EPOCH: EXPLICIT PROJECTION ON CONCENTRIC SHELLS FOR MULTI-ORDER MOLECULAR GEOMETRY

Anonymous authors
Paper under double-blind review

ABSTRACT

For many molecular properties, accurate 3D prediction depends on geometric information across multiple angular scales, from coarse global shape to fine directional detail around local environments. Modern equivariant networks expose this structure through spherical harmonics up to a chosen maximum order $l_{\max}$, but tensor-product cost grows as $\mathcal{O}(l_{\max}^6)$ per interaction, pushing practical models to a low-order operating point even when chemistry suggests finer angular content matters. We argue that raising $l_{\max}$ is not the only access path to angular structure: finite directional sampling can empirically preserve a useful mid-frequency angular band beyond the practical low-order operating point, at cost linear in the number of probes rather than polynomial in equivariant order. We instantiate this view as EPOCH (Explicit Projection on Concentric Shells), a shared geometric descriptor obtained by sampling a shellwise zonal-kernel response field, precomputed once from coordinates and inserted into scalar, angular, coordinate-equivariant, and tensor-equivariant backbones through lightweight backbone-aware fusion. On QM9, comparisons across SchNet, DimeNet++, EGNN, NequIP, and the strictly local equivariant Allegro show that this shared descriptor helps within heterogeneous backbone families, with the clearest gains concentrated on angular-sensitive electronic targets. Parameter-matched HOMO controls show that preserving response-to-direction correspondence improves over capacity-matched variants, while also clarifying that the effect is target-dependent. Order ablations on a tensor-equivariant backbone further show that EPOCH continues to help as native equivariant order increases, and that augmenting a low-order backbone with EPOCH can often match or remain competitive with a higher-order baseline at a fraction of its parameter and compute cost. Cross-dataset, cross-task evidence on MD17 further shows that the same design improves force prediction. Together, these results indicate that explicit spherical sampling can provide a portable, target-dependent complement to tensor-product equivariance rather than a replacement for it.

1 INTRODUCTION

For many molecular properties, accurate 3D prediction depends not only on distances and composition, but also on angular structure across multiple scales. Quantities such as orbital energies, polarizability, and dipole moments depend not only on which atoms are present, but also on how those atoms are arranged in space. Coarse molecular shape, directional orientation, and fine local angular detail can all change the electronic response, even when two molecules have similar 2D connectivity. The machine learning challenge is therefore not merely to use geometry, but to expose its directional content at the right resolution to the predictor, a recurring theme in recent molecular representation learning (Sheshanarayana & You, 2025).

One principled route to finer angular structure in modern 3D molecular modeling is to raise the maximum spherical-harmonic order $l_{\max}$ in equivariant features. In $\mathrm{SO}(3)$-based architectures, $l_{\max}$ directly controls how much angular content can be represented internally (Batzner et al., 2022). However, the cost of doing so grows steeply with order, because higher-order tensor products dominate both parameter count and per-step compute (Batatia et al., 2025). As a result, practical models usually operate at relatively low order even when chemistry suggests that finer angular structure may matter. This tension motivates recent efforts to scale higher-degree equivariant representations more

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.

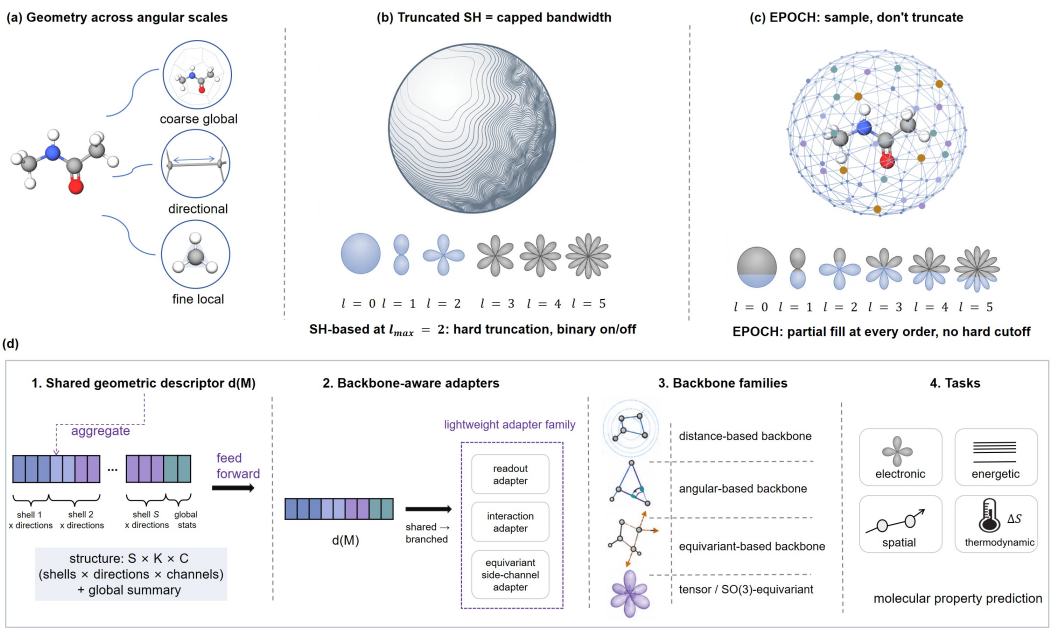

Figure 1: Overview of the representation view behind EPOCH. **(a)** Molecular targets depend on geometry across multiple angular scales, from coarse global shape to directional organization and fine local angular detail. **(b)** A truncated spherical-harmonic basis imposes a hard bandwidth cutoff: once $l_{\max}$ is fixed, higher-order angular content is simply absent. **(c)** EPOCH instead samples a shellwise directional response field, so each probe mixes contributions across orders and the representation is not tied to one hard truncation boundary. **(d)** The resulting shared descriptor is aggregated once, passed through a lightweight backbone-aware adapter family, and attached across heterogeneous backbone classes, including strictly local equivariant backbones such as Allegro, and across molecular prediction tasks.

effectively or to revisit spherical sampling inside the interaction design (Liao et al., 2024; An et al., 2025).

The question in this paper is therefore simple: when practical equivariant models operate at low order, can explicit directional sampling recover geometric information that the backbone does not fully expose? We argue that the answer is yes. Increasing equivariant order is not the only access path to angular structure; a spherical response field can also be probed through a finite set of directional samples. In practice, this retains a useful angular band at much lower marginal cost than increasing equivariant order directly (Figure 1b,c), and it suggests a shared geometric representation that can be attached across heterogeneous backbones with backbone-aware fusion.

We address this gap with EPOCH (Explicit Projection on Concentric Shells), a shared multi-order geometric descriptor built from shellwise directional projection responses. For each molecule, EPOCH evaluates a kernelized angular response across multiple radial shells and directional probes, producing an explicit summary of coarse-to-fine angular structure. In continuous terms, the underlying response field gives the construction a harmonic interpretation; in practical terms, EPOCH is a finite directional sampling of that field, precomputed once from molecular coordinates. EPOCH is used as a complementary side channel and inserted through lightweight backbone-aware fusion modules that respect each model's native state space.

This portability also gives a second analysis axis beyond raw accuracy: in NequIP, varying $l_{\max}$ lets the same descriptor serve as a controlled probe of representational overlap.

Our work makes four contributions. First, we propose EPOCH, a shared descriptor that explicitly encodes molecular geometry through shellwise directional sampling, with a harmonic interpretation that supports its practical angular complementarity beyond the lowest-order practical equivariant regime. Second, we design a lightweight backbone-aware fusion scheme and evaluate it across scalar,

angular, message-passing equivariant, tensor-equivariant, and strictly local equivariant backbones, including Allegro. Third, on QM9, parameter-matched HOMO controls provide a representative mechanism diagnostic showing that directional organization can contribute beyond the adapter pathway, while the target-dependent results clarify the scope of this interpretation. Fourth, EPOCH remains useful as native equivariant order rises in NequIP, lets lower-order models match higher-order baselines at $400\times$–$7800\times$ lower incremental FLOP cost, and transfers from QM9 property prediction to MD17 force prediction.

2 Related Work

3D molecular representation learning. Neural models for molecular property prediction have progressively moved from purely topological message passing to architectures that encode 3D geometry directly. SchNet introduced continuous-filter convolutions over interatomic distances, establishing a strong radial baseline for atomistic learning (Schütt et al., 2017). DimeNet then coupled distances and angles through directional message passing on triplets (Gasteiger et al., 2020b), and later models such as SphereNet and GemNet pushed toward richer distance-angle-torsion and higher-order geometric interactions (Liu et al., 2022; Gasteiger et al., 2021). These models show that richer geometry matters, but they realize it through architecture-specific interaction rules. Our work is complementary: rather than proposing another backbone-specific mechanism, we ask whether one explicit descriptor family can remain portable across heterogeneous backbones.

Equivariant and angular geometric modeling. Another line of work strengthens geometric inductive bias through Euclidean equivariance. EGNN showed that useful equivariant coordinate updates can be implemented with relatively simple message passing (Satorras et al., 2021); PaiNN introduced scalar-vector feature coupling (Schütt et al., 2021); and NequIP organized features through irreducible representations of $\mathrm{SO}(3)$, making angular order an explicit modeling knob (Batzner et al., 2022). Subsequent models expanded this design space through higher-order equivariant message passing, scalable local equivariant representations, and equivariant transformers (Batatia et al., 2022; Musaelian et al., 2023; Liao & Smidt, 2022; Liao et al., 2024). MACE combines message passing with high body-order equivariant interactions, whereas Allegro deliberately remains strictly local and avoids message passing (Batatia et al., 2022; Musaelian et al., 2023). This contrast makes Allegro a useful control for our descriptor-side study: it tests whether complementarity persists when strong local equivariant angular capacity is available without receptive-field expansion through message passing. Recent design-space analysis clarified the tradeoff surface among body order, equivariant degree, and computational efficiency (Batatia et al., 2025), while a recent preprint revisited spherical sampling inside equivariant interaction design (An et al., 2025). The closest conceptual neighbor to our NequIP study places spherical sampling *inside* an equivariant interaction mechanism, whereas EPOCH is precomputed from coordinates and injected *outside* the backbone as a descriptor-side side channel. This distinction makes EPOCH a portable descriptor-side signal that can attach across heterogeneous backbones and serve as a common probe of representational complementarity as native equivariant order increases; Appendix A.1 summarizes this positioning.

Designed geometric features and complementary 3D signals. A closely related question is whether explicitly designed geometric signals can complement learned backbone representations. Classical descriptor families such as SOAP and MBTR summarize molecular geometry through explicit kernelized or many-body constructions rather than through learned equivariant message passing (Bartók et al., 2013; Huo & Rupp, 2022). Those descriptors already encode angular information, but through compressed invariant summaries. EPOCH is designed around a different representational choice: explicit directional structure is preserved in the descriptor presented to the model rather than contracted away before learning. Related efforts have also injected richer angle and torsion statistics into molecular models or encouraged models to preserve geometric structure through pretraining (Wang et al., 2022; Liu et al., 2021; Zhou et al., 2023). Our goal is narrower and more comparative: whether one explicit descriptor family can expose multi-order angular structure, transfer across heterogeneous backbones, and support a common analysis of representational complementarity.

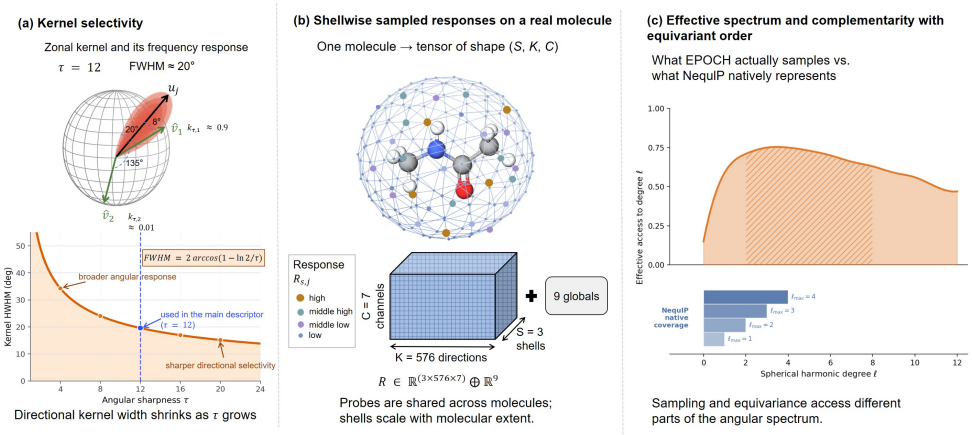

Figure 2: How EPOCH relates to angular structure beyond a fixed low-order equivariant budget. **(a)** The zonal kernel defines directional selectivity in real space and an associated frequency response across harmonic degrees, so the effective angular emphasis is controlled by kernel shape rather than by a hard harmonic truncation. **(b)** On a real molecule, shared probes evaluated across shells produce a tensor of shape (S, K, C) together with a small set of global summary statistics. **(c)** In the implemented setting, the resulting descriptor retains a practical mid-order angular band that overlaps only partially with the native coverage of practical low-order NequIP settings, so descriptor-side sampling and backbone equivariance operate as complementary representation channels across scalar, message-passing equivariant, and strictly local equivariant backbones.

3 METHOD

3.1 PROBLEM SETUP

Let a molecule be represented by atom types and coordinates $\mathcal{M} = \{(z_i, \mathbf{r}_i)\}_{i=1}^{n}$, and let f_θ denote a 3D molecular backbone that predicts a target property from $\mathcal{M}$. Existing backbones differ substantially in how they internalize geometry: some rely mainly on distances, some explicitly model angular interactions, and some maintain equivariant tensor features. Our goal is not to replace these native mechanisms with a single universal architecture. Instead, we seek an explicit geometric side channel $d(\mathcal{M})$ that can be attached across heterogeneous backbones and used to test which geometric content remains useful when exposed directly to the predictor (Figure 1d).

The method is organized in two layers. We first define a shared geometric descriptor family as a continuous shellwise projection field, since the main scientific object of the paper is the descriptor itself rather than any one backbone-specific implementation. We then derive a practical discretization and attach it to each backbone through a lightweight, backbone-aware fusion rule matched to the backbone's native state space. This separation fixes the scope of the claim: the geometric content is shared across backbones, while the fusion rule is only the local instantiation required for compatibility.

3.2 FROM PROJECTION FIELD TO DESCRIPTOR

EPOCH builds a shared geometric descriptor through a simple directional probing operation (Figure 2b). After centering and canonical alignment, we place a fixed set of probe directions on the unit sphere and reuse that same set for every molecule. For each direction and each radial shell, we ask how much atom-weighted signal lies in that directional cone and radial range. The answer is a soft kernel response rather than a hard count. Repeating this over shells and directions produces a shellwise directional response tensor. That tensor is then handed to the backbone as an explicit geometric side channel.

Two choices distinguish this picture from a standard low-order harmonic descriptor. First, angular resolution is controlled by the kernel width and the directional grid, not by a predeclared truncation order $l_{\max}$. Second, radial resolution is controlled by the shell scales, so one shared probing pattern exposes both coarse and fine geometry. Here *projection* refers to directional response onto fixed

probes rather than to truncation in a harmonic basis: each probe captures a soft directional response that mixes contributions across multiple harmonic orders, in contrast to the hard low-order cutoff used in practical equivariant settings (Figure 1b,c).

Continuous formulation. We begin with a continuous object, because the representation claim of EPOCH is about what geometric content is exposed, not about one particular cached tensor. For a unit direction $\mathbf{u} \in \mathbb{S}^2$, define the shellwise projection response

$$F_s(\mathbf{u}) = \sum_{i=1}^{n} w_i \, g_s(\|\tilde{\mathbf{r}}_i\|) \, k_\tau\left(\frac{\tilde{\mathbf{r}}_i}{\|\tilde{\mathbf{r}}_i\|}, \mathbf{u}\right), \tag{1}$$

where $\tilde{\mathbf{r}}_i$ denotes centered and canonically aligned coordinates, g_s is a shell-specific radial window, w_i is an atom weight, and k_τ is an angular kernel. In our implementation, g_s is a Gaussian shell and

$$k_\tau(\mathbf{v}, \mathbf{u}) = \exp\big(\tau(\mathbf{v}^\top \mathbf{u} - 1)\big), \tag{2}$$

which is a zonal kernel depending only on the angle between $\mathbf{v}$ and $\mathbf{u}$.

Each term maps directly to the directional-probing picture above: $\mathbf{u}$ is a probe direction, g_s selects a shell, k_τ softens the directional cone, and w_i specifies which fixed atom-dependent signal is being aggregated. In the QM9 and MD17 instantiations, w_i is implemented through a small family of fixed chemistry-aware weighting channels summarized in Appendix A.1. The shell windows are centered at a small number of radial scales. The kernel width τ controls angular concentration, with larger values producing sharper directional responses (Figure 2a).

Eq. 1 therefore defines a shellwise spherical response field. Because k_τ is zonal, $F_s(\mathbf{u})$ admits a spherical-harmonic decomposition, so the descriptor should be read as exposing angular structure across multiple orders rather than as selecting a fixed list of low-order moments.

The role of this continuous view is interpretive: Eq. 1 specifies what the practical descriptor samples, not a closed-form harmonic architecture. The implemented descriptor evaluates the field at finitely many directions and shells, so it preserves only a limited range of angular detail jointly determined by the probe grid and kernel width. In the implemented setting, that range covers a practical mid-order band beyond the very low-order regime emphasized by practical equivariant settings (Figure 2c), while avoiding the steep tensor-product cost of substantially higher equivariant orders. Appendix A.3 sketches this interpretation and its empirical boundary.

This construction is related in spirit to classical kernel-based geometric descriptors such as SOAP and MBTR (Bartók et al., 2013; Huo & Rupp, 2022); Section 2 discusses the representational distinction.

Relation to equivariance. Equivariant backbones and explicit projection descriptors address different parts of the geometry problem. Equivariance constrains how features transform under rigid motions; our projection field determines which geometric responses are made explicit to the model. The two are therefore related but not interchangeable. For non-equivariant models, the descriptor acts as an explicit geometric supplement. For equivariant models such as NequIP, it becomes a controlled external signal for testing representational overlap. If higher native $l_{\max}$ already captured the relevant angular content, descriptor-side gains should largely disappear. Our NequIP study is designed precisely to test this prediction.

3.3 Discrete Descriptor Construction

The practical descriptor used in experiments is the finite-sampling counterpart of Eq. 1. We first center each molecule by its coordinate mean and align it to a canonical frame whose axes are the principal axes of a weighted covariance matrix. Axis signs are then fixed by low-order odd moments, implemented here through cubic-moment sign disambiguation, to remove the remaining reflection ambiguity. We then evaluate the response on a finite directional set $\mathcal{U} = \{\mathbf{u}_j\}_{j=1}^{K} \subset \mathbb{S}^2$ and a finite collection of shell windows $\{g_s\}_{s=1}^{S}$. In the descriptor family used here, $\mathcal{U}$ is a fixed polar–azimuth sampling grid with $K = k^2$ shared directions. This yields the shellwise sampled responses

$$R_{s,j} = \sum_{i=1}^{n} w_i \, g_s(\|\tilde{\mathbf{r}}_i\|) \, \exp\left(\tau\left(\frac{\tilde{\mathbf{r}}_i^\top \mathbf{u}_j}{\|\tilde{\mathbf{r}}_i\|} - 1\right)\right). \tag{3}$$

In words, Eq. 3 says that each sampled entry records how strongly atoms in the molecule respond to one shared direction on one shell under one weighting channel. The implemented descriptor surface then expands this sampled field into a small multi-channel tensor by applying several atom-dependent weighting rules and companion summary statistics to the same directional grid. The default QM9/MD17 instantiation uses four fixed directional labels, seven dynamic shellwise response channels, and nine global summary features (Figure 2b); concrete channel definitions are deferred to Appendix A.1.

The implemented EPOCH descriptor should therefore be read as a finite sampling of that field. In the descriptor family used throughout this paper, the probe set is fixed once and reused across molecules, while shell centers are scaled by the molecule-level radial extent. The default surface used for the main experiments employs a fixed 24×24 angular grid (K=576 directional samples), S=3 shell scales, a normalized descriptor radius of 10.0, and angular sharpness τ=12.0. The resulting representation is finite and lightweight, but its role is not to emulate one predeclared low-order harmonic truncation. Instead, it retains an explicit directional tensor through descriptor construction and a practical angular band that can complement the native representations of heterogeneous 3D backbones. Backbone-specific hidden sizes, gating constants, and variant-specific knobs are deferred to the appendix.

Canonical alignment is used only to stabilize the finite directional sampling: the continuous operator in Eq. 1 is frame-independent, but canonicalization makes sampled responses comparable across molecules.

3.4 Backbone-Specific Adapters

Different molecular backbones maintain different geometric state spaces, so a single hard-coded fusion rule would be artificial. We instead use a constrained adapter family:

$$\mathbf{h}' = A_b(\mathbf{h}, d(\mathcal{M})), \tag{4}$$

where $\mathbf{h}$ is a native backbone state and A_b is a lightweight adapter chosen for backbone family b. We share the descriptor family and its geometric content, not a single universal insertion rule. Because these backbones expose geometry through different native state spaces, forcing one common fusion interface would confound the representation comparison with avoidable architectural mismatch.

The adapters are therefore deliberately minimal and transformation-compatible. For scalar backbones such as SchNet, the descriptor is fused with scalar graph or node states. For angular interaction models such as DimeNet++, it modulates scalar interaction states that already carry directional context. For coordinate-aware or equivariant models such as EGNN, NequIP, and Allegro, we restrict the descriptor to invariant summaries or compatible scalar side channels rather than overwriting non-scalar equivariant features. In this paper, these implementations fall into three practical patterns: readout residuals, interaction-level affine or gating terms, and invariant side heads.

For the NequIP `summary_residual` instantiation used in the order-complementarity study, the adapter takes a graph-level invariant state $\mathbf{g}$ and a descriptor summary $\mathbf{s}$, forms a hidden descriptor state

$$\mathbf{h}_{\mathrm{desc}} = \phi([\mathbf{g}; \mathbf{s}]), \tag{5}$$

then predicts a bounded scalar correction

$$\hat{y} = \hat{y}_{\mathrm{base}} + \alpha\, \sigma(a(\mathbf{h}_{\mathrm{desc}}) + b)\, r(\mathbf{h}_{\mathrm{desc}}), \tag{6}$$

where $\hat{y}_{\mathrm{base}}$ is the plain backbone prediction, $r(\mathbf{h}_{\mathrm{desc}})$ is a learned scalar residual, σ is a gate, and α is a residual scale. In the reported instantiation, $a(\cdot)$ and $r(\cdot)$ are lightweight scalar MLP heads applied to $\mathbf{h}_{\mathrm{desc}}$. This form captures the design principle shared across the reported adapters: the descriptor supplies a lightweight scalar correction path without replacing the backbone's native geometric state space. Appendix A.1 records the concrete settings used in each backbone family.

4 Experiments

4.1 Experimental Questions and Evaluation Protocol

Our experiments are organized around four questions. **Q1:** Does one shared descriptor family help across heterogeneous 3D backbones, including a strictly local high-order equivariant architecture

Table 1: QM9 test MAE for baseline versus baseline + EPOCH comparisons. E denotes EPOCH. Target units are D for μ, a_0^3 for α, a_0^2 for r^2, cal/mol K for C_v, and meV otherwise. Favorable rows across SchNet, EGNN, DimeNet++, NequIP, and the strictly local Allegro address Q1, while the concentration of larger gains on angular-sensitive electronic targets addresses Q2.

Target	SchNet			EGNN			DimeNet++			NequIP			Allegro		
	base	+E	Δ	base	+E	Δ	base	+E	Δ	base	+E	Δ	base	+E	Δ
μ	0.033	0.026	**-0.007**	0.029	0.023	**-0.006**	0.033	0.031	**-0.002**	0.025	0.022	**-0.003**	0.093	0.090	**-0.004**
α	0.235	0.166	**-0.069**	0.071	0.069	**-0.002**	0.057	0.057	**<0**	0.052	0.045	**-0.008**	0.065	0.060	**-0.005**
HOMO	41.000	30.984	**-10.016**	29.000	28.453	**-0.548**	24.433	21.987	**-2.446**	29.501	27.208	**-2.293**	48.787	48.771	**-0.016**
LUMO	34.000	26.630	**-7.370**	25.000	23.939	**-1.061**	19.422	18.749	**-0.672**	24.427	20.703	**-3.724**	39.909	39.907	**-0.002**
gap	63.000	50.037	**-12.963**	48.000	47.476	**-0.524**	32.364	32.346	**-0.018**	44.824	42.163	**-2.661**	75.615	74.980	**-0.636**
r^2	0.073	–	–	0.106	0.081	**-0.025**	0.371	0.394	**0.023**	0.236	0.257	**0.021**	0.588	0.565	**-0.023**
ZPVE	1.700	1.433	**-0.267**	1.550	1.521	**-0.029**	1.562	1.536	**-0.026**	1.532	1.244	**-0.288**	1.273	1.262	**-0.011**
U_0	14.000	13.677	**-0.323**	11.000	9.259	**-1.741**	10.184	9.926	**-0.258**	15.313	5.969	**-9.344**	8.853	6.996	**-1.857**
U	19.000	14.353	**-4.647**	12.000	10.302	**-1.698**	10.505	10.157	**-0.348**	7.677	7.328	**-0.349**	8.888	7.052	**-1.836**
H	14.000	10.258	**-3.742**	12.000	9.822	**-2.178**	10.680	10.389	**-0.291**	7.619	7.345	**-0.274**	8.943	6.768	**-2.175**
G	14.000	10.913	**-3.087**	12.000	9.676	**-2.324**	7.875	7.482	**-0.393**	9.018	7.926	**-1.092**	9.473	7.853	**-1.620**
C_v	0.0330	0.0296	**-0.0034**	0.0310	0.0317	**0.0007**	0.0280	0.0272	**-0.0008**	0.0267	0.0247	**-0.0020**	0.0251	0.0250	**-0.0001**

Note. E denotes EPOCH and $\Delta = \mathrm{MAE}_{\text{Backbone+EPOCH}} - \mathrm{MAE}_{\text{Backbone}}$; negative values indicate improvement. Values are rounded to three decimal places, except the C_v row, which retains four to resolve small changes. Allegro energy-like values are converted to meV. SchNet and EGNN entries are literature reference values; DimeNet++, NequIP, and Allegro entries are reproduced paired comparisons. SchNet r^2 is not paired because its legacy headline and shared-protocol rerun are not directly comparable (Appendix A.6).

without message passing, when fused through backbone-aware adapters? **Q2:** Are gains selective to angular-sensitive targets, and do matched controls support a directional mechanism? **Q3:** Does descriptor-side complementarity persist as native equivariant order increases? **Q4:** Can lower-order sampling-augmented models match higher-order baselines at substantially lower parameter and compute cost?

These questions test whether a shared descriptor exposes geometric content beyond practical equivariant orders. Sections 4.2–4.5 address them through cross-backbone QM9 comparisons, mechanism controls, NequIP order ablations, and cross-dataset force prediction, respectively. Unless otherwise noted, headline values are four-seed averages.

We use QM9 as the controlled main testbed and report mean absolute error in the native unit of each target. Within each reproduced implementation family, the baseline and EPOCH variant share the split, training protocol, and evaluation path. SchNet and EGNN cells in Table 1 are literature reference values, whereas DimeNet++, NequIP, and Allegro report paired reproductions. Appendix A.1 records the concrete implementation settings.

We implemented EPOCH on five representative 3D backbone families: SchNet (Schütt et al., 2017), DimeNet++ (Gasteiger et al., 2020a), EGNN (Satorras et al., 2021), NequIP (Batzner et al., 2022), and Allegro (Musaelian et al., 2023). They span scalar distance-based encoders, angular interaction models, message-passing equivariant architectures, and a strictly local high-order equivariant architecture without message passing. For each backbone, the fusion rule follows the adapter principle in Section 3: scalar backbones use readout-side fusion, angular backbones use interaction-aware fusion, and equivariant backbones use transformation-compatible side channels. The descriptor's geometric content is shared; the fusion rule is backbone-specific. We therefore report each comparison as *Backbone*, *Backbone + EPOCH*, and $\Delta = \mathrm{MAE}_{\text{Backbone+EPOCH}} - \mathrm{MAE}_{\text{Backbone}}$, where negative values indicate improvement.

4.2 Cross-Backbone Evidence

Table 1 is the main QM9 test of whether one descriptor family helps across heterogeneous backbones rather than only within one architecture class.

The table supports the cross-backbone claim: favorable rows appear in SchNet, EGNN, DimeNet++, NequIP, and Allegro. In particular, Allegro is strictly local and does not use message passing (Musaelian et al., 2023); its favorable matched comparisons show that EPOCH remains useful with high local equivariant angular capacity. The descriptor therefore contributes across scalar,

Table 2: Parameter-matched mechanism controls on QM9 HOMO with SchNet. Removing directional organization progressively increases the test MAE.

Variant	HOMO MAE (meV)
Full EPOCH	32.43
Directional mean	34.77
Shuffled response	35.37
Distance-only	35.97
Global-only	41.94

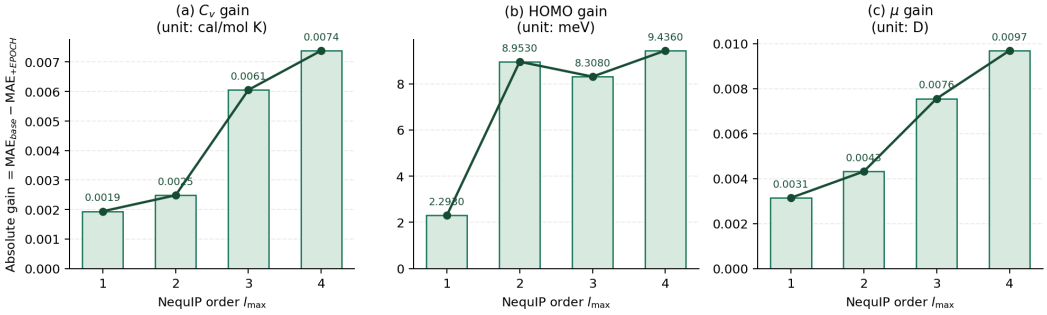

Figure 3: NequIP order analysis on three representative QM9 targets, anchored at the same $l_{\max} = 1$ rows used in Table 1. Each panel reports absolute paired gain $= \mathrm{MAE}_{base} - \mathrm{MAE}_{+EPOCH}$ in that target's native unit, so larger positive values indicate stronger improvement from EPOCH. Across C_v, HOMO, and μ, the descriptor-side gain remains positive as native equivariant order increases.

angular, coordinate-equivariant, tensor-equivariant, and strictly local equivariant families when fused through backbone-aware adapters.

The gains are selective rather than uniform: the clearest improvements appear on electronically sensitive targets such as U_0, LUMO, gap, and polarizability, while already stable quantities often move little. This target concentration is consistent with a geometric mechanism rather than generic capacity, and agrees with the separate SchNet–SOAP comparison in Appendix A.5 (Table 5).

The small and inconsistent r^2 changes mark a boundary: EPOCH is most useful when explicit directional geometry is the missing signal (Appendix A.6). We next isolate the contribution of directional organization before testing whether this complementarity persists as native equivariant order rises.

4.3 Mechanism Isolation

To separate directional organization from generic descriptor capacity, we evaluated parameter-matched controls on QM9 HOMO with SchNet. All variants use the same split, backbone, descriptor dimensionality, adapter, optimizer, seed, and training budget; the controls progressively remove cross-direction variation, response-to-direction correspondence, the directional dimension, and dynamic directional channels while retaining the fusion path.

In this representative HOMO diagnostic, progressively removing directional organization increases the test error, while the global-only variant performs worst. The result indicates a contribution from directional structure beyond the adapter pathway.

We next test whether this complementarity persists as native equivariant order increases.

4.4 NequIP as an Order-Complementarity Probe

The NequIP study tests whether explicit sampling remains useful as native order increases (Figure 3) and compares its cost with raising native order (Table 3). Varying $l_{\max}$ changes native angular

Table 3: Cross-order NequIP efficiency comparisons computed from the same corrected-family MAE values used in Figure 3. Each row compares a lower-order EPOCH-augmented candidate against the next-order plain baseline, directly addressing Q4. Here MAE_c denotes the EPOCH-augmented lower-order candidate and MAE_b the plain higher-order baseline. The final column is explicitly marginal: descriptor-side incremental FLOPs versus the incremental FLOPs required to raise the backbone from l to $l{+}1$.

Target	Comparison candidate	Baseline ($l{+}1$)	MAE_c	MAE_b	$\Delta\downarrow$	Params	Incremental FLOPs
μ	NequIP $l{=}1$ + EPOCH	NequIP $l{=}2$	0.0223	0.0229	-0.0006	**5.22M** / 10.28M	**0.302M** / +125.23M (414×)
μ	NequIP $l{=}2$ + EPOCH	NequIP $l{=}3$	0.0186	0.0285	-0.0099	**10.28M** / 17.22M	**0.302M** / +303.97M (1007×)
HOMO	NequIP $l{=}1$ + EPOCH	NequIP $l{=}2$	0.0272	0.0359	-0.0087	**5.22M** / 10.28M	**0.302M** / +125.23M (414×)
HOMO	NequIP $l{=}2$ + EPOCH	NequIP $l{=}3$	0.0269	0.0328	-0.0059	**10.28M** / 17.22M	**0.302M** / +303.97M (1007×)
C_v	NequIP $l{=}1$ + EPOCH	NequIP $l{=}2$	0.0247	0.0252	-0.0005	**5.22M** / 10.28M	**0.039M** / +125.23M (3211×)
C_v	NequIP $l{=}2$ + EPOCH	NequIP $l{=}3$	0.0227	0.0270	-0.0043	**10.28M** / 17.22M	**0.039M** / +303.97M (7794×)

capacity while keeping the descriptor fixed; all comparisons use the same corrected family, split, training protocol, and evaluation path.

In this family, EPOCH enters after equivariant feature aggregation as a descriptor-conditioned scalar correction. The order study therefore tests complementarity at the scalar prediction interface, rather than injecting EPOCH into the native equivariant channels.

Figure 3 answers Q3 directly: across the thermodynamic, electronic, and dipolar targets C_v, HOMO, and μ, EPOCH improves the same $l{=}1$ anchor used in Table 1, and gains remain positive through $l{=}4$.

Table 3 shows that a lower-order model with EPOCH matches or approaches the next-order baseline at a fraction of the cost, with a $400\times$–$7800\times$ gap in incremental FLOPs. The margin is tighter for μ and clearer for HOMO and C_v.

4.5 Cross-Dataset Experiments on Force Prediction

To test transfer beyond equilibrium scalar regression, we evaluate SchNet on the original MD17 force-prediction task. Forces are computed as negative gradients of the predicted energy through the coordinate-dependent descriptor, canonicalization, shell responses, and adapter. Unless otherwise noted, $k{=}48$; the sampling-density ablation is in Appendix A.7.

Four-seed-average force MAE improves for all ten MD17 molecules (Appendix A.8), showing transfer from equilibrium property prediction to force-level modeling under geometric perturbation.

The MD17 results provide the primary transfer test; an MD22 host–guest experiment is reported as supplementary evidence in Appendix A.9.

5 Conclusion

We presented EPOCH, a shared geometric side channel for 3D molecular learning. Its shellwise directional responses provide a finite-resolution, harmonic-interpretable summary of multi-scale geometry without requiring higher-order tensor products. Across QM9, EPOCH complements scalar, angular, coordinate-equivariant, tensor-equivariant, and strictly local equivariant backbones, including Allegro, with the clearest gains on angular-sensitive electronic targets. A parameter-matched HOMO diagnostic supports a contribution from directional organization beyond the adapter pathway, while the target-dependent results clarify the scope of this effect. On MD17, EPOCH also improves SchNet force prediction.

NequIP cross-order comparisons show that a lower-order backbone with EPOCH can match or outperform the next-order baseline at $400\times$–$7800\times$ lower incremental FLOP cost (Table 3). Together, these results indicate that explicit geometric side information complements, rather than replaces, native equivariant backbone design.

AI USE STATEMENT

Language-model tools supported literature search and source organization, early-stage idea exploration, experimental workflow planning, implementation-script drafting, experiment-log organization, result summarization, and manuscript editing for grammar, presentation, and clarity. They were not used as authors or as components of the EPOCH model, training pipeline, or scientific decision procedure.

The authors independently checked all AI-assisted material and verified the experimental configurations, code paths, data provenance, numerical results, citations, and scientific interpretations against the underlying source archives, scripts, and execution records. The authors take full responsibility for the final manuscript, reported results, and all scientific claims.

REPRODUCIBILITY STATEMENT

The descriptor construction and adapter principle are specified in Section 3, the evaluation protocol and matched-comparison policy in Section 4, and the backbone-specific settings, additional controls, and ablations in Appendix A. The anonymous code repository submitted with this paper, available at `https://anonymous.4open.science/api/repo/epoch-6188`, and the supplementary attachments contain the EPOCH descriptor, training and evaluation scripts, and instructions for reproducing the principal QM9 and MD17 experiments. The experiments use the public QM9 and MD17 benchmarks; the reported results are based on the fixed splits, seeds, and configurations documented in the paper and appendix.

REFERENCES

Junyi An, Xinyu Lu, Chao Qu, Yunfei Shi, Peijia Lin, Qianwei Tang, Licheng Xu, Fenglei Cao, and Yuan Qi. Equivariant spherical transformer for efficient molecular modeling. *arXiv preprint arXiv:2505.23086*, 2025.

Kendall Atkinson and Weimin Han. *Spherical Harmonics and Approximations on the Unit Sphere: An Introduction*, volume 2044 of *Lecture Notes in Mathematics*. Springer, 2012. doi: 10.1007/978-3-642-25983-8.

Albert P Bartók, Risi Kondor, and Gábor Csányi. On representing chemical environments. *Physical Review B—Condensed Matter and Materials Physics*, 87(18):184115, 2013.

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*, 2022.

Ilyes Batatia, Simon Batzner, Dávid Péter Kovács, Albert Musaelian, Gregor N. C. Simm, Ralf Drautz, Christoph Ortner, Boris Kozinsky, and Gábor Csányi. The design space of E(3)-equivariant atom-centred interatomic potentials. *Nature Machine Intelligence*, 7:56–67, 2025.

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 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, NeurIPS*, 2020a.

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

Johannes Gasteiger, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional graph neural networks for molecules. *Advances in neural information processing systems*, 34:6790–6802, 2021.

Haoyan Huo and Matthias Rupp. Unified representation of molecules and crystals for machine learning. *Machine Learning: Science and Technology*, 3(4):045017, 2022.

Yi-Lun Liao and Tess Smidt. Equiformer: Equivariant graph attention transformer for 3d atomistic graphs. *arXiv preprint arXiv:2206.11990*, 2022.

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.

S Liu, H Wang, W Liu, J Lasenby, H Guo, and J Tang. Pre-training molecular graph representation with 3d geometry. arxiv 2021. *arXiv preprint arXiv:2110.07728*, 10, 2021.

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.

Albert Musaelian, Simon Batzner, Anders Johansson, Lixin Sun, Cameron J. Owen, Mordechai Kornbluth, and Boris Kozinsky. Learning local equivariant representations for large-scale atomistic dynamics. *Nature Communications*, 14(1):579, 2023.

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

Kristof Schütt, Pieter-Jan Kindermans, Huziel Enoc Sauceda Felix, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. *Advances in neural information processing systems*, 30, 2017.

Kristof Schütt, Oliver Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial properties and molecular spectra. In *International conference on machine learning*, pp. 9377–9388. PMLR, 2021.

Rahul Sheshanarayana and Fengqi You. Molecular representation learning: Cross-domain foundations and future frontiers. *Digital Discovery*, 4:2298–2335, 2025. doi: 10.1039/D5DD00170F.

Limei Wang, Yi Liu, Yuchao Lin, Haoran Liu, and Shuiwang Ji. Comenet: Towards complete and efficient message passing for 3d molecular graphs. In *Advances in Neural Information Processing Systems*, 2022.

Gengmo Zhou, Zhifeng Gao, Qiankun Ding, Hang Zheng, Hongteng Xu, Zhewei Wei, Linfeng Zhang, and Guolin Ke. Uni-mol: A universal 3d molecular representation learning framework. In *International Conference on Learning Representations*, 2023.

A Experimental Implementation Details

Because the backbone architectures and codebases differ, EPOCH is attached through backbone-specific adapters rather than through one shared insertion point. This section summarizes the main training and fusion settings used for the reported backbone families. Across all backbone families, the descriptor branch is trained from scratch together with the backbone; learning-rate warmup may be used, but warm-start initialization is not part of the reported setup.

Unless otherwise noted, the reported headline values are four-seed averages.

A.1 Shared Descriptor Definition

Unless otherwise noted, the QM9 experiments use the cached `chemist_shells_v1` descriptor family. The descriptor is precomputed from mean-centered molecular coordinates after principal-axis canonicalization with a cubic-moment sign disambiguation step. The surface uses a fixed polar–azimuth 24×24 directional grid, corresponding to $K{=}576$ shared directional probes, together with $S{=}3$ shell scales $[0.6, 1.0, 1.4]$, shell-width scale 0.35, descriptor radius 10.0, angular sharpness $\tau{=}12.0$, and an outer electrostatic channel evaluated at radius multiplier 1.8. Per direction, the cached surface contains four fixed directional labels, seven dynamic shellwise response channels, and nine global summary features.

Before turning to backbone-specific settings, Table 4 summarizes the conceptual position of EPOCH relative to representative descriptor-side, interaction-side, and equivariant approaches. Its early placement is intentional: this distinction is used repeatedly in the main text, so we state it here before the appendix shifts to backbone-specific implementation details. The point of the table is not to rank these methods universally, but to make explicit the design axis used throughout the paper: EPOCH is a portable descriptor-side signal, whereas methods such as DimeNet++ and EST place geometric processing inside the backbone interaction rule.

Table 4: Conceptual positioning of EPOCH relative to representative geometric, interaction-side, and equivariant approaches. "Portable" means the geometric signal can be reused across heterogeneous backbone families without redesigning the backbone's core interaction rule.

Method	Explicit direction	Multi-shell	Learnable	Sampling location	Portable
SOAP	No	Yes	No	Descriptor-side	Yes
MBTR	No	Yes	No	Descriptor-side	Yes
DimeNet++	Partial	No	Yes	Interaction-side	No
NequIP	Implicit	No	Yes	Backbone-internal	No
Allegro	Implicit	No	Yes	Backbone-internal	No
EST (An et al., 2025)	Yes	No	Yes	Interaction-side	No
EPOCH	Yes	Yes	No	Descriptor-side	Yes

Two contrasts matter for the discussion in Section 2. Relative to classical descriptors such as SOAP and MBTR, EPOCH keeps explicit directional structure rather than compressing geometry directly into invariant summaries. Relative to EST-style interaction-side spherical sampling, EPOCH moves the geometric construction outside the equivariant interaction block, which is why the same descriptor family can be reused across SchNet, DimeNet++, EGNN, NequIP, and Allegro instead of being tied to one backbone family.

A.2 Backbone-Specific Instantiations

SchNet. Backbone setting: $n_{\text{atom}}{=}128$, $n_{\text{interactions}}{=}6$, cutoff 5.0, batch size 100, learning rate 5×10^{-4}. EPOCH setting: a descriptor-conditioned residual correction at scalar readout built on this same implementation family.

EGNN. Backbone setting: 7 message-passing layers and radius cutoff 5.0. EPOCH setting: a lightweight scalar adapter on the EGNN message/readout pathway within the same implementation family.

DimeNet++. Backbone setting: hidden channels 128, 4 interaction blocks, interaction embedding size 64, basis embedding size 8, output embedding channels 256, 7 spherical bases, 6 radial bases,

cutoff 5.0, at most 32 neighbors, batch size 32, learning rate 10^{-3}, and 3000 warmup steps. EPOCH setting: a lightweight residual head on the scalar prediction built on the same implementation family.

NequIP. All reported NequIP order comparisons are drawn from the same implementation family. Backbone setting: $l_{\max} = 1$, 5 interaction layers, 128 feature channels, 8 radial bases, and radius cutoff 5.0. EPOCH setting: `summary_residual` with descriptor hidden width 128, `desc_gate_bias` $= -2.0$, and `desc_residual_scale` $= 1.0$.

Allegro. Allegro is evaluated as a strictly local, non-message-passing high-order $E(3)$-equivariant backbone (Musaelian et al., 2023). EPOCH is attached through the descriptor data contract of shellwise directional responses and global summary features. For each QM9 target, the descriptor-augmented run is compared with its corresponding reproduced no-descriptor Allegro baseline; all twelve target-matched comparisons favor Allegro+EPOCH. The main table records the direction of these matched comparisons as Δ_{m}.

A.3 Proof Sketch for the Harmonic Interpretation

We provide an interpretive derivation for the three main-text statements in Section 3: **(i)** the underlying continuous response field admits a harmonic decomposition, **(ii)** the implemented descriptor is obtained by sampling that field rather than by analytically truncating it to a fixed $l_{\max}$, and **(iii)** finite directional sampling preserves only an implementation-dependent portion of the available angular content.

The identities used here are standard facts from harmonic analysis on the sphere, including zonal-kernel Legendre expansions and the spherical-harmonic addition theorem (Atkinson & Han, 2012). We include the derivation to clarify the continuous parent object behind EPOCH.

Start from the continuous shellwise response

$$F_s(\mathbf{u}) = \sum_{i=1}^{n} w_i g_s(\|\tilde{\mathbf{r}}_i\|) k_\tau\left(\frac{\tilde{\mathbf{r}}_i}{\|\tilde{\mathbf{r}}_i\|}, \mathbf{u}\right).$$

Because k_τ is zonal, it depends only on the inner product $t = \mathbf{v}^\top \mathbf{u}$, so it admits a Legendre expansion

$$k_\tau(t) = \sum_{\ell=0}^{\infty} a_\ell P_\ell(t).$$

Substituting this into F_s yields

$$F_s(\mathbf{u}) = \sum_{i=1}^{n} w_i g_s(\|\tilde{\mathbf{r}}_i\|) \sum_{\ell=0}^{\infty} a_\ell P_\ell\left(\frac{\tilde{\mathbf{r}}_i}{\|\tilde{\mathbf{r}}_i\|}^\top \mathbf{u}\right).$$

Applying the spherical harmonic addition theorem,

$$P_\ell(\hat{\mathbf{r}}_i^\top \mathbf{u}) = \frac{4\pi}{2\ell+1} \sum_{m=-\ell}^{\ell} Y_{\ell m}(\hat{\mathbf{r}}_i) \overline{Y_{\ell m}(\mathbf{u})},$$

we obtain

$$F_s(\mathbf{u}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \left[\frac{4\pi a_\ell}{2\ell+1} \sum_{i=1}^{n} w_i g_s(\|\tilde{\mathbf{r}}_i\|) Y_{\ell m}(\hat{\mathbf{r}}_i)\right] \overline{Y_{\ell m}(\mathbf{u})}.$$

This is the coefficient form underlying the main-text harmonic interpretation, up to the usual real/complex spherical-harmonic convention.

The coefficient form above gives the main-text interpretation directly: because k_τ has nonzero Legendre weights at infinitely many ℓ, the continuous response field F_s is a multi-order harmonic object rather than one restricted to a predeclared finite $l_{\max}$. The implemented descriptor $\{R_{s,j}\}_{s,j}$ is obtained by evaluating this same field at a finite directional set $\{u_j\}_{j=1}^{K}$, so it is a *sampling* of a multi-order harmonic field, not a closed-form truncation of one.

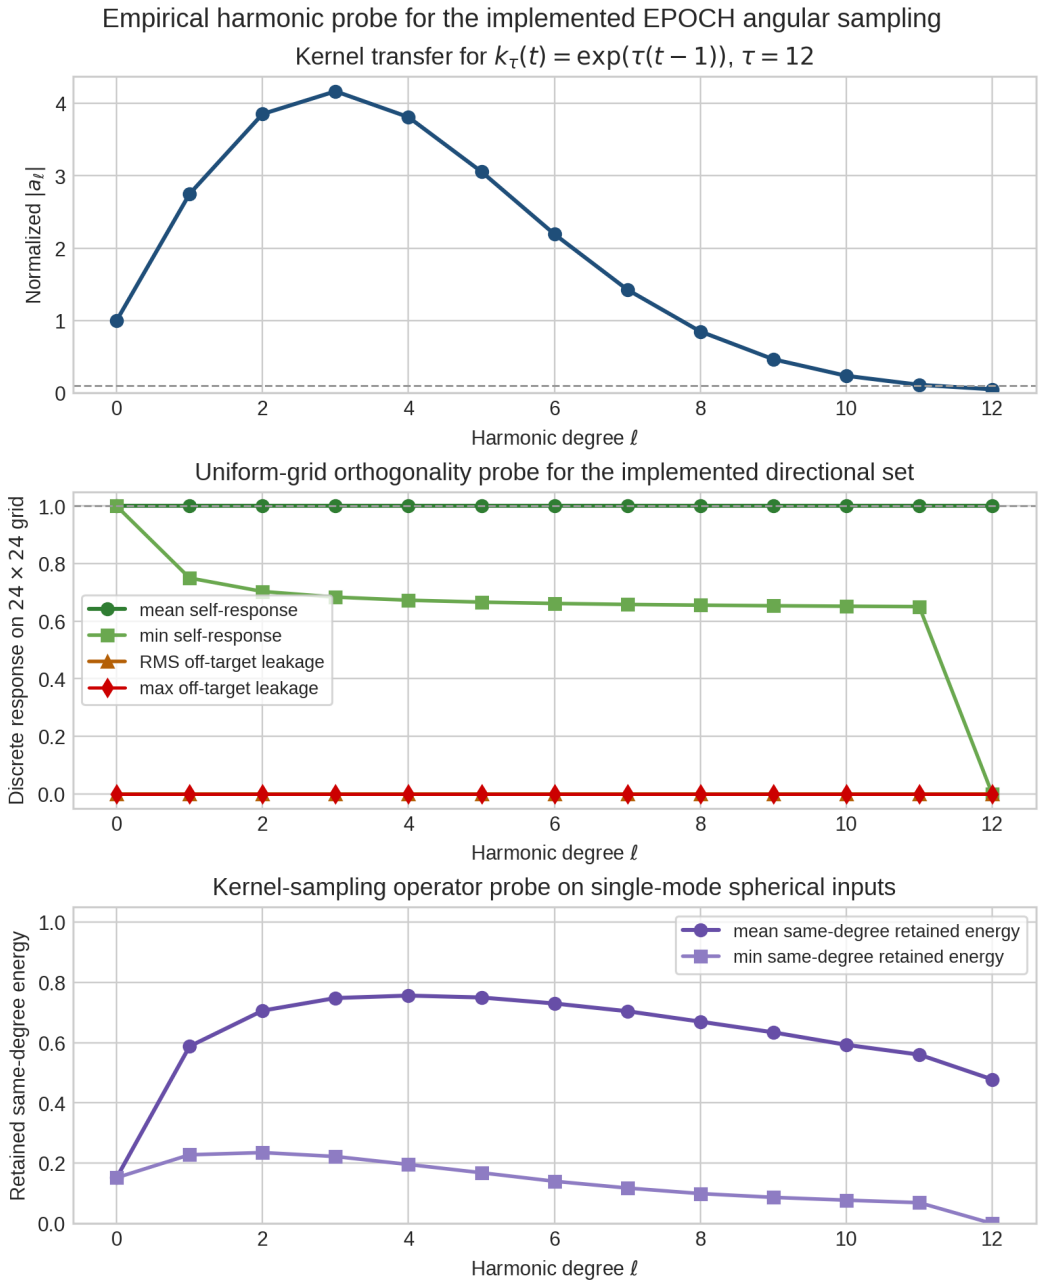

Figure 4: Top: normalized Legendre weights of the τ=12 zonal kernel. Middle: discrete orthogonality of real spherical harmonics on the implemented 24×24 directional grid. Bottom: same-degree retained energy when individual spherical-harmonic modes are passed through the implemented kernel-sampling operator. The practical picture is finite-resolution and gradually decaying, not exact preservation of arbitrarily high orders.

Two points bound this interpretation. First, the continuous field has nonzero Legendre weights at infinitely many ℓ, but a given molecule's realized coefficients still depend on its shell-weighted atomic distribution and may vanish by symmetry. Second, finite-K sampling preserves only an empirical, implementation-dependent angular range, which we characterize directly in Figure 4 rather than via an asymptotic recovery argument.

Table 5: SchNet C_v comparison against a SOAP descriptor baseline on QM9.

Model	C_v MAE (cal/mol K)
SchNet	0.0330
SchNet + SOAP	0.031736
SchNet + EPOCH	**0.0296**

Empirical harmonic probe. To ground this boundary in the implemented setting, we ran a small operator-level probe on the exact 24×24 directional grid and τ=12.0 used in the main descriptor surface. The probe is reported in Figure 4. Two observations matter. First, the fixed directional grid itself remains numerically well behaved through the tested regime: the discrete real-spherical-harmonic Gram matrix stays nearly block-orthogonal up to the plotted degrees, so the grid is not the immediate bottleneck at low order. Second, the kernel-sampling operator shows gradual rather than abrupt decay across degrees: the normalized Legendre weights remain substantial through the mid-order regime and the same-degree retained energy stays strong through roughly the range most relevant to our paired equivariant experiments before tapering at higher degree. Because the probe is evaluated only up to the plotted range, it should be read as evidence against an abrupt low-order cutoff in the implemented descriptor. In the implemented setting, the retained angular content is concentrated in a mid-order band roughly spanning $\ell \approx 2$ to 8, which lies largely outside the native coverage of practical NequIP settings at $l_{\max} \leq 4$ (Figure 2c). This is the empirical basis for the practical-band interpretation used in the main text.

A.4 Compute Summary

The reproduced experiments were run on internal GPU servers and cluster nodes. All NequIP experiments reported in this paper used NVIDIA A100 80 GB GPUs. The remaining backbone families, including SchNet, EGNN, DimeNet++, Allegro, and the MD17 SchNet study, were run on NVIDIA L40 40 GB GPUs. The incremental FLOP discussion in the NequIP order analysis should therefore be read relative to the A100-based NequIP setting reported in the main text.

A.5 Descriptor Baseline Comparison on QM9 C_v

To situate the reported SchNet C_v gain against a standard hand-crafted geometric descriptor, we ran a small descriptor-side comparison on the same SchNet C_v setting used in the main text. Table 5 shows that EPOCH yields a larger reduction in MAE than adding SOAP in this matched setting.

A.6 QM9 r^2 Reporting Note for SchNet

The SchNet r^2 row in Table 1 is reported without a paired Δ because the legacy headline of $0.073\,a_0^2$ (Schütt et al., 2017) differs substantially from our shared-protocol reproduction. Reported QM9 values for this target span a broad range, and prior work notes that its accuracy is sensitive to the loss and training protocol (Satorras et al., 2021; Schütt et al., 2021). We therefore retain the legacy value for reference but avoid mixing it with our reproduced baseline.

Under the shared protocol used for the other targets, SchNet reaches r^2 MAE $= 0.3529$ and SchNet+EPOCH reaches 0.3200 ($\Delta = -0.0329$). This matched comparison is reported here, while the main table leaves the legacy cell unpaired.

A.7 MD17 Sampling Resolution Ablation

We also tested how directional-grid resolution affects the MD17 SchNet setting used in the main text. Table 6 reports an aspirin-only ablation with $k \in \{12, 24, 48, 96, 192\}$, where k denotes the per-axis directional-grid resolution rather than the total number of probes. Performance improves with denser directional coverage and begins to saturate: force MAE continues to improve modestly beyond k=48, while energy MAE is already near saturation by k=96. These results motivate k=48 as a practical default for the main experiments.

Table 6: Directional sampling resolution ablation on MD17 aspirin with SchNet+EPOCH. Larger k means denser directional sampling.

k	Best val. loss	Test energy MAE	Test forces MAE
12	0.07401	0.14122	0.20078
24	0.07453	0.13969	0.19908
48	0.07000	0.13719	0.19342
96	0.06997	0.13691	0.19093
192	0.06651	0.13766	0.18777

Table 7: Cross-dataset, cross-task transfer evidence on MD17 with SchNet, reported as four-seed-average test force MAE.

Molecule	SchNet	SchNet + EPOCH	$\Delta \downarrow$
`aspirin`	0.2802	**0.1934**	**-0.0868**
`ethanol`	0.0578	**0.0547**	**-0.0031**
`uracil`	0.1509	**0.1082**	**-0.0428**
`azobenzene`	0.2781	**0.1783**	**-0.0998**
`paracetamol`	0.2207	**0.1407**	**-0.0800**
`salicylic_acid`	0.2199	**0.1194**	**-0.1005**
`toluene`	0.1565	**0.0863**	**-0.0702**
`naphthalene`	0.2026	**0.0909**	**-0.1117**
`malonaldehyde`	0.0997	**0.0991**	**-0.0007**
`benzene`	0.2092	**0.1997**	**-0.0095**

A.8 MD17 Transfer Results

Table 7 gives the four-seed-average force MAE values underlying the transfer statement in the main text. SchNet+EPOCH improves over SchNet on all ten molecules.

A.9 Supplementary MD22 Transfer Test

To probe transfer in a chemically and structurally larger regime, we evaluated the 148-atom MD22 `buckyball-catcher` host–guest system. From the benchmark's 600-configuration training budget, we fixed a seed-42 split of 540 training and 60 validation structures; the remaining 5,502 configurations formed the test set. SchNet and SchNet+EPOCH used the same three interaction blocks, 128 hidden channels, optimizer, energy–force loss weights, training budget, and validation-loss checkpoint rule. The EPOCH run used the fixed MD17 descriptor setting with a 48×48 directional grid. This experiment provides supplementary transfer evidence.

Table 8: Supplementary MD22 `buckyball-catcher` energy–force test with SchNet (540/60/5,502 train/validation/test split). Energy MAE is in kcal/mol and force MAE in kcal/mol/Å; the checkpoint is selected by minimum validation loss.

Model	Energy MAE ↓	Δ	Force MAE ↓	Δ
SchNet	1.3629	–	0.8427	–
SchNet + EPOCH	**1.1716**	**-0.1913**	**0.8306**	**-0.0121**