1.399	1.252	30.18	1.170	3.76×10^{-5}	1.26×10^{-6}
	SLEM	2.384	0.132	0.324	0.132	9.74×10^{-2}	6.89×10^{-4}
	EACO	1.762	**0.126**	**0.268**	**0.125**	1.61×10^{-5}	9.73×10^{-7}
HfO_2	DeepH-E3	1.887	1.693	87.25	1.284	6.09×10^{-5}	3.90×10^{-6}
	MACE-H	1.566	1.604	87.15	1.196	8.09×10^{-5}	3.12×10^{-6}
	SLEM	2.127	0.392	2.096	0.384	2.43×10^{-1}	7.19×10^{-4}
	EACO	1.497	**0.326**	**1.674**	**0.319**	1.86×10^{-5}	5.04×10^{-7}

- **Computational and Parameter Efficiency.** EACO groups all degree-ℓ features containing the magnetic orders $\pm s$ and applies one channel-mixing operation to each group. If the number of channels per degree is bounded, both $C_{\mathrm{in}}(s)$ and $C_{\mathrm{out}}(s)$ scale as $\mathcal{O}(\ell_{\max} - s + 1)$. The channel-mixing cost is $\sum_{s=0}^{\ell_{\max}} C_{\mathrm{in}}(s)C_{\mathrm{out}}(s) = \mathcal{O}(\ell_{\max}^3)$, which is comparable to that of eSCN (Passaro & Zitnick, 2023) and SLEM (Zhouyin et al., 2025). For T layers and $|\mathcal{E}|$ directed edges, these operations cost $\mathcal{O}(T|\mathcal{E}|\ell_{\max}^3)$ up to channel factors. In addition, for $s > 0$, an $\mathrm{SO}(2)$ block contains two dense matrices, whereas EACO applies one matrix to store the coefficients of input–output channel pairs. Detailed analysis is provided in Appendix B.8.

3.2 EACO-Based $\mathrm{O}(3)$-Equivariant MPNN

As illustrated in Figure 1, plugging EACO into an existing EquiMPNN architecture (e.g., the SLEM model (Zhouyin et al., 2025) shown in Appendix B.6) leads to an $\mathrm{O}(3)$-equivariant MPNN. The model first constructs the cutoff graph and initializes the pair, hidden, edge, and node features. Then, it applies Equation (4) for T layers and uses EACO in every direction-dependent linear map. Finally, the node and edge readouts predict coupled operator components, which are converted to orbital blocks by Equation (2). The full forward procedure can be found in Appendix B.6.4.

For the above EquiMPNN, EACO, equivariant linear maps, gated nonlinearities, residual connections, and sum aggregation are equivariant, while the Clebsch–Gordan output is a fixed intertwiner. In addition, atomic types, distances, radial weights, and cutoff coefficients are invariant. Therefore, we have the following corollary, whose proof is in Appendix B.7.

Corollary 1 ($\mathrm{O}(3)$ equivariance of the EACO-based MPNN)**.** *Let all direction-dependent linear maps of a typical EquiMPNN use the convolution kernels characterized by Proposition 1. Then the coupled node and edge outputs are $\mathrm{O}(3)$-equivariant, and the reconstructed quantum operator satisfies Equation (5) for every $\boldsymbol{Q} \in \mathrm{O}(3)$.*

4 Experiments

To demonstrate the superiority of EACO, we first evaluate its Hamiltonian prediction accuracy and $\mathrm{O}(3)$ equivariance across Si, MoS_2, Graphene, GaN, and HfO_2. We then study mirror-symmetry selection rules in graphene to verify the physical effectiveness of EACO. Finally, we examine the role of data augmentation, the cost of angular operations, and training convergence on graphene. Representative results are shown below, and more experimental details are in Appendix D.

4.1 Quantum Operator Prediction

Datasets and Baselines. The five material datasets are from (Zhouyin et al., 2025). Each provides an atomic geometry and DFT Hamiltonian blocks in a localized atomic-orbital basis. For the bulk datasets, Si, GaN, and HfO_2, each contains 150 training and 30 validation structures. The monolayer datasets have training/validation/test splits of 270/90/90 for graphene and 300/100/100 for MoS_2. These materials span different orbital bases, including g orbitals in HfO_2. For each dataset, we compare EACO with SLEM (Zhouyin et al., 2025), DeepH-E3 (Gong et al., 2023), and MACE-H (Qian et al., 2026). SLEM and EACO use four layers, with matched channel widths, and are trained for 1,000 epochs with Adam and batch size one. We use $\mathcal{L}_{\mathrm{H}} = \frac{1}{2}(d_{\mathrm{on}} + d_{\mathrm{off}})$, where d_{on} and d_{off} are defined over the valid on-site and off-site entries, respectively, and each of them is the sum of the MAE and RMSE losses. DeepH-E3 and MACE-H use their native Hamiltonian losses and optimization schedules for up to 1,500 epochs. All methods use the same data splits, orbital bases, and cutoffs. For each run, we select the checkpoint with the lowest validation loss.

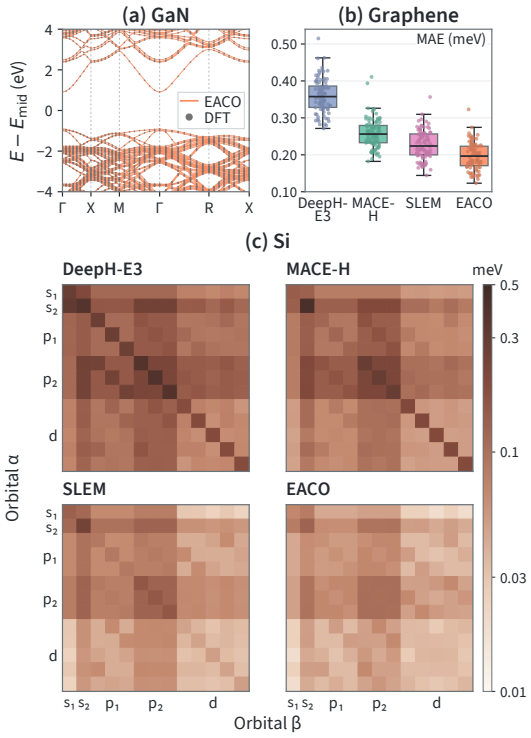

Figure 2: **Hamiltonian prediction results on GaN, graphene, and Si.** (a) GaN band structures from EACO (orange curves) and DFT (gray dots) along the supercell path. (b) The boxplot of Hamiltonian MAEs for 90 graphene validation structures per method. (c) Si orbital-pair MAEs, averaged over atom-pair blocks in 30 validation structures, with a shared logarithmic color scale.

Evaluation. For each method, we report i) its model size (the number of parameters in millions), ii) its MAE (meV) over valid Hamiltonian matrix entries for on-site, off-site, and all blocks, and iii) its equivariance error under inversion, reflections, and a fixed rotation composed with inversion. For the predicted Hamiltonian $\widehat{\boldsymbol{H}}$ and a transformation $\boldsymbol{Q} \in \mathrm{O}(3)$ with $\det \boldsymbol{Q} = -1$, we define

$$\boldsymbol{\delta}_{H,\boldsymbol{Q}}(\mathcal{G}) = \widehat{\boldsymbol{H}}[\boldsymbol{Q}\mathcal{G}] - \boldsymbol{D}(\boldsymbol{Q})\widehat{\boldsymbol{H}}[\mathcal{G}]\boldsymbol{D}(\boldsymbol{Q})^{\top}, \quad \epsilon_{H,\boldsymbol{Q}}(\mathcal{G}) = \|\boldsymbol{\delta}_{H,\boldsymbol{Q}}(\mathcal{G})\|_{F,\mathrm{valid}} \,/\, \|\widehat{\boldsymbol{H}}[\mathcal{G}]\|_{F,\mathrm{valid}}. \quad (14)$$

Here, $\boldsymbol{D}(\boldsymbol{Q})$ is the orbital representation of $\boldsymbol{Q}$, and the norm $\|\cdot\|_{F,\mathrm{valid}}$ excludes padded entries. Given the evaluated structure–transformation pairs $\mathcal{D} = \{(\mathcal{G}, \boldsymbol{Q})\}$, we report the mean absolute

Table 3: **Mirror equivariance and electron–phonon coupling in graphene.** Each mode uses eight displaced structure pairs for mirror residuals and a primitive-cell q grid with shared DFT states for EPC. g_{f} and g_{a} denote forbidden and allowed channels. Insets show cropped periodic graphene patches and the prescribed frozen displacements for the mirror test. Arrows show amplified out-of-plane displacements; those below 25% of the maximum amplitude are omitted in both panels.

Method	ZA			ZO		
	$\epsilon_{\sigma_h}^{\max}\downarrow$	g_{f}-RMS(meV)↓	g_{a}-Err(%)↓	$\epsilon_{\sigma_h}^{\max}\downarrow$	g_{f}-RMS(meV)↓	g_{a}-Err(%)↓
DeepH-E3	**0**	3.66×10^{-11}	0.280	**0**	1.97×10^{-11}	0.437
MACE-H	**0**	7.11×10^{-11}	0.260	**0**	2.47×10^{-11}	0.363
SLEM	1.47×10^{-4}	0.129	0.219	1.59×10^{-4}	0.324	0.220
EACO	**0**	$\mathbf{2.74\times10^{-11}}$	**0.100**	**0**	$\mathbf{1.04\times10^{-11}}$	**0.113**

residual $\Delta_{\mathrm{imp}}H$ in meV over all valid matrix entries pooled across $\mathcal{D}$ and the maximum relative residual $\epsilon_{H,\mathrm{imp}}^{\max} = \max_{\mathcal{G},\boldsymbol{Q}\in\mathcal{D}} \epsilon_{H,\boldsymbol{Q}}(\mathcal{G})$.

Table 2 shows that EACO outperforms the baselines on Si, graphene, GaN, and HfO_2 in overall Hamiltonian MAE. Compared with SLEM, EACO reduces overall MAE on all five materials with 17.2–29.6% fewer parameters. In particular, MAE decreases by 22.1% on Si, 13.5% on graphene, and 17.0% on HfO_2. EACO improves off-site predictions on all five materials and on-site predictions on the three bulk materials. On MoS_2, EACO achieves competitive performance with 0.384M parameters, compared with 1.340M for DeepH-E3 and 2.794M for MACE-H. Figure 2(a) shows that EACO closely reproduces the GaN valence and conduction bands along the supercell path. Figure 2(b) shows EACO has the lowest median MAE across validation structures of graphene. Figure 2(c) shows lower errors across Si's s, p, and d orbital groups.

Under improper transformations, SLEM has maximum relative errors of 10^{-4}–10^{-3}. Its absolute equivariance residuals are comparable in magnitude to its prediction MAEs, indicating a violation of $\mathrm{O}(3)$ equivariance. In contrast, EACO is $\mathrm{O}(3)$-equivariant: it reduces $\Delta_{\mathrm{imp}}H$ by three to four orders of magnitude compared with SLEM, and its maximum relative error stays below 10^{-6} on all five materials, with residuals comparable to those of DeepH-E3 and MACE-H.

4.2 Mirror-Symmetry Selection Rules in Graphene

We use graphene to test the physical effectiveness of EACO. For out-of-plane vibrations, mirror symmetry forbids first-order electron–phonon coupling (EPC) within the π sector but allows coupling between π and σ states (Mariani & von Oppen, 2008; Alidoosti et al., 2022). We compare the checkpoints in Table 2 for the acoustic (ZA) and optical (ZO) modes. We measure the mirror error, the RMS forbidden coupling, and the relative error in the allowed coupling against DFT. The EPC comparisons use common DFT electronic states and phonon data. Table 3 shows zero mirror residuals for EACO in both modes, matching DeepH-E3 and MACE-H. SLEM has residuals of 10^{-4} and forbidden couplings of 0.129meV (ZA) and 0.324meV (ZO). EACO suppresses both couplings below 3×10^{-11}meV. It means that **the $\mathrm{O}(3)$ equivariance matters for the selection rule.**

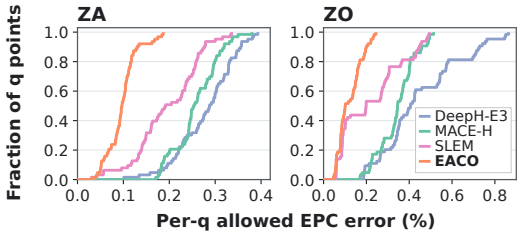

Figure 3: Cumulative distributions of per-q allowed EPC errors.

Notably, EACO achieves the lowest relative error in the allowed π–σ channel. Its errors are 0.100% for ZA and 0.113% for ZO. Compared with SLEM, EACO reduces these errors by 54.5% and 48.4%, respectively. Since the electronic states and phonon data are shared, these gains reflect a more accurate learned Hamiltonian response. In Figure 3, EACO's curves rise at smaller errors in both modes, showing that the improvement extends across the q grid. **EACO therefore preserves the selection rule while improving the prediction of the allowed coupling.**

4.3 Structural Constraints and Computational Efficiency

Figure 4 compares SLEM, SLEM with data augmentation, and EACO on graphene and Si. Specifically, the data augmentation is achieved by inverting each training structure with probability $1/2$ and transforming its cell and operator labels consistently. It lowers the MAE of SLEM from 0.227meV to 0.220meV on graphene and from 0.069meV to 0.056meV on Si. However, the errors under improper transformations remain at 10^{-4} for SLEM. EACO keeps both below 2×10^{-7}. This result means that the data augmentation cannot replace our local $\mathrm{O}(2)$ constraint.

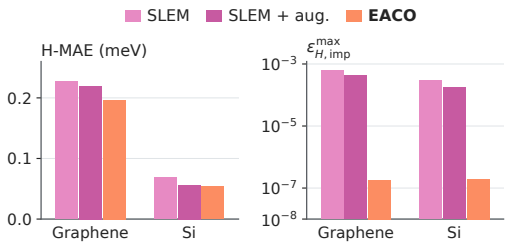

Figure 4: **Effect of inversion augmentation.** Hamiltonian MAE (left) and equivariance error (right) on graphene and Si.

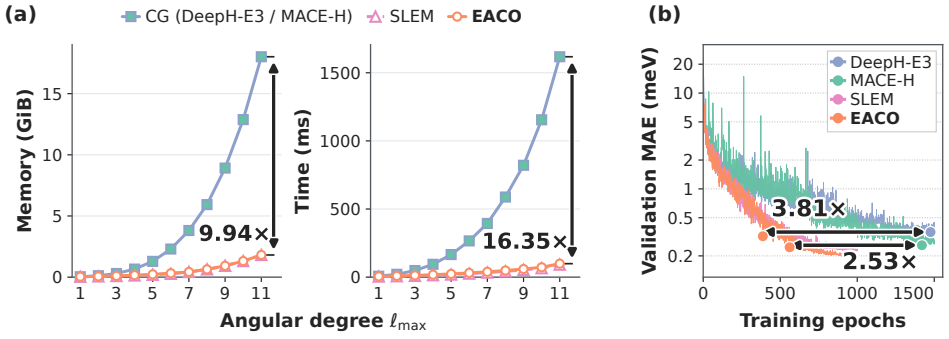

Figure 5: **Angular-operator efficiency and training convergence.** (a) Peak GPU memory and forward-plus-backward time as $\ell_{\max}$ increases, with 4,096 edges and 16 channels per degree and parity. Measurements use FP32 on one idle A100; points are medians of three independent runs. CG denotes the shared angular operator of DeepH-E3 and MACE-H. (b) Graphene validation MAE at every training epoch, without smoothing. Arrows compare the first epochs reaching each baseline's best recorded MAE; the ratios measure reductions in epochs.

Figure 5(a) shows that EACO retains SLEM's low angular-operator cost as $\ell_{\max}$ increases. At $\ell_{\max} = 11$, EACO reduces peak memory and forward-plus-backward time by factors of 9.94 and 16.35 relative to the CG operator used by DeepH-E3 and MACE-H. The comparison includes direction encoding, Wigner matrices, and feature rotations. Thus, the added reflection constraint preserves the computational advantage of the $O(\ell^3_{\max})$ local operator. Figure 5(b) shows that SLEM and EACO reach lower MAEs within 1000 epochs than DeepH-E3 and MACE-H achieve within 1,500 epochs. EACO matches the best recorded MAEs of DeepH-E3 and MACE-H in 387 and 561 epochs, respectively. The baselines reach those values at 1,476 and 1,420 epochs. These correspond to 3.81× and 2.53× reductions in epochs, and EACO finishes with the lowest MAE.

5 Conclusion and Future Work

In this study, we proposed EACO for efficient $\mathrm{O}(3)$-equivariant quantum operator prediction. We proved that local $\mathrm{O}(2)$ equivariance is necessary and sufficient for global $\mathrm{O}(3)$ equivariance in the edge-aligned construction. The local reflection constraint reduces kernel parameters while retaining strict locality and $\mathcal{O}(\ell^3_{\max})$ angular complexity. Experiments on five materials demonstrate its prediction accuracy and equivariance. The graphene study further shows that EACO preserves the mirror selection rule while improving allowed electron–phonon coupling predictions.

In future work, we will investigate explicit many-body feature interactions in the local $\mathrm{O}(2)$ frame. We aim to construct reflection-equivariant products of local features while keeping the spatial neighborhoods fixed. We will analyze how their cost grows with angular degree and body order. We will also study when higher body order improves data efficiency and when it limits transfer to new atomic environments.

AI USE STATEMENT

In this work, we used LLMs to refine the language of the paper, assist with proof checking, and support parts of the code implementation.

ETHICS STATEMENT

This work studies quantum operator prediction using public materials datasets and first-principles calculations. It does not involve human participants, personal data, or animal experiments. The method is intended to support computational materials research by reducing the cost of electronic-structure modeling. We acknowledge the sources of the datasets and baseline implementations used in the experiments.

REPRODUCIBILITY STATEMENT

Appendix B provides the model architecture, algorithms, and complete proofs of the theoretical results. Section 4 describes the experimental comparisons and evaluation metrics. Appendix D specifies the data splits, orbital bases, cutoffs, and model and training configurations.

REFERENCES

Mohammad Alidoosti, Davoud Nasr Esfahani, and Reza Asgari. σ_h symmetry and electron–phonon interaction in two-dimensional crystalline systems. *Physical Review B*, 106(4):045301, 2022.

Ilyes Batatia, Dávid Péter Kovács, Gregor N. C. Simm, Christoph Ortner, and Gábor Csányi. MACE: Higher order equivariant message passing neural networks for fast and accurate force fields. In *Advances in Neural Information Processing Systems*, volume 35, pp. 11423–11436. Curran Associates, Inc., 2022.

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.

Johannes Brandstetter, Rob Hesselink, Elise van der Pol, Erik J. Bekkers, and Max Welling. Geometric and physical quantities improve E(3) equivariant message passing. In *International Conference on Learning Representations*, 2022.

Kamal Choudhary, Kevin F. Garrity, Vinit Sharma, Adam J. Biacchi, Angela R. Hight Walker, and Francesca Tavazza. High-throughput density functional perturbation theory and machine learning predictions of infrared, piezoelectric, and dielectric responses. *npj Computational Materials*, 6 (1):64, 2020.

Maarten de Jong, Wei Chen, Henry Geerlings, Mark Asta, and Kristin Aslaug Persson. A database to enable discovery and design of piezoelectric materials. *Scientific Data*, 2(1):150053, 2015.

Fabian B. Fuchs, Daniel E. Worrall, Volker Fischer, and Max Welling. SE(3)-Transformers: 3D roto-translation equivariant attention networks. In *Advances in Neural Information Processing Systems*, volume 33, pp. 1970–1981. Curran Associates, Inc., 2020.

Johannes Gasteiger, Shankari Giri, Johannes T. Margraf, and Stephan Günnemann. Fast and uncertainty-aware directional message passing for non-equilibrium molecules. In *Machine Learning for Molecules Workshop at NeurIPS 2020*, 2020a.

Johannes Gasteiger, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs. In *International Conference on Learning Representations*, 2020b.

Johannes Gasteiger, Florian Becker, and Stephan Günnemann. GemNet: Universal directional graph neural networks for molecules. In *Advances in Neural Information Processing Systems*, volume 34, pp. 6790–6802. Curran Associates, Inc., 2021.

Mario Geiger and Tess Smidt. e3nn: Euclidean neural networks. *arXiv preprint arXiv:2207.09453*, 2022.

Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1263–1272. PMLR, 2017.

Xiaoxun Gong, He Li, Nianlong Zou, Runzhang Xu, Wenhui Duan, and Yong Xu. General framework for E(3)-equivariant neural network representation of density functional theory Hamiltonian. *Nature Communications*, 14(1):2848, 2023.

P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Physical Review*, 136(3B):B864–B871, 1964.

Timothy H. Hsieh, Hsin Lin, Junwei Liu, Wenhui Duan, Arun Bansil, and Liang Fu. Topological crystalline insulators in the SnTe material class. *Nature Communications*, 3(1):982, 2012.

Lin Huang, Chengxiang Huang, Ziang Wang, Yiyue Du, Chu Wang, Haocheng Lu, Yunyang Li, Xiaoli Liu, Arthur Jiang, and Jia Zhang. E2Former-V2: On-the-fly equivariant attention with linear activation memory. In *International Conference on Machine Learning*, 2026.

Manasa Kaniselvan, Benjamin Kurt Miller, Meng Gao, Juno Nam, and Daniel S. Levine. Learning from the electronic structure of molecules across the periodic table. In *International Conference on Learning Representations*, 2026.

Seongsu Kim, Nayoung Kim, Dongwoo Kim, and Sungsoo Ahn. High-order equivariant flow matching for density functional theory Hamiltonian prediction. In *Advances in Neural Information Processing Systems*, volume 38. Curran Associates, Inc., 2025.

Seongsu Kim, Chanhui Lee, Yoonho Kim, Seongjun Yun, Honghui Kim, Nayoung Kim, Changyoung Park, Sehui Han, Sungbin Lim, and Sungsoo Ahn. Machine learning Hamiltonians are accurate energy-force predictors. In *International Conference on Machine Learning*, 2026.

W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Physical Review*, 140(4A):A1133–A1138, 1965.

Risi Kondor, Zhen Lin, and Shubhendu Trivedi. Clebsch–Gordan nets: A fully Fourier space spherical convolutional neural network. In *Advances in Neural Information Processing Systems*, volume 31, pp. 10138–10147. Curran Associates, Inc., 2018.

Konstantin N. Kudin, Gustavo E. Scuseria, and Eric Cancès. A black-box self-consistent field convergence algorithm: One step closer. *The Journal of Chemical Physics*, 116(19):8255–8261, 2002.

He Li, Zun Wang, Nianlong Zou, Meng Ye, Runzhang Xu, Xiaoxun Gong, Wenhui Duan, and Yong Xu. Deep-learning density functional theory Hamiltonian for efficient ab initio electronic-structure calculation. *Nature Computational Science*, 2(6):367–377, 2022.

Tiancheng Li, Wentao Li, Anyang Peng, Jianming Xue, Linfeng Zhang, Duo Zhang, and Han Wang. DPA4: Pushing the accuracy-cost frontier of interatomic potentials with EMFA SO(2) convolution. *arXiv preprint arXiv:2606.02419*, 2026.

Yunyang Li, Lin Huang, Zhihao Ding, Chu Wang, Xinran Wei, Han Yang, Zun Wang, Chang Liu, Yu Shi, Peiran Jin, Tao Qin, Mark Gerstein, and Jia Zhang. E2Former: An efficient and equivariant transformer with linear-scaling tensor products. In *Advances in Neural Information Processing Systems*, volume 38. Curran Associates, Inc., 2025a.

Yunyang Li, Zaishuo Xia, Lin Huang, Xinran Wei, Han Yang, Sam Harshe, Zun Wang, Chang Liu, Jia Zhang, Bin Shao, and Mark B. Gerstein. Enhancing the scalability and applicability of Kohn-Sham Hamiltonians for molecular systems. In *International Conference on Learning Representations*, 2025b.

Yi-Lun Liao and Tess Smidt. Equiformer: Equivariant graph attention transformer for 3D atomistic graphs. In *International Conference on Learning Representations*, 2023.

Yi-Lun Liao, Brandon Wood, Abhishek Das, and Tess Smidt. EquiformerV2: Improved equivariant transformer for scaling to higher-degree representations. In *International Conference on Learning Representations*, 2024.

Yi-Lun Liao, Alexander J. Hoffman, Sabrina C. Shen, Alexandre Duval, Sam Walton Norwood, and Tess Smidt. EquiformerV3: Scaling efficient, expressive, and general SE(3)-equivariant graph attention transformers. *arXiv preprint arXiv:2604.09130*, 2026.

Yi Liu, Limei Wang, Meng Liu, Yuchao Lin, Xuan Zhang, Bora Oztekin, and Shuiwang Ji. Spherical message passing for 3D molecular graphs. In *International Conference on Learning Representations*, 2022.

Erpai Luo, Xinran Wei, Lin Huang, Yunyang Li, Han Yang, Zaishuo Xia, Zun Wang, Chang Liu, Bin Shao, and Jia Zhang. Efficient and scalable density functional theory Hamiltonian prediction through adaptive sparsity. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 41368–41390. PMLR, 2025.

Eros Mariani and Felix von Oppen. Flexural phonons in free-standing graphene. *Physical Review Letters*, 100(7):076801, 2008.

Saro Passaro and C. Lawrence Zitnick. Reducing SO(3) convolutions to SO(2) for efficient equivariant GNNs. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 27420–27438. PMLR, 2023.

M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos. Iterative minimization techniques for *ab initio* total-energy calculations: Molecular dynamics and conjugate gradients. *Reviews of Modern Physics*, 64(4):1045–1097, 1992.

Chen Qian, Valdas Vitartas, James R. Kermode, and Reinhard J. Maurer. Equivariant electronic Hamiltonian prediction with many-body message passing. *npj Computational Materials*, 12(1): 169, 2026.

K. T. Schütt, H. E. Sauceda, P.-J. Kindermans, A. Tkatchenko, and K.-R. Müller. SchNet–a deep learning architecture for molecules and materials. *The Journal of Chemical Physics*, 148(24): 241722, 2018.

K. T. Schütt, M. Gastegger, A. Tkatchenko, K.-R. Müller, and R. J. Maurer. Unifying machine learning and quantum chemistry with a deep neural network for molecular wavefunctions. *Nature Communications*, 10(1):5024, 2019.

Kristof T. Schütt, Oliver T. Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial properties and molecular spectra. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9377–9388. PMLR, 2021.

Martin Schwade, Shaoming Zhang, Frederik Vonhoff, Frederico P. Delgado, and David A. Egger. Physics-informed Hamiltonian learning for large-scale optoelectronic property prediction. *Nature Communications*, 17(1):2652, 2026.

Li-kun Shi, Dong Zhang, Kai Chang, and Justin C. W. Song. Geometric photon-drag effect and nonlinear shift current in centrosymmetric crystals. *Physical Review Letters*, 126(19):197402, 2021.

Edbert J. Sie, Clara M. Nyby, C. D. Pemmaraju, Su Ji Park, Xiaozhe Shen, Jie Yang, Matthias C. Hoffmann, B. K. Ofori-Okai, Renkai Li, Alexander H. Reid, Stephen Weathersby, Ehren Mannebach, Nathan Finney, Daniel Rhodes, Daniel Chenet, Abhinandan Antony, Luis Balicas, James Hone, Thomas P. Devereaux, Tony F. Heinz, Xijie Wang, and Aaron M. Lindenberg. An ultrafast symmetry switch in a Weyl semimetal. *Nature*, 565(7737):61–66, 2019.

José M. Soler, Emilio Artacho, Julian D. Gale, Alberto García, Javier Junquera, Pablo Ordejón, and Daniel Sánchez-Portal. The SIESTA method for ab initio order-N materials simulation. *Journal of Physics: Condensed Matter*, 14(11):2745–2779, 2002.

Zechen Tang, He Li, Peize Lin, Xiaoxun Gong, Gan Jin, Lixin He, Hong Jiang, Xinguo Ren, Wenhui Duan, and Yong Xu. A deep equivariant neural network approach for efficient hybrid density functional calculations. *Nature Communications*, 15(1):8815, 2024.

Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds. *arXiv preprint arXiv:1802.08219*, 2018.

Oliver T. Unke, Mihail Bogojeski, Michael Gastegger, Mario Geiger, Tess Smidt, and Klaus-Robert Müller. SE(3)-equivariant prediction of molecular wavefunctions and electronic densities. In *Advances in Neural Information Processing Systems*, volume 34, pp. 14434–14447. Curran Associates, Inc., 2021.

Zun Wang, Chang Liu, Nianlong Zou, He Zhang, Xinran Wei, Lin Huang, Lijun Wu, and Bin Shao. Infusing self-consistency into density functional theory Hamiltonian prediction via deep equilibrium models. In *Advances in Neural Information Processing Systems*, volume 37, pp. 89652–89681. Curran Associates, Inc., 2024.

Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco Cohen. 3D steerable CNNs: learning rotationally equivariant features in volumetric data. In *Advances in Neural Information Processing Systems*, volume 31, pp. 10402–10413. Curran Associates, Inc., 2018.

B. M. Wojek, M. H. Berntsen, V. Jonsson, A. Szczerbakow, P. Dziawa, B. J. Kowalski, T. Story, and O. Tjernberg. Direct observation and temperature control of the surface dirac gap in a topological crystalline insulator. *Nature Communications*, 6(1):8463, 2015.

Qiunan Xu, Yang Zhang, Klaus Koepernik, Wujun Shi, Jeroen van den Brink, Claudia Felser, and Yan Sun. Comprehensive scan for nonmagnetic Weyl semimetals with nonlinear optical response. *npj Computational Materials*, 6(1):32, 2020.

Aiqin Yang, Dian Jin, Mingkang Liu, Daye Zheng, Qi Wang, Qiangqiang Gu, and Jian-Hua Jiang. Accelerating discovery of infrared nonlinear optical materials with large shift current via high-throughput screening. *npj Computational Materials*, 12(1):188, 2026.

Shi Yin, Xinyang Pan, Fengyan Wang, and Lixin He. TraceGrad: a framework learning expressive SO(3)-equivariant non-linear representations for electronic-structure Hamiltonian prediction. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 72364–72392. PMLR, 2025.

Shi Yin, Zujian Dai, Xinyang Pan, and Lixin He. Advancing universal deep learning for electronic-structure Hamiltonian prediction of materials. In *International Conference on Learning Representations*, 2026.

Haiyang Yu, Zhao Xu, Xiaofeng Qian, Xiaoning Qian, and Shuiwang Ji. Efficient and equivariant graph networks for predicting quantum Hamiltonian. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 40412–40424. PMLR, 2023.

Haiyang Yu, Yuchao Lin, Xuan Zhang, Xiaofeng Qian, and Shuiwang Ji. Efficient prediction of SO(3)-equivariant Hamiltonian matrices via SO(2) local frames. In *International Conference on Machine Learning*, 2026.

He Zhang, Chang Liu, Zun Wang, Xinran Wei, Siyuan Liu, Nanning Zheng, Bin Shao, and Tie-Yan Liu. Self-consistency training for density-functional-theory Hamiltonian prediction. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 59329–59357. PMLR, 2024.

Yang Zhong, Hongyu Yu, Mao Su, Xingao Gong, and Hongjun Xiang. Transferable equivariant graph neural networks for the Hamiltonians of molecules and solids. *npj Computational Materials*, 9(1):182, 2023.

Yang Zhong, Rui Wang, Xingao Gong, and Hongjun Xiang. A universal spin–orbit-coupled Hamiltonian model for accelerated quantum material discovery. *Nature Machine Intelligence*, 8(3): 403–414, 2026.

Zhanghao Zhouyin, Zixi Gan, Shishir K. Pandey, Linfeng Zhang, and Qiangqiang Gu. Learning local equivariant representations for quantum operators. In *International Conference on Learning Representations*, 2025.

C. Lawrence Zitnick, Abhishek Das, Adeesh Kolluru, Janice Lan, Muhammed Shuaibi, Anuroop Sriram, Zachary Ulissi, and Brandon Wood. Spherical channels for modeling atomic interactions. In *Advances in Neural Information Processing Systems*, volume 35, pp. 8054–8067. Curran Associates, Inc., 2022.

Appendices

CONTENTS

A Background on Equivariant Quantum Operator Learning **16**
A.1 Quantum Operators in an Atom-Centered Basis . . . 16
A.1.1 Kohn–Sham equations and the atomic-orbital basis . . . 16
A.1.2 Operator blocks and their transformation law . . . 16
A.2 Orthogonal Symmetry and Irreducible Representations . . . 17
A.2.1 Irreducible representations of $\mathrm{O}(3)$ . . . 17
A.2.2 Real spherical harmonics . . . 18
A.3 Clebsch–Gordan Decomposition of Operator Blocks . . . 18
A.3.1 Coupled components of an orbital block . . . 18
A.4 Equivariant Message Passing and Harmonic Convolution . . . 19
A.5 Edge-Aligned Local Frames and $\mathrm{SO}(2)$ Convolution . . . 19
A.5.1 The rotation stabilizer . . . 19
A.5.2 Paired magnetic orders and local kernels . . . 20
A.6 The Full $\mathrm{O}(2)$ Stabilizer . . . 20

B Method Details and Theoretical Analysis **22**
B.1 Alignment and Representation Conventions . . . 22
B.2 Proof of Local and Global Equivariance . . . 22
B.3 Restriction of $\mathrm{O}(3)$ Irreps to the Local Stabilizer . . . 23
B.4 Complete $\mathrm{O}(2)$ Kernel Space . . . 23
B.4.1 Positive magnetic orders . . . 23
B.4.2 The zero-order block and bias . . . 23
B.5 Correspondence with Parity-Allowed Clebsch–Gordan Couplings . . . 24
B.6 Detailed Message-Passing Architecture . . . 24
B.6.1 Initialization . . . 24
B.6.2 Layer updates . . . 25
B.6.3 Quantum-operator readout . . . 25
B.6.4 Forward procedure . . . 25
B.7 Equivariance of the Complete Network . . . 26
B.8 Parameter Count and Computational Complexity . . . 26
B.9 Blockwise Evaluation . . . 27

C Why $\mathrm{O}(3)$ Equivariance Matters for Quantum Operators **28**
C.1 Parity Content of the Hamiltonian Labels . . . 28
C.2 Symmetry Constraints on Physical Observables . . . 29

D Experimental Details and Supplementary Evaluation **30**
D.1 Datasets and Training . . . 30
D.2 Graphene Responses . . . 30
D.3 Angular-Operator Benchmark . . . 31

A Background on Equivariant Quantum Operator Learning

This appendix reviews the operator representation in Section 2 and the group actions used in Section 3. We first describe quantum operators in a real atomic-orbital basis. We then introduce irreducible representations, Clebsch–Gordan decomposition, and edge-aligned convolutions. Throughout this appendix, we consider spinless scalar operators.

A.1 Quantum Operators in an Atom-Centered Basis

A.1.1 Kohn–Sham equations and the atomic-orbital basis

In Kohn–Sham density functional theory (DFT), auxiliary orbitals $\psi_a(\boldsymbol{r})$ describe the electron density $n_{\mathrm{e}}(\boldsymbol{r})$ at position $\boldsymbol{r}$. They satisfy (Hohenberg & Kohn, 1964; Kohn & Sham, 1965)

$$H_{\mathrm{KS}}[n_{\mathrm{e}}]\psi_a(\boldsymbol{r}) = \varepsilon_a \psi_a(\boldsymbol{r}), \qquad n_{\mathrm{e}}(\boldsymbol{r}) = \sum_a f_a |\psi_a(\boldsymbol{r})|^2. \tag{15}$$

Here, H_{KS} is the Kohn–Sham Hamiltonian operator, ε_a is the energy of orbital a, and f_a is its occupation. The Hamiltonian depends on the density, so these equations are usually solved by self-consistent field iterations.

Let $\phi_{in\ell m}$ denote a real atomic orbital centered at position $\boldsymbol{r}_i$. As in Section 2, n is the radial index, ℓ is the angular degree, and $m \in \{-\ell, \dots, \ell\}$ is the magnetic quantum number. We use $\alpha = (n, \ell, m)$ as a compact orbital index. A linear combination of atomic orbitals (LCAO) gives

$$\psi_a(\boldsymbol{r}) = \sum_{i,\alpha} C_{i\alpha,a}\phi_{i\alpha}(\boldsymbol{r}), \qquad \phi_{in\ell m}(\boldsymbol{r}) = R_{in\ell}(\|\boldsymbol{r}-\boldsymbol{r}_i\|_2) Y_m^{(\ell)}\left(\frac{\boldsymbol{r}-\boldsymbol{r}_i}{\|\boldsymbol{r}-\boldsymbol{r}_i\|_2}\right). \tag{16}$$

Here, $C_{i\alpha,a}$ is an expansion coefficient, $R_{in\ell}$ is a radial function, and $Y_m^{(\ell)}$ is a real spherical harmonic. Projecting Equation (15) onto this basis gives

$$\boldsymbol{HC} = \boldsymbol{SC}\boldsymbol{\varepsilon}, \qquad H^{ij}_{\alpha\beta} = \langle \phi_{i\alpha}|H_{\mathrm{KS}}|\phi_{j\beta}\rangle, \qquad S^{ij}_{\alpha\beta} = \langle \phi_{i\alpha}|\phi_{j\beta}\rangle. \tag{17}$$

The indices α, β label atomic orbitals. The matrices $\boldsymbol{H}$ and $\boldsymbol{S}$ are the Hamiltonian and overlap matrices, $\boldsymbol{C} = [C_{i\alpha,a}]$, and $\boldsymbol{\varepsilon} = \mathrm{diag}(\varepsilon_a)$. Given $\boldsymbol{H}$ and $\boldsymbol{S}$, this generalized eigenvalue problem yields the orbital energies. Learning $\boldsymbol{H}$ provides a surrogate for the self-consistent calculation (Gong et al., 2023; Yu et al., 2023).

An orbital shell of degree ℓ contains $2\ell+1$ angular components. Its natural inversion parity is $(-1)^\ell$. Table 4 lists common shells and the largest coupled degree in a same-shell block. The last column follows from the Clebsch–Gordan decomposition in Appendix A.3.

Table 4: Angular degrees, dimensions, and natural parities of atomic-orbital shells. A block between two degree-ℓ shells contains coupled components up to degree 2ℓ.

Shell	ℓ	Components	Natural parity	Largest coupled degree
s	0	1	$+1$	0
p	1	3	-1	2
d	2	5	$+1$	4
f	3	7	-1	6
g	4	9	$+1$	8

A.1.2 Operator blocks and their transformation law

For a one-electron operator O, the atom-pair orbital block is

$$\boldsymbol{O}^{ij}_{n_1\ell_1,n_2\ell_2} = [\langle \phi_{in_1\ell_1m_1}|O|\phi_{jn_2\ell_2m_2}\rangle]_{m_1,m_2}. \tag{18}$$

This is the $(2\ell_1+1)\times(2\ell_2+1)$ block defined in Equation (1). Setting O to the identity gives an overlap block. On-site blocks couple orbitals at the same atomic center, while off-site blocks couple different centers. For periodic structures, the atom indices include the selected periodic images. For a Hermitian operator,

$$\boldsymbol{O}^{ji}_{n_2\ell_2,n_1\ell_1} = \left(\boldsymbol{O}^{ij}_{n_1\ell_1,n_2\ell_2}\right)^\dagger, \tag{19}$$

where † denotes the conjugate transpose. It reduces to the transpose for the real blocks considered here.

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an atomistic graph with atom set $\mathcal{V}$ and directed edge set $\mathcal{E}$. Each node i has atomic type z_i and position $\boldsymbol{r}_i$. For an orthogonal transformation $\boldsymbol{Q} \in \mathrm{O}(3)$, let $\boldsymbol{Q}\mathcal{G}$ replace each position by $\boldsymbol{Q}\boldsymbol{r}_i$. The cell vectors transform in the same way for periodic structures. Denote by $\boldsymbol{P}_{\ell,p}(\boldsymbol{Q})$ the representation matrix acting on an orbital shell of degree ℓ and parity p. For a scalar operator that is even under inversion,

$$\boldsymbol{O}^{ij}_{n_1\ell_1,n_2\ell_2}[\boldsymbol{Q}\mathcal{G}] = \boldsymbol{P}_{\ell_1,p_1}(\boldsymbol{Q})\boldsymbol{O}^{ij}_{n_1\ell_1,n_2\ell_2}[\mathcal{G}]\boldsymbol{P}_{\ell_2,p_2}(\boldsymbol{Q})^\top, \qquad p_k = (-1)^{\ell_k}, \quad k \in \{1,2\}. \tag{20}$$

A rotation mixes, for example, the p_x, p_y, and p_z orbitals. Their matrix entries therefore change with the coordinate frame. An equivariant predictor must preserve this transformation law, as required by Equation (5) (Unke et al., 2021; Gong et al., 2023; Yu et al., 2023).

A.2 Orthogonal Symmetry and Irreducible Representations

The orthogonal group is

$$\mathrm{O}(3) = \{\boldsymbol{Q} \in \mathbb{R}^{3\times 3} \mid \boldsymbol{Q}^\top\boldsymbol{Q} = \boldsymbol{I}_3\}, \tag{21}$$

where $\boldsymbol{I}_d$ denotes the $d \times d$ identity matrix. Transformations with determinant $+1$ form the proper rotation group $\mathrm{SO}(3)$. Those with determinant -1 include plane reflections and spatial inversion. Using relative atomic positions removes dependence on a common translation. The remaining coordinate transformations are described by $\mathrm{O}(3)$.

A representation assigns a matrix $\boldsymbol{P}(\boldsymbol{Q})$ to each transformation, with $\boldsymbol{P}(\boldsymbol{Q}_1\boldsymbol{Q}_2) = \boldsymbol{P}(\boldsymbol{Q}_1)\boldsymbol{P}(\boldsymbol{Q}_2)$ and $\boldsymbol{P}(\boldsymbol{I}_3) = \boldsymbol{I}$. Here, $\boldsymbol{I}$ has the dimension of the representation space. We consider finite-dimensional continuous representations and use real orthogonal matrices.

Definition A.1 (Equivariance)**.** Let V_{in} and V_{out} be input and output spaces with representation matrices $\boldsymbol{P}_{\mathrm{in}}$ and $\boldsymbol{P}_{\mathrm{out}}$. A map $f : V_{\mathrm{in}} \to V_{\mathrm{out}}$ is $\mathrm{O}(3)$-equivariant if

$$f(\boldsymbol{P}_{\mathrm{in}}(\boldsymbol{Q})\boldsymbol{x}) = \boldsymbol{P}_{\mathrm{out}}(\boldsymbol{Q})f(\boldsymbol{x}) \quad \text{for every } \boldsymbol{Q} \in \mathrm{O}(3),\ \boldsymbol{x} \in V_{\mathrm{in}}. \tag{22}$$

When the output action is the identity, this condition reduces to invariance. A linear equivariant map is called an intertwiner.

A.2.1 Irreducible representations of $\mathrm{O}(3)$

An irreducible representation (irrep) has no nonzero proper invariant subspace. Following Section 2, we denote an $\mathrm{O}(3)$ irrep and its representation space by $V_{\ell,p}$. It has dimension $2\ell+1$, angular degree $\ell \in \{0, 1, \ldots\}$, and inversion parity $p \in \{\pm 1\}$. Let $\boldsymbol{D}^{(\ell)}(\boldsymbol{R})$ be the real degree-ℓ Wigner matrix for a proper rotation $\boldsymbol{R}$. Then

$$\boldsymbol{P}_{\ell,p}(\boldsymbol{Q}) = p^{(1-\det\boldsymbol{Q})/2}\boldsymbol{D}^{(\ell)}((\det\boldsymbol{Q})\boldsymbol{Q})\,. \tag{23}$$

The argument $(\det\boldsymbol{Q})\boldsymbol{Q}$ is always a proper rotation. In particular, $\boldsymbol{P}_{\ell,p}(-\boldsymbol{I}_3) = p\boldsymbol{I}_{2\ell+1}$. Thus, irreps with the same ℓ and opposite p have identical rotation actions but different inversion actions.

Table 5 gives low-degree examples. A polar vector has type $(1,-1)$, while an axial vector has type $(1,+1)$. Similarly, a rotation-invariant scalar with $\ell = 0$ can be even or odd under inversion. Only type $(0,+1)$ is invariant under all of $\mathrm{O}(3)$.

Table 5: Common geometric quantities and their $\mathrm{O}(3)$ representation types.

Type	ℓ	p	Dimension	Examples
Scalar	0	$+1$	1	energy, charge density
Pseudoscalar	0	-1	1	scalar chirality
Polar vector	1	-1	3	position, displacement, electric field
Axial vector	1	$+1$	3	angular momentum, magnetic field
Symmetric traceless tensor	2	$+1$	5	electric quadrupole moment

A.2.2 Real spherical harmonics

Let $\mathbb{S}^2$ be the unit sphere in $\mathbb{R}^3$. For a unit direction $\boldsymbol{u} \in \mathbb{S}^2$, collect the degree-$\ell$ real spherical harmonics as $\boldsymbol{Y}^{(\ell)}(\boldsymbol{u}) = [Y_m^{(\ell)}(\boldsymbol{u})]_{m=-\ell}^{\ell}$. With integral normalization, they satisfy

$$\int_{\mathbb{S}^2} Y_m^{(\ell)}(\boldsymbol{u}) Y_{m'}^{(\ell')}(\boldsymbol{u})\, \mathrm{d}\Omega(\boldsymbol{u}) = \delta_{\ell\ell'}\delta_{mm'}, \tag{24}$$

where $\mathrm{d}\Omega$ is the surface-area measure and δ is the Kronecker delta. Their transformation law is

$$\boldsymbol{Y}^{(\ell)}(\boldsymbol{Q}\boldsymbol{u}) = \boldsymbol{P}_{\ell,(-1)^\ell}(\boldsymbol{Q})\boldsymbol{Y}^{(\ell)}(\boldsymbol{u}). \tag{25}$$

The parity follows from $Y_m^{(\ell)}(-\boldsymbol{u}) = (-1)^\ell Y_m^{(\ell)}(\boldsymbol{u})$. These fixed functions encode the edge direction in Equation (3) (Thomas et al., 2018; Weiler et al., 2018; Geiger & Smidt, 2022). General feature channels may have either inversion parity at a given angular degree.

A.3 Clebsch–Gordan Decomposition of Operator Blocks

A tensor product combines two representation spaces. For $\boldsymbol{x}_1 \in V_{\ell_1,p_1}$ and $\boldsymbol{x}_2 \in V_{\ell_2,p_2}$, the product $\boldsymbol{x}_1 \otimes \boldsymbol{x}_2$ transforms by $\boldsymbol{P}_{\ell_1,p_1}(\boldsymbol{Q}) \otimes \boldsymbol{P}_{\ell_2,p_2}(\boldsymbol{Q})$. The Clebsch–Gordan (CG) change of basis separates this product into irreducible components.

Theorem A.2 (Clebsch–Gordan decomposition)**.** *For real* $\mathrm{O}(3)$ *irreps,*

$$V_{\ell_1,p_1} \otimes V_{\ell_2,p_2} \cong \bigoplus_{\ell_3=|\ell_1-\ell_2|}^{\ell_1+\ell_2} V_{\ell_3,p_1p_2}. \tag{26}$$

Each allowed degree ℓ_3 *appears once, with output parity* $p_3 = p_1p_2$.

For $m_3 \in \{-\ell_3, \dots, \ell_3\}$, the corresponding component is

$$[\boldsymbol{x}_1 \otimes \boldsymbol{x}_2]_{\ell_3,p_3,m_3} = \sum_{m_1=-\ell_1}^{\ell_1} \sum_{m_2=-\ell_2}^{\ell_2} c_{\ell_1,m_1,\ell_2,m_2}^{\ell_3,m_3} x_{1,m_1} x_{2,m_2}, \tag{27}$$

where x_{1,m_1} and x_{2,m_2} are entries of the two input vectors. The coefficients $c_{\ell_1,m_1,\ell_2,m_2}^{\ell_3,m_3}$ are fixed, as in Equation (2). They can be chosen to give a real orthogonal change of basis. The network learns channel weights and coupled components, while the CG coefficients remain fixed (Kondor et al., 2018; Geiger & Smidt, 2022).

A.3.1 Coupled components of an orbital block

Let $\mathrm{Hom}(V, W)$ denote the space of linear maps from V to W, and let V^* be the dual representation space. A shell-pair block maps the column space to the row space, so

$$\mathrm{Hom}(V_{\ell_2,p_2}, V_{\ell_1,p_1}) \cong V_{\ell_1,p_1} \otimes V_{\ell_2,p_2}^*. \tag{28}$$

The representation matrices are real and orthogonal, which identifies the dual with the original representation. For a fixed shell pair $(n_1\ell_1, n_2\ell_2)$, the coupled coefficients are

$$o_{\ell_3,p_3,m_3}^{ij} = \sum_{m_1=-\ell_1}^{\ell_1} \sum_{m_2=-\ell_2}^{\ell_2} c_{\ell_1,m_1,\ell_2,m_2}^{\ell_3,m_3} O_{n_1\ell_1m_1,n_2\ell_2m_2}^{ij}. \tag{29}$$

The shell-pair indices on o are suppressed, as in the main text. Different shell pairs have separate coupled components. Since orbital shells have $p_k = (-1)^{\ell_k}$, the block decomposes as

$$\boldsymbol{O}_{n_1\ell_1,n_2\ell_2}^{ij} \longleftrightarrow \bigoplus_{\ell_3=|\ell_1-\ell_2|}^{\ell_1+\ell_2} \boldsymbol{o}_{\ell_3,(-1)^{\ell_1+\ell_2}}^{ij}. \tag{30}$$

All coupled degrees in a fixed shell-pair block have the same parity. For example,

$$\begin{aligned} s\text{–}p: &\quad V_{0,+1} \otimes V_{1,-1} \cong V_{1,-1}, \\ p\text{–}p: &\quad V_{1,-1} \otimes V_{1,-1} \cong V_{0,+1} \oplus V_{1,+1} \oplus V_{2,+1}, \\ p\text{–}d: &\quad V_{1,-1} \otimes V_{2,+1} \cong V_{1,-1} \oplus V_{2,-1} \oplus V_{3,-1}. \end{aligned} \tag{31}$$

The network predicts $\hat{\boldsymbol{o}}^{ij}_{\ell_3,p_3}$ with these representation types. The inverse CG transformation then gives $\widehat{\boldsymbol{O}}^{ij}_{n_1\ell_1,n_2\ell_2}$ through Equation (2). This is the Wigner–Eckart output layer used in the main text (Gong et al., 2023; Yu et al., 2023).

A.4 Equivariant Message Passing and Harmonic Convolution

For a directed edge $i \to j$, let $\boldsymbol{r}_{ij} = \boldsymbol{r}_j - \boldsymbol{r}_i$, $r_{ij} = \|\boldsymbol{r}_{ij}\|_2$, and $\hat{\boldsymbol{r}}_{ij} = \boldsymbol{r}_{ij}/r_{ij}$. The neighbor set $\mathcal{N}(i)$ contains atoms within the cutoff radius. As in Section 3, the input $\boldsymbol{x}_{ij}$ may collect node, edge, or hidden features. With explicit $\mathrm{O}(3)$ type labels, its channels are written as

$$\boldsymbol{x}^{(t)}_{ij} = \bigoplus_{\ell=0}^{\ell_{\max}} \bigoplus_{p\in\{\pm1\}} \bigoplus_{c=1}^{C^{\mathrm{in}}_{\ell,p}} \boldsymbol{x}^{(t)}_{ij,\ell,p,c}, \qquad \boldsymbol{x}^{(t)}_{ij,\ell,p,c} \in V_{\ell,p}, \tag{32}$$

where t is the layer index and $\ell_{\max}$ is the maximum angular degree. As in Section 3.1, $C^{\mathrm{in}}_{\ell,p}$ counts input channels of type (ℓ, p), and c indexes these channels.

Suppressing layer, channel, and parity indices gives the angular message in Equation (3):

$$\boldsymbol{m}_{i,\ell_{\mathrm{out}}} = \sum_{j\in\mathcal{N}(i)} \sum_{\ell_{\mathrm{in}},\ell_f} w^{\ell_{\mathrm{out}}}_{\ell_{\mathrm{in}},\ell_f}(\boldsymbol{z}_{ij})[\boldsymbol{x}_{ij,\ell_{\mathrm{in}}} \otimes \boldsymbol{Y}^{(\ell_f)}(\hat{\boldsymbol{r}}_{ij})]_{\ell_{\mathrm{out}}}. \tag{33}$$

Here, ℓ_{in}, ℓ_f, and ℓ_{out} are the input, filter, and output degrees. The vector $\boldsymbol{m}_{i,\ell_{\mathrm{out}}}$ is the aggregated message. The learnable function w takes a rotation-invariant pair embedding $\boldsymbol{z}_{ij}$. This embedding contains atom types and distances and may include learned scalar features. The spherical harmonics are fixed.

The CG projection requires $|\ell_{\mathrm{in}} - \ell_f| \le \ell_{\mathrm{out}} \le \ell_{\mathrm{in}} + \ell_f$. For full $\mathrm{O}(3)$ equivariance, let p_{in} and p_{out} denote the input and output inversion parities. If the scalar weights are $\mathrm{O}(3)$-invariant, the allowed couplings satisfy

$$|\ell_{\mathrm{in}} - \ell_f| \le \ell_{\mathrm{out}} \le \ell_{\mathrm{in}} + \ell_f, \qquad p_{\mathrm{out}} = p_{\mathrm{in}}(-1)^{\ell_f}. \tag{34}$$

The parity condition is an additional requirement beyond rotation equivariance. In particular, the degree-zero features used as invariant weights must have even parity.

Linear channel maps, gated nonlinearities, and neighbor sums preserve equivariance when they respect these representation types (Thomas et al., 2018; Unke et al., 2021; Batzner et al., 2022; Liao & Smidt, 2023). Such operations update the node, edge, hidden, and pair features in Equation (4). In standard neighbor-to-neighbor message passing, the spatial dependence can expand with network depth. The strictly localized updates used here are described in Appendix B.6.

Remark (Angular complexity)**.** With bounded channel multiplicities and degree cutoffs proportional to $\ell_{\max}$, a conventional full CG contraction has angular cost $\mathcal{O}(\ell^6_{\max})$. This cost concerns the angular operation on an edge, separately from the size of the spatial neighborhood. Edge-aligned convolutions reduce it by evaluating the operation in a local frame (Passaro & Zitnick, 2023; Liao et al., 2024).

A.5 Edge-Aligned Local Frames and SO(2) Convolution

Let $\boldsymbol{u} = \hat{\boldsymbol{r}}_{ij} \in \mathbb{S}^2$ be a unit edge direction and $\boldsymbol{e}_z = (0, 0, 1)^\top$ the fixed reference direction. We choose a rotation $\boldsymbol{F}_{\boldsymbol{u}} \in \mathrm{SO}(3)$ satisfying

$$\boldsymbol{F}_{\boldsymbol{u}}\boldsymbol{e}_z = \boldsymbol{u}. \tag{35}$$

Applying $\boldsymbol{P}_{\mathrm{in}}(\boldsymbol{F}_{\boldsymbol{u}})^{-1}$ expresses the input in the edge-aligned frame. The local kernel acts at $\boldsymbol{e}_z$, and $\boldsymbol{P}_{\mathrm{out}}(\boldsymbol{F}_{\boldsymbol{u}})$ transforms the output back. This is the align–convolve–rotate procedure used by edge-aligned models (Zitnick et al., 2022; Passaro & Zitnick, 2023; Liao et al., 2024).

A.5.1 The rotation stabilizer

The alignment rotation is not unique. For any rotation $\boldsymbol{R}$ about $\boldsymbol{e}_z$, the matrix $\boldsymbol{F}_{\boldsymbol{u}}\boldsymbol{R}$ also maps $\boldsymbol{e}_z$ to $\boldsymbol{u}$. These rotations form the stabilizer

$$\mathcal{S}^+_{\boldsymbol{e}_z} = \{\boldsymbol{R} \in \mathrm{SO}(3) \mid \boldsymbol{R}\boldsymbol{e}_z = \boldsymbol{e}_z\} \cong \mathrm{SO}(2). \tag{36}$$

Theorem A.3 (Reference-direction form of a kernel). *Let $\boldsymbol{K}_{\boldsymbol{u}} : V_{\text{in}} \to V_{\text{out}}$ satisfy*

$$\boldsymbol{K}_{\boldsymbol{R}\boldsymbol{u}}\boldsymbol{P}_{\text{in}}(\boldsymbol{R}) = \boldsymbol{P}_{\text{out}}(\boldsymbol{R})\boldsymbol{K}_{\boldsymbol{u}} \quad \textit{for every } \boldsymbol{R} \in \mathrm{SO}(3),\ \boldsymbol{u} \in \mathbb{S}^2. \tag{37}$$

Then its value $\boldsymbol{K}_{\boldsymbol{e}_z}$ at the reference direction determines every other value:

$$\boldsymbol{K}_{\boldsymbol{u}} = \boldsymbol{P}_{\text{out}}(\boldsymbol{F}_{\boldsymbol{u}})\boldsymbol{K}_{\boldsymbol{e}_z}\boldsymbol{P}_{\text{in}}(\boldsymbol{F}_{\boldsymbol{u}})^{-1}. \tag{38}$$

This construction is independent of the alignment choice if and only if

$$\boldsymbol{P}_{\text{out}}(\boldsymbol{g})\boldsymbol{K}_{\boldsymbol{e}_z} = \boldsymbol{K}_{\boldsymbol{e}_z}\boldsymbol{P}_{\text{in}}(\boldsymbol{g}) \quad \textit{for every } \boldsymbol{g} \in \mathcal{S}^+_{\boldsymbol{e}_z}. \tag{39}$$

Indeed, setting $\boldsymbol{u} = \boldsymbol{e}_z$ and $\boldsymbol{R} = \boldsymbol{F}_{\boldsymbol{u}'}$ in Equation (37) gives the kernel at any direction $\boldsymbol{u}'$. Two alignment choices differ by a rotation in $\mathcal{S}^+_{\boldsymbol{e}_z}$. Substituting them into Equation (38) gives Equation (39). This standard construction explains the local $\mathrm{SO}(2)$ constraint (Weiler et al., 2018; Passaro & Zitnick, 2023). Theorem 1 extends it to the full $\mathrm{O}(3)$ action.

A.5.2 Paired magnetic orders and local kernels

Within $V_{\ell,p}$, let W_s be the real subspace associated with magnetic orders $m = \pm s$, where $0 \le s \le \ell$. The space W_0 has dimension one, while each W_s with $s > 0$ has dimension two. For a rotation about $\boldsymbol{e}_z$ through angle φ, the action on a positive-order subspace is

$$\boldsymbol{R}_s(\varphi) = \begin{pmatrix} \cos(s\varphi) & -\sin(s\varphi) \\ \sin(s\varphi) & \cos(s\varphi) \end{pmatrix}, \qquad s > 0, \tag{40}$$

in a real cosine–sine basis. The rotation acts as the identity on W_0. Thus,

$$V_{\ell,p}\downarrow_{\mathcal{S}^+_{\boldsymbol{e}_z}} \cong W_0 \oplus W_1 \oplus \cdots \oplus W_\ell, \tag{41}$$

where $\downarrow$ denotes restriction to a subgroup. This rotation action does not depend on the inversion parity p.

Channels from different degrees that contain the same pair $\pm s$ therefore have the same local rotation action. Let $C_{\text{in}}(s)$ and $C_{\text{out}}(s)$ be the numbers of these input and output channels. If $C^a_{\ell,p}$ is the multiplicity of type (ℓ, p) in V_a, then $C_a(s) = \sum_{\ell=s}^{\ell_{\max}} \sum_{p\in\{\pm1\}} C^a_{\ell,p}$ for $a \in \{\text{in}, \text{out}\}$.

Theorem A.4 (Real $\mathrm{SO}(2)$ intertwiners). *A linear map that is equivariant under rotations about $\boldsymbol{e}_z$ does not mix different paired orders s. For $s > 0$, its general block is*

$$\boldsymbol{K}_{\boldsymbol{e}_z}(s) = \boldsymbol{A}_s \otimes \boldsymbol{I}_2 + \boldsymbol{B}_s \otimes \boldsymbol{J}_2, \qquad \boldsymbol{J}_2 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \tag{42}$$

where $\boldsymbol{A}_s, \boldsymbol{B}_s \in \mathbb{R}^{C_{\text{out}}(s)\times C_{\text{in}}(s)}$ are learnable channel-mixing matrices. Here, $\otimes$ denotes the Kronecker product. At $s = 0$, the block is an arbitrary real channel-mixing matrix.

Both $\boldsymbol{I}_2$ and $\boldsymbol{J}_2$ commute with every $\boldsymbol{R}_s(\varphi)$, and they span all real matrices with this property. If the channel count per degree is bounded, $C_{\text{in}}(s)$ and $C_{\text{out}}(s)$ grow at most as $\ell_{\max} - s + 1$. The angular cost is therefore

$$\sum_{s=0}^{\ell_{\max}} C_{\text{in}}(s)C_{\text{out}}(s) = \mathcal{O}\left(\sum_{s=0}^{\ell_{\max}} (\ell_{\max} - s + 1)^2\right) = \mathcal{O}(\ell^3_{\max}). \tag{43}$$

This grouping gives the cubic angular cost of edge-aligned $\mathrm{SO}(2)$ convolutions (Passaro & Zitnick, 2023).

A.6 The Full $\mathrm{O}(2)$ Stabilizer

For the full group $\mathrm{O}(3)$, the stabilizer also contains reflections:

$$\mathcal{S}_{\boldsymbol{e}_z} = \{\boldsymbol{Q} \in \mathrm{O}(3) \mid \boldsymbol{Q}\boldsymbol{e}_z = \boldsymbol{e}_z\} \cong \mathrm{O}(2). \tag{44}$$

Its rotation subgroup is $\mathcal{S}^+_{\boldsymbol{e}_z}$. As in Section 3.1, choose $\boldsymbol{Q}_{\boldsymbol{e}_z} = \mathrm{diag}(1, -1, 1)$, the reflection through the xz-plane. For every rotation $\boldsymbol{R}$ about $\boldsymbol{e}_z$,

$$\boldsymbol{Q}_{\boldsymbol{e}_z}\boldsymbol{R}\boldsymbol{Q}^{-1}_{\boldsymbol{e}_z} = \boldsymbol{R}^{-1}. \tag{45}$$

Every element of $\mathcal{S}_{\boldsymbol{e}_z}$ is either $\boldsymbol{R}$ or $\boldsymbol{R}\boldsymbol{Q}_{\boldsymbol{e}_z}$. This gives the semidirect product $\mathcal{S}_{\boldsymbol{e}_z} = \mathcal{S}^+_{\boldsymbol{e}_z} \rtimes \{\boldsymbol{I}_3, \boldsymbol{Q}_{\boldsymbol{e}_z}\}$. The symbol $\rtimes$ records that conjugating a rotation by the reflection reverses its angle.

At $s = 0$, there are two one-dimensional irreps, distinguished by the reflection sign. For each $s > 0$, there is one two-dimensional real irrep up to a change of basis. Its rotation action is $\boldsymbol{R}_s(\varphi)$, and a canonical reflection action is

$$\boldsymbol{\Sigma}_s = \begin{cases} 1, & s = 0, \\ \mathrm{diag}(1, -1), & s > 0. \end{cases} \tag{46}$$

For $s > 0$, these matrices satisfy $\boldsymbol{\Sigma}_s \boldsymbol{R}_s(\varphi) \boldsymbol{\Sigma}_s^{-1} = \boldsymbol{R}_s(-\varphi)$. Table 6 lists the irreps.

Table 6: Real irreps of $\mathrm{O}(2)$. The reflection column refers to the generating reflection $\boldsymbol{Q}_{\boldsymbol{e}_z}$.

Type	Dimension	Rotation action	Reflection action
$s = 0$, sign $+1$	1	1	$+1$
$s = 0$, sign -1	1	1	-1
$s > 0$	2	$\boldsymbol{R}_s(\varphi)$	$\boldsymbol{\Sigma}_s$

For $s > 0$, the change of basis $\boldsymbol{J}_2$ sends $\boldsymbol{\Sigma}_s$ to $-\boldsymbol{\Sigma}_s$ and leaves the rotation action unchanged. These two reflection matrices therefore describe equivalent $\mathrm{O}(2)$ irreps. However, their relative sign matters when the channels inherit fixed bases from $\mathrm{O}(3)$ irreps. For a channel of type (ℓ, p), Equation (11) gives the reflection action $\tau_{\ell,p}\boldsymbol{\Sigma}_s$, with $\tau_{\ell,p} = p(-1)^\ell$.

An $\mathrm{O}(2)$-equivariant local kernel must satisfy both the rotation and reflection constraints. The relation $\boldsymbol{\Sigma}_s \boldsymbol{J}_2 = -\boldsymbol{J}_2 \boldsymbol{\Sigma}_s$ allows the $\boldsymbol{I}_2$ term for equal reflection signs and the $\boldsymbol{J}_2$ term for opposite signs. These are the channel constraints in Proposition 1. Their derivation is given in Appendix B.4.

Remark (Reflection and spatial inversion)**.** The reflection $\boldsymbol{Q}_{\boldsymbol{e}_z}$ fixes $\boldsymbol{e}_z$ and belongs to $\mathcal{S}_{\boldsymbol{e}_z}$. Spatial inversion $-\boldsymbol{I}_3$ sends $\boldsymbol{e}_z$ to $-\boldsymbol{e}_z$ and does not belong to this stabilizer. Both have determinant -1. After edge alignment, a global improper transformation induces a local reflection through the residual transformation in Equation (8). This is why the local $\mathrm{O}(2)$ constraint covers global reflections and spatial inversion.

B METHOD DETAILS AND THEORETICAL ANALYSIS

This appendix gives the proofs and implementation details for Section 3. We first prove the local kernel constraints, then describe the network updates, parameter counts, and blockwise evaluation.

B.1 ALIGNMENT AND REPRESENTATION CONVENTIONS

We use the input and output spaces V_{in} and V_{out} from Section 3.1. They are finite direct sums of real $\mathrm{O}(3)$ irreps, with representation matrices $\boldsymbol{P}_{\text{in}}$ and $\boldsymbol{P}_{\text{out}}$. Let $\boldsymbol{e}_z = (0,0,1)^\top$ be the reference direction and $\boldsymbol{u} \in \mathbb{S}^2$ a unit edge direction, where $\mathbb{S}^2$ is the unit sphere in $\mathbb{R}^3$. We choose $\boldsymbol{F}_{\boldsymbol{u}} \in \mathrm{SO}(3)$ such that $\boldsymbol{F}_{\boldsymbol{u}}\boldsymbol{e}_z = \boldsymbol{u}$, with $\boldsymbol{F}_{\boldsymbol{e}_z} = \boldsymbol{I}_3$. Here, $\boldsymbol{I}_d$ is the $d \times d$ identity matrix. Any dependence of the kernel on invariant inputs is held fixed in the proofs below.

The stabilizer $\mathcal{S}_{\boldsymbol{e}_z}$ consists of all orthogonal transformations that fix $\boldsymbol{e}_z$. Its rotation subgroup is $\mathcal{S}^+_{\boldsymbol{e}_z} = \mathcal{S}_{\boldsymbol{e}_z} \cap \mathrm{SO}(3) \cong \mathrm{SO}(2)$. Two alignment choices for the same direction satisfy $\boldsymbol{F}'_{\boldsymbol{u}} = \boldsymbol{F}_{\boldsymbol{u}}\boldsymbol{R}$ for some $\boldsymbol{R} \in \mathcal{S}^+_{\boldsymbol{e}_z}$. If the reference kernel $\boldsymbol{K}_{\boldsymbol{e}_z}$ intertwines this subgroup, then

$$\begin{aligned}\boldsymbol{P}_{\text{out}}(\boldsymbol{F}'_{\boldsymbol{u}})\boldsymbol{K}_{\boldsymbol{e}_z}\boldsymbol{P}_{\text{in}}(\boldsymbol{F}'_{\boldsymbol{u}})^{-1} &= \boldsymbol{P}_{\text{out}}(\boldsymbol{F}_{\boldsymbol{u}})\boldsymbol{P}_{\text{out}}(\boldsymbol{R})\boldsymbol{K}_{\boldsymbol{e}_z}\boldsymbol{P}_{\text{in}}(\boldsymbol{R})^{-1}\boldsymbol{P}_{\text{in}}(\boldsymbol{F}_{\boldsymbol{u}})^{-1} \\ &= \boldsymbol{P}_{\text{out}}(\boldsymbol{F}_{\boldsymbol{u}})\boldsymbol{K}_{\boldsymbol{e}_z}\boldsymbol{P}_{\text{in}}(\boldsymbol{F}_{\boldsymbol{u}})^{-1}. \end{aligned} \tag{47}$$

Thus, Equation (7) does not depend on the alignment choice. A continuous choice of $\boldsymbol{F}_{\boldsymbol{u}}$ over the whole sphere is not required.

As in Equation (11), we choose the generating reflection

$$\boldsymbol{Q}_{\boldsymbol{e}_z} = \mathrm{diag}(1,-1,1). \tag{48}$$

It fixes $\boldsymbol{e}_z$ and reflects through the xz-plane. Together with $\mathcal{S}^+_{\boldsymbol{e}_z}$, it generates the full stabilizer $\mathcal{S}_{\boldsymbol{e}_z} \cong \mathrm{O}(2)$.

B.2 PROOF OF LOCAL AND GLOBAL EQUIVARIANCE

We use the following property of the residual transformation.

Lemma B.1 (Residual stabilizer element)**.** *For $\boldsymbol{Q} \in \mathrm{O}(3)$ and $\boldsymbol{u} \in \mathbb{S}^2$, let*

$$\boldsymbol{g}(\boldsymbol{Q},\boldsymbol{u}) = \boldsymbol{F}^{-1}_{\boldsymbol{Qu}}\boldsymbol{Q}\boldsymbol{F}_{\boldsymbol{u}}. \tag{49}$$

Then $\boldsymbol{g}(\boldsymbol{Q},\boldsymbol{u}) \in \mathcal{S}_{\boldsymbol{e}_z}$ and $\det \boldsymbol{g}(\boldsymbol{Q},\boldsymbol{u}) = \det \boldsymbol{Q}$. Every element of $\mathcal{S}_{\boldsymbol{e}_z}$ occurs as a residual transformation.

Proof. The alignment condition gives

$$\boldsymbol{g}(\boldsymbol{Q},\boldsymbol{u})\boldsymbol{e}_z = \boldsymbol{F}^{-1}_{\boldsymbol{Qu}}\boldsymbol{Q}\boldsymbol{F}_{\boldsymbol{u}}\boldsymbol{e}_z = \boldsymbol{F}^{-1}_{\boldsymbol{Qu}}\boldsymbol{Q}\boldsymbol{u} = \boldsymbol{e}_z. \tag{50}$$

Both alignment matrices have determinant $+1$, so the determinant of the residual is $\det \boldsymbol{Q}$. For any $\boldsymbol{g} \in \mathcal{S}_{\boldsymbol{e}_z}$, setting $\boldsymbol{Q} = \boldsymbol{g}$ and $\boldsymbol{u} = \boldsymbol{e}_z$ gives $\boldsymbol{g}(\boldsymbol{g},\boldsymbol{e}_z) = \boldsymbol{g}$, since $\boldsymbol{F}_{\boldsymbol{e}_z} = \boldsymbol{I}_3$. □

Proof of Theorem 1. First assume that $\boldsymbol{K}_{\boldsymbol{e}_z}$ satisfies the local constraint in Equation (10). For any $\boldsymbol{Q}$ and $\boldsymbol{u}$, write $\boldsymbol{g} = \boldsymbol{g}(\boldsymbol{Q},\boldsymbol{u})$. By Equation (49),

$$\boldsymbol{Q}\boldsymbol{F}_{\boldsymbol{u}} = \boldsymbol{F}_{\boldsymbol{Qu}}\boldsymbol{g}. \tag{51}$$

Substituting this relation into Equation (7) gives

$$\begin{aligned}\boldsymbol{K}_{\boldsymbol{Qu}}\boldsymbol{P}_{\text{in}}(\boldsymbol{Q}) &= \boldsymbol{P}_{\text{out}}(\boldsymbol{F}_{\boldsymbol{Qu}})\boldsymbol{K}_{\boldsymbol{e}_z}\boldsymbol{P}_{\text{in}}(\boldsymbol{F}^{-1}_{\boldsymbol{Qu}}\boldsymbol{Q}) \\ &= \boldsymbol{P}_{\text{out}}(\boldsymbol{F}_{\boldsymbol{Qu}})\boldsymbol{K}_{\boldsymbol{e}_z}\boldsymbol{P}_{\text{in}}(\boldsymbol{g})\boldsymbol{P}_{\text{in}}(\boldsymbol{F}_{\boldsymbol{u}})^{-1} \\ &= \boldsymbol{P}_{\text{out}}(\boldsymbol{F}_{\boldsymbol{Qu}})\boldsymbol{P}_{\text{out}}(\boldsymbol{g})\boldsymbol{K}_{\boldsymbol{e}_z}\boldsymbol{P}_{\text{in}}(\boldsymbol{F}_{\boldsymbol{u}})^{-1} \\ &= \boldsymbol{P}_{\text{out}}(\boldsymbol{Q})\boldsymbol{K}_{\boldsymbol{u}}. \end{aligned} \tag{52}$$

This proves sufficiency.

Conversely, assume the global constraint in Equation (9). Set $\boldsymbol{u} = \boldsymbol{e}_z$ and choose any $\boldsymbol{Q} = \boldsymbol{g} \in \mathcal{S}_{\boldsymbol{e}_z}$. Since $\boldsymbol{g}\boldsymbol{e}_z = \boldsymbol{e}_z$ and $\boldsymbol{F}_{\boldsymbol{e}_z} = \boldsymbol{I}_3$, we obtain

$$\boldsymbol{K}_{\boldsymbol{e}_z}\boldsymbol{P}_{\text{in}}(\boldsymbol{g}) = \boldsymbol{P}_{\text{out}}(\boldsymbol{g})\boldsymbol{K}_{\boldsymbol{e}_z}. \tag{53}$$

This is the local $\mathrm{O}(2)$ constraint. It also implies the rotation constraint, so Equation (47) proves independence from the alignment choice. □

B.3 Restriction of O(3) Irreps to the Local Stabilizer

For an irrep $V_{\ell,p}$, let W_s be the subspace associated with magnetic orders $m = \pm s$, where $0 \le s \le \ell$. The zero-order space W_0 is one-dimensional, while W_s is two-dimensional for $s > 0$. The rotation action on each W_s is independent of p. The reflection action retains this parity information.

Lemma B.2 (Local reflection action)**.** *In a real cosine–sine basis,*

$$\boldsymbol{P}_{\ell,p}(\boldsymbol{Q}_{\boldsymbol{e}_z})\big|_{W_s} = \tau_{\ell,p}\boldsymbol{\Sigma}_s, \qquad \tau_{\ell,p} = p(-1)^\ell, \qquad \boldsymbol{\Sigma}_s = \begin{cases} 1, & s = 0, \\ \mathrm{diag}(1,-1), & s > 0. \end{cases} \tag{54}$$

Here, $\tau_{\ell,p}$ is the local reflection sign.

Proof. Real spherical harmonics of degree ℓ have inversion parity $(-1)^\ell$. Reflection through the xz-plane preserves the cosine and zero-order components and reverses the sine component. Thus, $\boldsymbol{P}_{\ell,(-1)^\ell}(\boldsymbol{Q}_{\boldsymbol{e}_z})|_{W_s} = \boldsymbol{\Sigma}_s$. Changing the inversion parity from $(-1)^\ell$ to p multiplies every determinant--1 action by $p(-1)^\ell$. This gives Equation (54). □

For $s = 0$, the sign $\tau_{\ell,p}$ distinguishes the two one-dimensional reflection actions. For $s > 0$, $\boldsymbol{\Sigma}_s$ and $-\boldsymbol{\Sigma}_s$ describe equivalent O(2) irreps, as discussed in Appendix A.6. Their relative sign still determines the allowed couplings between channels in the fixed O(3) bases. A consistent change of basis conjugates all these matrices and leaves the constraints unchanged.

B.4 Complete O(2) Kernel Space

We now prove Proposition 1. Different paired orders have different rotation actions, so a linear intertwiner cannot mix them. We treat the positive and zero orders separately.

B.4.1 Positive magnetic orders

For $s > 0$, let $C_{\mathrm{in}}(s)$ and $C_{\mathrm{out}}(s)$ count the input and output channels containing orders $\pm s$. Fix an input channel a and an output channel b. By Theorem A.4, their 2×2 block has the form

$$[\boldsymbol{K}_{\boldsymbol{e}_z}(s)]_{ba} = A_{s,ba}\boldsymbol{I}_2 + B_{s,ba}\boldsymbol{J}_2, \qquad \boldsymbol{J}_2 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \tag{55}$$

where $A_{s,ba}$ and $B_{s,ba}$ are real coefficients. Let $\tau^{\mathrm{in}}_{s,a}$ and $\tau^{\mathrm{out}}_{s,b}$ be the reflection signs of these channels. Reflection equivariance requires

$$\tau^{\mathrm{out}}_{s,b}\boldsymbol{\Sigma}_s[\boldsymbol{K}_{\boldsymbol{e}_z}(s)]_{ba} = [\boldsymbol{K}_{\boldsymbol{e}_z}(s)]_{ba}\tau^{\mathrm{in}}_{s,a}\boldsymbol{\Sigma}_s. \tag{56}$$

Using $\boldsymbol{\Sigma}_s\boldsymbol{J}_2 = -\boldsymbol{J}_2\boldsymbol{\Sigma}_s$, we obtain

$$(\tau^{\mathrm{out}}_{s,b} - \tau^{\mathrm{in}}_{s,a})A_{s,ba} = 0, \qquad (\tau^{\mathrm{out}}_{s,b} + \tau^{\mathrm{in}}_{s,a})B_{s,ba} = 0. \tag{57}$$

Thus, $A_{s,ba}$ is allowed for equal signs, while $B_{s,ba}$ is allowed for opposite signs. Collecting the channel pairs gives the masks in Equation (12) and the kernel in Equation (13). Across the full feature spaces, s ranges up to $\ell_{\max}$. An individual degree-ℓ channel contributes only when $s \le \ell$.

B.4.2 The zero-order block and bias

At $s = 0$, both rotation actions are the identity. Writing $K_{0,ba} = [\boldsymbol{K}_{\boldsymbol{e}_z}(0)]_{ba}$, the reflection constraint becomes

$$\tau^{\mathrm{out}}_{0,b}K_{0,ba} = K_{0,ba}\tau^{\mathrm{in}}_{0,a}. \tag{58}$$

Hence a weight is allowed only between equal reflection signs. For an affine map, let b_b be the constant bias of output channel b. The bias must satisfy

$$\tau^{\mathrm{out}}_{0,b}b_b = b_b, \tag{59}$$

so it is allowed only in the $+1$ sector. Positive-order channels admit no constant bias because no nonzero vector is fixed by every local rotation.

Rotations and $\boldsymbol{Q}_{\boldsymbol{e}_z}$ generate $\mathcal{S}_{\boldsymbol{e}_z}$. The blocks above therefore satisfy the constraint for every stabilizer element. Conversely, each rotation and reflection condition is necessary, so no other linear blocks are allowed. This proves Proposition 1.

B.5 Correspondence with Parity-Allowed Clebsch–Gordan Couplings

The reflection masks reduce the number of coefficients in an $\mathrm{SO}(2)$ kernel. We show that they retain every angular coupling allowed by $\mathrm{O}(3)$. The invariant inputs are held fixed in this comparison.

Proposition B.3 (Correspondence with Clebsch–Gordan couplings)**.** *Consider an input irrep $V_{\ell_{\mathrm{in}},p_{\mathrm{in}}}$ and an output irrep $V_{\ell_{\mathrm{out}},p_{\mathrm{out}}}$. Let $q = \min(\ell_{\mathrm{in}}, \ell_{\mathrm{out}})$. The allowed spherical-harmonic filter degrees form the set*

$$\mathcal{L}_{\mathrm{CG}} = \left\{ \ell_f : |\ell_{\mathrm{out}} - \ell_{\mathrm{in}}| \le \ell_f \le \ell_{\mathrm{out}} + \ell_{\mathrm{in}}, \quad (-1)^{\ell_f} = p_{\mathrm{out}} p_{\mathrm{in}} \right\}. \tag{60}$$

If the two local reflection signs are equal, both $|\mathcal{L}_{\mathrm{CG}}|$ and the number of EACO coefficients equal $q+1$. Otherwise, both equal q. Evaluation at $\boldsymbol{e}_z$ gives an invertible change of basis between the allowed CG kernels and the EACO kernels.

Proof. The triangle interval contains $2q+1$ integers. Its endpoints have parity $(-1)^{\ell_{\mathrm{out}}+\ell_{\mathrm{in}}}$. Equality of the local reflection signs is equivalent to

$$p_{\mathrm{out}} p_{\mathrm{in}} = (-1)^{\ell_{\mathrm{out}}+\ell_{\mathrm{in}}}. \tag{61}$$

In this case, the parity constraint selects $q+1$ filter degrees. EACO also retains one coefficient for each $s = 0, \dots, q$. For opposite signs, the constraint selects q filter degrees, and EACO retains one coefficient for each $s = 1, \dots, q$.

For each allowed filter degree, the CG contraction with $\boldsymbol{Y}^{(\ell_f)}$ gives a globally equivariant kernel. These kernels are linearly independent because spherical harmonics of different degrees are orthogonal. By Theorem 1, evaluation at $\boldsymbol{e}_z$ is a bijection between the global kernel space and the local intertwiner space. Proposition 1 gives the dimension of the latter space, which equals the number of allowed CG kernels. Thus, both constructions give bases of the same space, related by an invertible change of basis. □

The parameter reduction therefore removes coefficients incompatible with reflection equivariance without removing an allowed $\mathrm{O}(3)$ coupling.

B.6 Detailed Message-Passing Architecture

We follow SLEM's strictly localized architecture (Zhouyin et al., 2025) and replace its angular maps with EACO. We use the graph notation and four feature types from Section 3. At layer $t = 0, \dots, T$, these are node features $\boldsymbol{n}_i^{(t)}$, edge features $\boldsymbol{e}_{ij}^{(t)}$, directed hidden features $\boldsymbol{h}_{ij}^{(t)}$, and invariant pair embeddings $\boldsymbol{z}_{ij}^{(t)}$.

For a unit direction $\boldsymbol{u}$, write an angular processing block as

$$\mathcal{B}^{(t)}(\boldsymbol{x}; \boldsymbol{u}) = \mathcal{L}^{(t)} \mathcal{A}^{(t)} \left(\boldsymbol{K}_{\boldsymbol{u}}^{(t)} \boldsymbol{x} \right). \tag{62}$$

Here, $\boldsymbol{K}_{\boldsymbol{u}}^{(t)}$ is an EACO kernel, $\mathcal{A}^{(t)}$ is a gated nonlinearity, and $\mathcal{L}^{(t)}$ mixes channels of the same irrep type. We suppress admissible biases and irrep-preserving input normalizations in the update equations. Invariant weights generated from pair embeddings are applied separately.

B.6.1 Initialization

For each directed edge, a two-body multilayer perceptron first computes

$$\widetilde{\boldsymbol{z}}_{ij}^{(0)} = f_{2\mathrm{b}}(\mathbf{1}(z_i) \,\|\, \mathbf{1}(z_j) \,\|\, \boldsymbol{B}(r_{ij})), \qquad \boldsymbol{z}_{ij}^{(0)} = u(r_{ij}) \widetilde{\boldsymbol{z}}_{ij}^{(0)}. \tag{63}$$

Here, $\mathbf{1}(z_i)$ is the one-hot encoding of the atom type, $\boldsymbol{B}$ is a trainable Bessel basis, and $\|$ denotes concatenation. The function u is the cutoff envelope. All these inputs are invariant under $\mathrm{O}(3)$.

Scalar weights from $\widetilde{\boldsymbol{z}}_{ij}^{(0)}$ multiply the fixed spherical harmonics $\boldsymbol{Y}^{(\ell)}(\hat{\boldsymbol{r}}_{ij})$ to initialize the hidden features. The edge features use weights from the cutoff-weighted embedding $\boldsymbol{z}_{ij}^{(0)}$. The initial node feature is the normalized sum of $\boldsymbol{e}_{ij}^{(0)}$ over $j \in \mathcal{N}(i)$. The sum is scaled by $\gamma_i = \bar{N}_{z_i}^{-1/2}$, where $\bar{N}_{z_i}$ is the fixed average neighbor count for atom type z_i. This average may be shared across atom types. Hidden and node features then use irrep-preserving normalization.

B.6.2 LAYER UPDATES

Let $\mathcal{B}_h^{(t)}$, $\mathcal{B}_e^{(t)}$, and $\mathcal{B}_n^{(t)}$ be instances of Equation (62) for the hidden, edge, and node updates. Their input and output irreps may differ. We use a tilde for intermediate features before scalar weighting or residual addition. The hidden and pair updates begin with

$$\widetilde{\boldsymbol{h}}_{ij}^{(t)} = \mathcal{B}_h^{(t)}\left(\boldsymbol{n}_i^{(t-1)} \parallel \boldsymbol{h}_{ij}^{(t-1)}; \hat{\boldsymbol{r}}_{ij}\right), \tag{64}$$

$$\widetilde{\boldsymbol{z}}_{ij}^{(t)} = u(r_{ij}) f_z^{(t)}\left(\mathbf{1}(z_i) \parallel \boldsymbol{z}_{ij}^{(t-1)} \parallel [\widetilde{\boldsymbol{h}}_{ij}^{(t)}]_{0,+1} \parallel \mathbf{1}(z_j)\right). \tag{65}$$

The projection $[\cdot]_{0,+1}$ extracts channels with angular degree zero and inversion parity $+1$. In the reported configurations, these even scalars enter the multilayer perceptron $f_z^{(t)}$ before pair-dependent weighting. Layer normalization of the invariant pair input is included in $f_z^{(t)}$.

Let $\mathcal{W}_h^{(t)}$, $\mathcal{W}_e^{(t)}$, and $\mathcal{W}_n^{(t)}$ generate invariant channel weights from the pair embedding. Each weight is shared by all magnetic components of its irrep channel. Let $\mathcal{R}_h^{(t)}$, $\mathcal{R}_e^{(t)}$, $\mathcal{R}_n^{(t)}$, and $\mathcal{R}_z^{(t)}$ denote residual updates. Each combines a new feature with the previous feature through scalar coefficients whose squared sum is one. An equivariant linear map matches the old and new feature spaces when needed. The hidden and pair features are then

$$\begin{aligned} \boldsymbol{h}_{ij}^{(t)} &= \mathcal{R}_h^{(t)}\left(\boldsymbol{h}_{ij}^{(t-1)}, \mathcal{W}_h^{(t)}(\boldsymbol{z}_{ij}^{(t-1)})\widetilde{\boldsymbol{h}}_{ij}^{(t)}\right), \\ \boldsymbol{z}_{ij}^{(t)} &= \mathcal{R}_z^{(t)}\left(\boldsymbol{z}_{ij}^{(t-1)}, \widetilde{\boldsymbol{z}}_{ij}^{(t)}\right). \end{aligned} \tag{66}$$

These two residual updates share their scalar coefficients. Normalization of the pair input to $\mathcal{W}_h^{(t)}$ is included in this weight map.

The edge update uses both endpoint features,

$$\begin{aligned} \widetilde{\boldsymbol{e}}_{ij}^{(t)} &= \mathcal{W}_e^{(t)}(\boldsymbol{z}_{ij}^{(t)})\mathcal{B}_e^{(t)}\left(\boldsymbol{n}_i^{(t-1)} \parallel \boldsymbol{h}_{ij}^{(t)} \parallel \boldsymbol{n}_j^{(t-1)}; \hat{\boldsymbol{r}}_{ij}\right), \\ \boldsymbol{e}_{ij}^{(t)} &= \mathcal{R}_e^{(t)}\left(\boldsymbol{e}_{ij}^{(t-1)}, \widetilde{\boldsymbol{e}}_{ij}^{(t)}\right). \end{aligned} \tag{67}$$

Finally, the node message and aggregation are

$$\begin{aligned} \boldsymbol{m}_{ij}^{(t)} &= \mathcal{W}_n^{(t)}(\boldsymbol{z}_{ij}^{(t)})\mathcal{B}_n^{(t)}\left(\boldsymbol{n}_i^{(t-1)} \parallel \boldsymbol{h}_{ij}^{(t)}; \hat{\boldsymbol{r}}_{ij}\right), \\ \boldsymbol{n}_i^{(t)} &= \mathcal{R}_n^{(t)}\left(\boldsymbol{n}_i^{(t-1)}, \gamma_i \sum_{j \in \mathcal{N}(i)} \boldsymbol{m}_{ij}^{(t)}\right). \end{aligned} \tag{68}$$

Here, $\boldsymbol{m}_{ij}^{(t)}$ is the message for the edge centered at i.

The environment of atom i contains i and its fixed cutoff neighbors. Initialization uses only this environment. The hidden, pair, and node updates at i use neither $\boldsymbol{n}_j^{(t-1)}$ nor $\boldsymbol{e}_{ij}^{(t-1)}$. Their dependence therefore stays within this environment at every layer. Only the edge update combines the environments of i and j. It does not feed back into hidden or node updates, so these two environments also remain fixed.

B.6.3 QUANTUM-OPERATOR READOUT

After T layers, linear readouts give $\hat{\boldsymbol{o}}^{ii} = \mathrm{Readout}(\boldsymbol{n}_i^{(T)})$ and $\hat{\boldsymbol{o}}^{ij} = \mathrm{Readout}(\boldsymbol{e}_{ij}^{(T)})$. These are the coupled components used in Section 3. For each shell pair $(n_1\ell_1, n_2\ell_2)$, the output contains one component for every degree

$$|\ell_1 - \ell_2| \le \ell_3 \le \ell_1 + \ell_2, \qquad p_3 = (-1)^{\ell_1+\ell_2}. \tag{69}$$

Different shell pairs have separate output channels. The fixed CG coefficients in Equation (2) convert these components to the predicted orbital blocks $\widehat{\boldsymbol{O}}^{ij}_{n_1\ell_1,n_2\ell_2}$. This output transformation adds no learnable parameters.

B.6.4 FORWARD PROCEDURE

Algorithm 2 summarizes the forward procedure, using Algorithm 1 for each angular block.

Algorithm 2 $\mathrm{O}(3)$-Equivariant MPNN

Require: Atomistic graph $\mathcal{G}$ and target orbital basis
Construct the cutoff graph, lengths r_{ij}, and directions $\hat{\boldsymbol{r}}_{ij}$
Initialize $\boldsymbol{z}_{ij}^{(0)}$, $\boldsymbol{h}_{ij}^{(0)}$, $\boldsymbol{e}_{ij}^{(0)}$, and $\boldsymbol{n}_i^{(0)}$
for $t = 1, \ldots, T$ **do**
  Compute $\widetilde{\boldsymbol{h}}_{ij}^{(t)}$ and $\widetilde{\boldsymbol{z}}_{ij}^{(t)}$ by Equation (64) and Equation (65)
  Update $\boldsymbol{h}_{ij}^{(t)}$ and $\boldsymbol{z}_{ij}^{(t)}$ by Equation (66)
  Update $\boldsymbol{e}_{ij}^{(t)}$ by Equation (67)
  Compute node messages and update $\boldsymbol{n}_i^{(t)}$ by Equation (68)
end for
Read out the coupled components $\hat{\boldsymbol{o}}^{ii}$ and $\hat{\boldsymbol{o}}^{ij}$
Assemble orbital blocks by Equation (2)
Ensure: Predicted quantum operator

B.7 Equivariance of the Complete Network

Proof of Corollary 1. Atom types, distances, radial functions, and cutoff coefficients are invariant under $\mathrm{O}(3)$. The initial pair embeddings are therefore invariant. Real spherical harmonics transform in $V_{\ell,(-1)^\ell}$, so scalar-weighted initialization gives equivariant hidden, edge, and node features.

Assume that these properties hold at layer $t-1$. Proposition 1 and Theorem 1 make each EACO operation equivariant. The channel maps mix only equal irrep types. Gates have type $(0,+1)$ and preserve the transformation law of the features they multiply. An odd scalar channel, if present, uses an odd activation. Feature normalization uses invariant norms, with additive shifts only in even scalar channels.

The projection $[\widetilde{\boldsymbol{h}}_{ij}^{(t)}]_{0,+1}$ is invariant. Thus, the pair update remains invariant, and all weights generated from it are invariant scalars. Residual addition, the fixed factors γ_i, and neighbor sums preserve equivariance. This completes the induction for all four feature types.

The final node and edge readouts are equivariant linear maps, with biases only in even scalar channels. CG reconstruction is a fixed intertwiner from the coupled components to the orbital basis. Therefore, the predicted blocks satisfy Equation (5) for every $\boldsymbol{Q} \in \mathrm{O}(3)$. □

B.8 Parameter Count and Computational Complexity

Let $C_a^{\pm}(s)$ count the order-s channels with reflection signs ± 1, for $a \in \{\mathrm{in}, \mathrm{out}\}$. Their total is $C_a(s) = C_a^{+}(s) + C_a^{-}(s)$. Using the type multiplicities $C_{\ell,p}^a$ defined in Section 3.1, we have

$$C_a(s) = \sum_{\ell=s}^{\ell_{\max}} \sum_{p\in\{\pm 1\}} C_{\ell,p}^a, \qquad a \in \{\mathrm{in}, \mathrm{out}\}. \tag{70}$$

For $s > 0$, each counted channel has two real components. At $s = 0$, it has one.

Let $P_s^{\mathrm{SO}(2)}$ and P_s^{EACO} count kernel weights at order s. For $s > 0$, each channel pair has two $\mathrm{SO}(2)$ coefficients and one EACO coefficient. Thus,

$$P_s^{\mathrm{SO}(2)} = 2C_{\mathrm{in}}(s)C_{\mathrm{out}}(s), \tag{71}$$

$$P_s^{\mathrm{EACO}} = (C_{\mathrm{out}}^{+}(s) + C_{\mathrm{out}}^{-}(s))(C_{\mathrm{in}}^{+}(s) + C_{\mathrm{in}}^{-}(s)) = C_{\mathrm{in}}(s)C_{\mathrm{out}}(s). \tag{72}$$

Thus, the positive-order kernel uses half as many weights. At zero order, only equal-sign weights remain:

$$P_0^{\mathrm{EACO}} = C_{\mathrm{out}}^{+}(0)C_{\mathrm{in}}^{+}(0) + C_{\mathrm{out}}^{-}(0)C_{\mathrm{in}}^{-}(0). \tag{73}$$

There are also $C_{\mathrm{out}}^{+}(0)$ admissible bias parameters. These counts concern the local kernel. The radial networks, gates, channel maps, and readouts contribute additional parameters to the complete model.

Assume that each $C^a_{\ell,p}$ is bounded independently of $\ell_{\max}$. Then $C_a(s) = \mathcal{O}(\ell_{\max} - s + 1)$, and the order-wise channel mixing costs

$$\sum_{s=0}^{\ell_{\max}} C_{\text{in}}(s) C_{\text{out}}(s) = \mathcal{O}\left(\sum_{s=0}^{\ell_{\max}} (\ell_{\max} - s + 1)^2 \right) = \mathcal{O}(\ell^3_{\max}). \tag{74}$$

Applying the alignment and inverse-alignment matrices to the features costs $\mathcal{O}(\sum_{\ell=0}^{\ell_{\max}} \sum_p (C^{\text{in}}_{\ell,p} + C^{\text{out}}_{\ell,p})(2\ell+1)^2) = \mathcal{O}(\ell^3_{\max})$. For T layers and $|\mathcal{E}|$ directed edges, these angular feature operations therefore cost

$$\mathcal{O}(T|\mathcal{E}|\ell^3_{\max}). \tag{75}$$

This count treats the frame matrices as given. Their geometry-dependent construction is shared with SLEM. The efficiency measurements include this construction, both feature rotations, and local channel mixing. The reflection constraints change only the local mixing step. With a fixed cutoff and bounded neighbor count, $|\mathcal{E}| = \mathcal{O}(|\mathcal{V}|)$, so these operations scale linearly with the number of atoms.

B.9 Blockwise Evaluation

Algorithm 1 rotates the input into the local frame before applying the order-wise blocks. For $s > 0$, let $\bar{\boldsymbol{x}}^{\pm}_s$ and $\bar{\boldsymbol{y}}^{\pm}_s$ collect the local input and output channels with reflection signs ± 1. Partition $\boldsymbol{A}_s$ and $\boldsymbol{B}_s$ by output and input signs, in that order. The equal-sign blocks are $\boldsymbol{A}^{++}_s$ and $\boldsymbol{A}^{--}_s$, and the opposite-sign blocks are $\boldsymbol{B}^{+-}_s$ and $\boldsymbol{B}^{-+}_s$. Then Equation (13) becomes

$$\begin{aligned} \bar{\boldsymbol{y}}^{+}_s &= (\boldsymbol{A}^{++}_s \otimes \boldsymbol{I}_2)\bar{\boldsymbol{x}}^{+}_s + (\boldsymbol{B}^{+-}_s \otimes \boldsymbol{J}_2)\bar{\boldsymbol{x}}^{-}_s, \\ \bar{\boldsymbol{y}}^{-}_s &= (\boldsymbol{B}^{-+}_s \otimes \boldsymbol{J}_2)\bar{\boldsymbol{x}}^{+}_s + (\boldsymbol{A}^{--}_s \otimes \boldsymbol{I}_2)\bar{\boldsymbol{x}}^{-}_s. \end{aligned} \tag{76}$$

The implementation selects these four submatrices from one dense coefficient matrix. The signs, channel partitions, $\boldsymbol{I}_2$, and $\boldsymbol{J}_2$ are fixed, while the stored coefficients are learned. The forward pass evaluates only allowed blocks, without multiplying entries masked to zero.

At $s = 0$, the corresponding expressions are

$$\bar{\boldsymbol{y}}^{+}_0 = \boldsymbol{A}^{++}_0 \bar{\boldsymbol{x}}^{+}_0 + \boldsymbol{b}^{+}_0, \qquad \bar{\boldsymbol{y}}^{-}_0 = \boldsymbol{A}^{--}_0 \bar{\boldsymbol{x}}^{-}_0, \tag{77}$$

where $\boldsymbol{b}^{+}_0$ is the bias for the positive-sign output channels. The two equal-sign weight blocks are stored separately. After collecting the local outputs, the inverse frame rotation returns them to the global frame.

Remark (Projection of an existing $\mathrm{SO}(2)$ kernel)**.** A compatible $\mathrm{SO}(2)$ kernel can be mapped to the $\mathrm{O}(2)$ kernel space by the Reynolds projection

$$\mathcal{P}_{\mathrm{O}(2)}(\boldsymbol{K}_{\boldsymbol{e}_z}) = \frac{1}{2}\left(\boldsymbol{K}_{\boldsymbol{e}_z} + \boldsymbol{P}_{\text{out}}(\boldsymbol{Q}_{\boldsymbol{e}_z})\boldsymbol{K}_{\boldsymbol{e}_z}\boldsymbol{P}_{\text{in}}(\boldsymbol{Q}_{\boldsymbol{e}_z})^{-1}\right). \tag{78}$$

It retains the $\boldsymbol{I}_2$ coefficients for equal signs and the $\boldsymbol{J}_2$ coefficients for opposite signs. An affine bias $\boldsymbol{b}$ is projected separately as $\frac{1}{2}(\boldsymbol{b} + \boldsymbol{P}_{\text{out}}(\boldsymbol{Q}_{\boldsymbol{e}_z})\boldsymbol{b})$. Rotation invariance already restricts this bias to zero order, and the projection retains only the positive-sign channels. Our reported models are trained from scratch and do not use this initialization.

C WHY O(3) EQUIVARIANCE MATTERS FOR QUANTUM OPERATORS

Full O(3) equivariance requires the transformation law in Equation (5) for rotations, reflections, and spatial inversion. This requirement applies to arbitrary atomic structures. When a reflection or inversion is also a symmetry of the structure, it constrains the operator itself. We examine the inversion parities present in the DFT labels, then explain how these symmetries constrain physical predictions.

C.1 PARITY CONTENT OF THE HAMILTONIAN LABELS

We analyze the DFT off-site Hamiltonian labels on the five training splits. A shell pair with angular degrees ℓ_1 and ℓ_2 has inversion parity $p = (-1)^{\ell_1+\ell_2}$. Different shell pairs can therefore contribute to the same coupled degree ℓ with opposite parities. For edge $i \to j$, let $\boldsymbol{o}^{ij}_{\ell,p,c} \in \mathbb{R}^{2\ell+1}$ denote a coupled component of its DFT label. Here, c indexes valid shell pairs in the reduced orbital-pair representation used for training. The power fraction of type (ℓ, p) is

$$P_{\ell,p} = \frac{\sum_{i\to j,c} \|\boldsymbol{o}^{ij}_{\ell,p,c}\|_2^2}{\sum_{\ell',p'} \sum_{i\to j,c} \|\boldsymbol{o}^{ij}_{\ell',p',c}\|_2^2}. \tag{79}$$

The edge sums run over all training structures of each material and exclude padded entries. The orthogonal CG transform preserves the squared Frobenius norm of the included orbital blocks (Appendix A.3).

Let $N_{\ell,p}$ count the valid scalar coefficients of type (ℓ, p) over these edges. The angular degrees containing both parities form the set

$$\mathcal{L}_{\text{mixed}} = \{\ell \mid N_{\ell,+1} > 0 \text{ and } N_{\ell,-1} > 0\}. \tag{80}$$

Their shares of label power and coefficient count are

$$P_{\text{mixed}} = \sum_{\ell\in\mathcal{L}_{\text{mixed}}} (P_{\ell,+1} + P_{\ell,-1}), \tag{81}$$

$$C_{\text{mixed}} = \frac{\sum_{\ell\in\mathcal{L}_{\text{mixed}}} (N_{\ell,+1} + N_{\ell,-1})}{\sum_{\ell,p} N_{\ell,p}}. \tag{82}$$

Figure 6 shows these two shares, separated by inversion parity. We also measure the smaller parity contributions and their norm fraction:

$$P_{\min} = \sum_{\ell\in\mathcal{L}_{\text{mixed}}} \min(P_{\ell,+1}, P_{\ell,-1}), \tag{83}$$

$$A_{\min} = \sqrt{P_{\min}}. \tag{84}$$

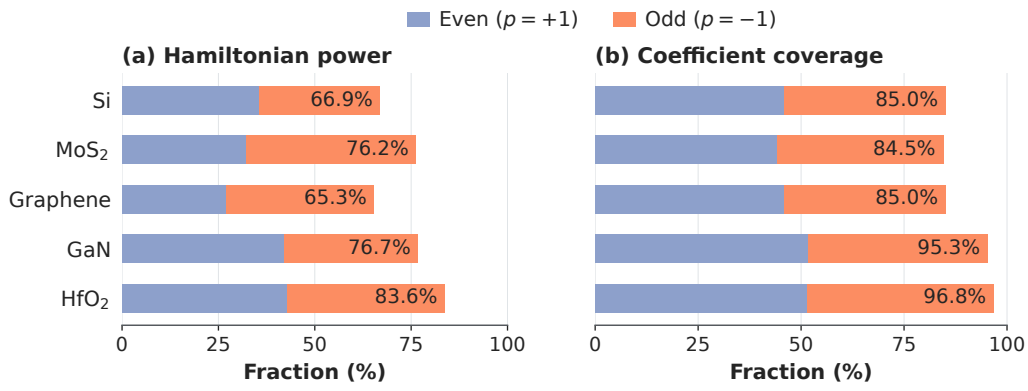

Figure 6: **Inversion-parity content of the DFT Hamiltonian labels.** For angular degrees containing both parities, we report (a) the share of Hamiltonian power, P_{mixed}, and (b) the share of valid coefficients, C_{mixed}. Colors separate the parity contributions, and numbers give the bar totals. Both panels use the reduced off-site labels from each training split.

Table 7: Parity statistics of the DFT training labels. The last four columns report percentages.

Material	Structures	Directed edges	$100P_{\mathrm{mixed}}$	$100P_{\mathrm{min}}$	$100A_{\mathrm{min}}$	$100C_{\mathrm{mixed}}$
Si	150	1,196,912	66.88	2.32	15.23	85.05
MoS_2	300	1,237,044	76.18	2.82	16.78	84.47
Graphene	270	875,054	65.31	1.43	11.95	85.05
GaN	150	3,042,872	76.68	2.38	15.43	95.27
HfO_2	150	3,266,004	83.63	2.98	17.28	96.75

Figure 6 and Table 7 show that degrees with both parities account for 65.31%–83.63% of the label power and 84.47%–96.75% of the valid coefficients. Even the smaller parity contributions carry 1.43%–2.98% of the total power, with norm fractions of 11.95%–17.28%. The $\mathrm{SO}(2)$ rotation action does not distinguish their inversion parities, as shown in Equation (41). EACO retains this distinction through the reflection sign $\tau_{\ell,p} = p(-1)^{\ell}$.

C.2 Symmetry Constraints on Physical Observables

A small matrix-element MAE does not enforce an exact symmetry constraint. Let $\boldsymbol{Q}$ be a reflection or inversion symmetry of $\mathcal{G}$. Its orbital action $\boldsymbol{D}(\boldsymbol{Q})$ includes permutations of atoms and periodic images. For spinless orbitals, $\boldsymbol{D}(\boldsymbol{Q})$ is orthogonal and satisfies $\boldsymbol{D}(\boldsymbol{Q})^2 = \boldsymbol{I}$, where $\boldsymbol{I}$ is the identity. The exact Hamiltonian satisfies

$$\boldsymbol{H}[\mathcal{G}] = \boldsymbol{D}(\boldsymbol{Q})\boldsymbol{H}[\mathcal{G}]\boldsymbol{D}(\boldsymbol{Q})^{\top}. \tag{85}$$

For a matrix $\boldsymbol{X}$ on this basis, define its symmetry-preserving and symmetry-breaking parts as

$$\Pi^{\pm}_{\boldsymbol{Q}}[\boldsymbol{X}] = \frac{1}{2}\left(\boldsymbol{X} \pm \boldsymbol{D}(\boldsymbol{Q})\boldsymbol{X}\boldsymbol{D}(\boldsymbol{Q})^{\top}\right). \tag{86}$$

The two projections are orthogonal. The symmetry constraint in Equation (85) gives

$$\Pi^{-}_{\boldsymbol{Q}}[\boldsymbol{H}[\mathcal{G}]] = \boldsymbol{0}. \tag{87}$$

For a prediction error $\Delta\boldsymbol{H} = \widehat{\boldsymbol{H}} - \boldsymbol{H}$, we have

$$\|\Delta\boldsymbol{H}\|_F^2 = \|\Pi^{+}_{\boldsymbol{Q}}[\Delta\boldsymbol{H}]\|_F^2 + \|\Pi^{-}_{\boldsymbol{Q}}[\Delta\boldsymbol{H}]\|_F^2. \tag{88}$$

Thus, a small total error does not force the symmetry-breaking part to vanish. This part can produce a forbidden coupling.

Graphene gives a direct example. Let η be the amplitude of an out-of-plane displacement pattern. Its first-order Hamiltonian response at equilibrium is $\boldsymbol{H}' = \partial\boldsymbol{H}/\partial\eta|_{\eta=0}$. The horizontal mirror σ_h sends η to $-\eta$, so

$$\boldsymbol{D}(\sigma_h)\boldsymbol{H}'\boldsymbol{D}(\sigma_h)^{\top} = -\boldsymbol{H}'. \tag{89}$$

Matrix elements of $\boldsymbol{H}'$ between states with the same mirror eigenvalue must therefore vanish (Mariani & von Oppen, 2008; Alidoosti et al., 2022). This forbids first-order π–π EPC, while allowing π–σ EPC. Section 4.2 tests both channels using common DFT electronic states and phonon data. EACO preserves the forbidden zero and improves the allowed coupling predictions.

The same constraints matter in other materials applications. Inversion symmetry forbids the conventional electric-dipole shift current under uniform illumination (Yang et al., 2026). At finite photon momentum, photon drag can still produce a current (Shi et al., 2021). High-throughput Hamiltonian studies screen noncentrosymmetric materials for strong nonlinear optical responses (Yang et al., 2026; Xu et al., 2020). Crystal symmetry determines which response-tensor components are allowed (Yang et al., 2026). In the SnTe material class, mirror symmetry protects surface Dirac points. A mirror-breaking Dirac mass m_{D} opens a gap of $2|m_{\mathrm{D}}|$ (Hsieh et al., 2012). Experiments on (Pb,Sn)Se found a surface Dirac gap of about $25\,\mathrm{meV}$, consistent with a mirror-breaking surface distortion (Wojek et al., 2015). Crystal symmetry also constrains piezoelectric and dielectric tensors (Choudhary et al., 2020). High-throughput DFT workflows use these relations to check numerical results (de Jong et al., 2015). In WTe_2, lattice shear can change inversion symmetry and annihilate Weyl points (Sie et al., 2019). These examples show that symmetry-breaking Hamiltonian terms can change physical responses or lift protected crossings. Small matrix-element errors alone do not preserve these properties.

D Experimental Details and Supplementary Evaluation

D.1 Datasets and Training

Datasets. We use the same five datasets as SLEM (Zhouyin et al., 2025). Each contains atomic structures and the corresponding DFT Hamiltonian blocks. The bulk datasets, Si, GaN, and HfO_2, each contain 150 training and 30 validation structures. The monolayer datasets have training/validation/test splits of 270/90/90 for graphene and 300/100/100 for MoS_2. Table 8 lists the orbital bases and cutoffs.

Table 8: **Orbital bases and cutoffs.** All methods use the same basis and cutoff within each material. The number before each orbital type is the number of radial functions.

Material	Orbital basis by element	$r_{\max}$ (Å)
Si	Si: $2s2p1d$	8.5
MoS_2	Mo: $3s2p2d$, S: $2s2p1d$	7.4085
Graphene	C: $2s2p1d$	6.4
GaN	Ga: $2s2p2d1f$, N: $2s2p1d$	10.5 (Ga), 8.5 (N)
HfO_2	Hf: $4s2p2d2f1g$, O: $2s2p1d$	8.5

Model and Training Configurations. Table 9 lists the core model configurations. SLEM and EACO use the same settings apart from the local convolution. DeepH-E3 and MACE-H use their native architectures and training settings[3] for up to 1,500 epochs. We train SLEM and EACO for 1,000 epochs with Adam, batch size one, and an initial learning rate of 5×10^{-3}. ReduceLROnPlateau reduces the learning rate by a factor of 0.8, down to 10^{-6}.

Table 9: **Model configurations.** Tuples list hidden-channel multiplicities for $\ell = 0, 1, \ldots$.

Setting	DeepH-E3	MACE-H	SLEM / EACO
Si	$(64, 32, 16, 8, 8)$	$(64, 64, 64)$	$(28, 24, 24, 8, 4)$
MoS_2	$(96, 48, 24, 12, 12)$	$(64, 64, 64, 64, 64, 64)$	$(21, 18, 18, 6, 3)$
Graphene	$(64, 32, 16, 8, 8)$	$(64, 64, 64, 64, 64, 64)$	$(21, 18, 18, 6, 3)$
GaN	$(117, 58, 29, 15, 15)$	$(64, 64, 64)$	$(30, 30, 39, 24, 18, 6, 3)$
HfO_2	$(107, 53, 27, 13, 13)$	$(64, 64, 64)$	$(20, 18, 24, 20, 17, 8, 6, 2, 1)$
Layers	3	3	4
Edge / pair width	64	64	64
Radial functions	128	128	18 (Si), 10 (others)

Edge / pair width refers to initial scalar edge features in DeepH-E3 and MACE-H, and invariant pair embeddings in SLEM and EACO.

D.2 Graphene Responses

Mirror Test. We use the graphene checkpoints in Table 2 and eight evenly spaced validation structures. We make each structure planar by setting its vertical coordinates to their mean. The ZA displacement follows a cosine at the smallest nonzero supercell wavevector. The ZO displacement assigns opposite signs to the two sublattices. Both patterns have zero mean and unit maximum amplitude. We apply paired displacements of ± 0.01 Å on a shared graph. For each mode, $\epsilon^{\max}_{\sigma_h}$ is the maximum residual in Equation (14) over the eight pairs, with $\boldsymbol{Q} = \sigma_h$. Model predictions and EPC calculations use float64. Mirror tests use deterministic CPU execution. For SLEM and EACO, Wigner matrices are constructed from direction cosines through trigonometric recurrences. A reported zero means exact equality of the computed matrices.

DFT Settings. Reference Hamiltonians use OpenMX with PBE, the C6.0-$s2p2d1$ basis, and the C_PBE19 pseudopotential. The integration cutoff is 300 Ry, and the electronic temperature is 300 K. The primitive cell has lattice constant 2.46 Å, vertical cell size 30 Å, and a $21 \times 21 \times 1$ electronic mesh. Phonons use Quantum ESPRESSO DFPT on a $10 \times 10 \times 1$ mesh. We use a $24 \times 24 \times 1$ electronic mesh and wavefunction and charge-density cutoffs of 40 and 320 Ry. Rotational and translational acoustic sum rules are enforced.

[3] DeepH-E3: https://github.com/Xiaoxun-Gong/DeepH-E3.
MACE-H: https://github.com/maurergroup/MACE-H.

EPC Evaluation. We fix the initial electronic wavevector at $\boldsymbol{k} = (1/3, 1/3, 0)$ and sample $\boldsymbol{q}$ on an $8 \times 8 \times 1$ grid. The zero-frequency ZA mode at Γ is excluded. Hamiltonian and overlap derivatives use central differences at displacement amplitudes $\eta \in \{0.005, 0.01, 0.02\}$ Å, weighted by η^2. Let ∂_ν denote the Fourier-transformed derivative contracted with the unit-normalized eigenvector of phonon mode ν. The normalized coupling in the symmetric orthogonalization convention is

$$g_{mn\nu}(\boldsymbol{k}, \boldsymbol{q}) = \sqrt{\frac{\hbar}{2M_{\mathrm{C}}\omega_{\boldsymbol{q}\nu}}}\, \boldsymbol{c}^{\dagger}_{m,\boldsymbol{k}+\boldsymbol{q}} \left[\partial_\nu \widehat{\boldsymbol{H}} - \frac{\varepsilon_{m,\boldsymbol{k}+\boldsymbol{q}} + \varepsilon_{n,\boldsymbol{k}}}{2} \partial_\nu \boldsymbol{S}^{\mathrm{DFT}}\right] \boldsymbol{c}_{n,\boldsymbol{k}}, \tag{90}$$

where $\boldsymbol{S}^{\mathrm{DFT}}$ is the overlap matrix, M_{C} is the carbon mass, and $\omega_{\boldsymbol{q}\nu}$ is the phonon angular frequency. The indices n, m denote the initial and final electronic bands, with DFT eigenvectors $\boldsymbol{c}$ and energies ε. All methods share the DFT states, energies, overlap derivatives, and phonon data. Only the Hamiltonian derivative is replaced by the model prediction. The initial states are the two Dirac bands. The final states are the two nearest π bands and four frontier σ bands. We report the RMS forbidden π–π coupling over these states and $\boldsymbol{q}$ points. For the allowed π–σ couplings $\boldsymbol{g}_{\mathrm{a}}$, we report $100\|\boldsymbol{g}_{\mathrm{a}} - \boldsymbol{g}_{\mathrm{a}}^{\mathrm{DFT}}\|_F / \|\boldsymbol{g}_{\mathrm{a}}^{\mathrm{DFT}}\|_F$ in percent.

D.3 Angular-Operator Benchmark

We fix 4,096 directed edges and vary $\ell_{\max}$ from 1 to 11. Input and output features use 16 channels of each parity at every degree $\ell = 0, \ldots, \ell_{\max}$. We compare the $\mathrm{SO}(2)$ and $\mathrm{O}(2)$ operators of SLEM and EACO with the CG operator shared by DeepH-E3 and MACE-H. The CG operator uses spherical harmonics up to degree $2\ell_{\max}$. It has the same number of independent linear coefficients as EACO. Measurements include direction encoding, Wigner matrix construction, and feature rotations. Radial modulation is disabled.

All tests run on an idle NVIDIA A100-SXM4-80GB in float32 with TF32 disabled. After at least 15 warm-up calls, we measure synchronized forward-plus-backward time and peak allocated GPU memory. The memory measurement includes resident tensors and intermediate allocations. Figure 5(a) reports medians over three independent runs for each setting.