CIRCULAR TOPOLOGICAL TRANSFORMS: GEOMETRIC CHIRALITY *with* COMPLETENESS

Anonymous authors
Paper under double-blind review

ABSTRACT

Geometric graphs such as point clouds, 3D shapes, and molecular conformations require representations that are both (1) $\mathrm{SE}(n)$-*complete*, that representations agree if and only if they differ only by translations and rotations, and (2) *chirality-aware*, that distinct geometries and their mirror images remain distinguishable. Existing approaches do not provide both properties at scale. $\mathrm{SE}(n)$-invariant geometric networks are limited by geometric Weisfeiler–Leman expressivity and capture chirality only locally, while popular topological directional transforms, such as the persistent homology transform (PHT), are injective (can distinguish any shape), but therefore not $\mathrm{SE}(n)$-complete. In this work, we propose *Circular Topological Transforms* (CIRCUiT^2), a principled methodology that turns admissible injective directional transforms such as the PHT to be both $\mathrm{SE}(n)$-complete and chirality-aware on shape classes admitting an equivariant frame. Our construction organizes directional information into ordered circular scans: proper rotations are absorbed by changes of frame, whereas reflections induce a deterministic reversal of the scan. We also show that two naive approaches (group averaging and sorting canonicalization) to symmetrizing directional transforms will fail to detect chirality. Applying our methodology to the PHT, CIRC-PHT achieves perfect chirality-sensitive retrieval on multiple molecular benchmarks and captures long-range chirality beyond the interaction radius of the evaluated equivariant deep neural network models.

1 INTRODUCTION

Geometric graphs represent molecular conformations (Gasteiger et al., 2020; Atz et al., 2021), physical systems (Sanchez-Gonzalez et al., 2020; Satorras et al., 2021), point clouds (Wang et al., 2019), and 3D shapes (Monti et al., 2017; Bronstein et al., 2017) through Euclidean coordinates (Bronstein et al., 2021; Joshi et al., 2023). Their representations should be *complete* (distinguishing shapes not related by translations and proper rotations) and *chirality-aware* (distinguishing mirror images). Chirality is essential in domains such as stereochemistry (enantiomer bioactivity and toxicity; Nguyen et al., 2006; Gaiński et al., 2023; Dumitrescu et al., 2025) and shape analysis (differentiating mirrored parts; Wang et al., 2025; Roell & Rieck, 2024). More broadly, many n-body physical systems exhibit $\mathrm{SE}(3)$ symmetry (Sanchez-Gonzalez et al., 2020; Satorras et al., 2021). We therefore seek $\mathrm{SE}(3)$-invariant representations that preserve completeness and chirality.

Chirality can arise from geometric arrangements of varying spatial extent (Cahn et al., 1966; Eliel & Wilen, 1994). In *central chirality*, handedness depends on the arrangement of a vertex's immediate neighbours, whereas in *axial*, *planar*, and *helical chirality*, it depends on the relative orientation of potentially distant substructures that may themselves be achiral. We term these *local* and *long-range chirality*, respectively. In Figure 1, two achiral terminal groups determine molecular handedness across a bridge, as in allenes and atropisomers. Increasing their separation makes chirality dependent on geometric information beyond local neighbourhoods. Thus, representations must be *reflection-sensitive* and, for long-range chirality, *capture global geometry*. Existing representations satisfy these requirements only partially.

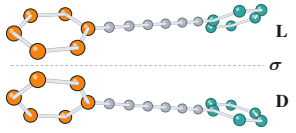

Figure 1: Long-range chirality from two achiral groups.

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.

Naive Invariantization Kills Chirality (Section 3)	
Group averaging collapses the directional height function	**Proposition 3.2**
Canonicalization of $O(n)$-equivariant transforms is $O(n)$-invariant	**Proposition 3.4**
Circular Topological Transforms (Section 4)	
Construction of the circle method $\mathfrak{C}(X, \Phi, \mathrm{Fr}, \mathrm{d})$	**Definition 4.2**
$\mathfrak{C}$ is H-invariant for subgroup $H \subseteq SO(n)$	**Theorem 4.4**
Reflection acts non-trivially on the canonical directional coordinates	**Theorem 4.6**
Equidistributed finite scans with admissible Φ separate distinct rotational orbits	**Theorem 4.11**

Figure 3: Overview of the theoretical results, each linked to its section.

Where current geometric descriptors fall short. Geometric neural networks encode spatial information and symmetry, but face three limitations. *Reflection sensitivity.* Coordinate-free message-passing networks often cannot detect geometric chirality (see Gaiński et al. (2023) for more discussions), while $\mathrm{E}(3)$-invariant models (EGNN, Satorras et al., 2021; GotenNet, Aykent & Xia, 2025) conflate mirror images by construction (Figure 2). *Completeness.* $\mathrm{SE}(3)$-equivariant models (AlphaFold2, Jumper et al., 2021; EquiformerV2, Liao et al., 2024) can distinguish mirror images, but equivariance alone does not guarantee completeness, as illustrated by the geometric WL hierarchy (Joshi et al., 2023). *Global geometry.* Radius-cutoff architectures access distant vertices through successive local interactions, introducing information bottlenecks (Alon & Yahav, 2020). Larger cutoffs incur cubic edge growth, over-squashing (Li et al., 2024), and require prior knowledge of the spatial range (Section 5.2).

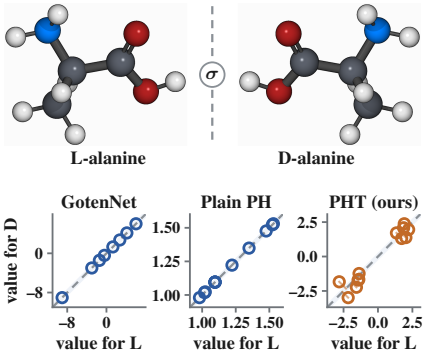

Figure 2: Descriptors of D- versus L-alanine. Only PHT distinguishes them.

Where current topological descriptors fall short. Persistent homology (PH) from topological data analysis (TDA) captures multiscale topological features such as components, loops, and voids (Edelsbrunner & Morozov, 2012). PH-based descriptors have previously been used to augment the expressivity of graph representations (PersLay, Carrière et al., 2020; RePHINE, Immonen et al., 2023; SpectRe, Ji et al., 2025). However, these discard Euclidean geometry. *Directional transforms* instead preserve geometry by computing topological descriptors along viewing directions on the shape. The persistent homology transform (PHT) and Euler characteristic transform (ECT) (Turner et al., 2014) are canonical examples. They are injective on suitable shape classes (Turner et al., 2014; Ghrist et al., 2018; Curry et al., 2022), capture global geometry, and distinguish mirror images (Figure 2). However, as noted in Munch (2025), due to their complete injectivity, PHT and ECT require shape alignments prior to being used in learning pipelines. There have been existing approaches to address this (Curry et al., 2022; Marsh et al., 2024; Paik, 2023). However, the constructions considered there either do not detect reflections or do not establish completeness modulo proper rigid motions. In contrast, our goal is to provide a general framework for constructing $\mathrm{SE}(n)$-complete representations from admissible injective directional transforms on broad classes of shapes. We discuss comparisons with existing literatures in more details in Appendix A.

Contributions. We introduce *circular topological transforms* that are $\mathrm{SE}(n)$-invariant, complete, and chirality-aware. Translation invariance can be achieved by centering the shape (see Section 2 and Appendix D). As this is a standard reduction, our main focus is on rotations and directional information. We summarize the theoretical results in Figure 3 and leave proofs and additional details to the Appendix. The outline of our contributions is as follows:

(1) Limits of naive invariantization. We theoretically characterize two failure modes. Rotation averaging collapses directional height functions to constants (Proposition 3.2), while sorting-based canonicalization of $\mathrm{O}(n)$-equivariant transforms discards handedness (Proposition 3.4).

(2) Chirality-preserving. In response to these limitations, we construct a shape-induced directional transform (Definition 4.2) that removes dependence on proper rotations, while reflections are

not thrown away. We scan the transform along labelled, oriented circles. The ordered scans are invariant to proper rotations and reverse under reflections (Theorems 4.4 and 4.6).

(3) Completeness under finite sampling. For admissible directional transforms (Definition 4.8) and an equivariant frame, we show equidistributed finite scans eventually separate any fixed pair of shapes in distinct rotational orbits (Theorem 4.11). Applying this to PHT, ECT, and the Betti curve transform (BCT) gives training-free descriptors CIRC-PHT, CIRC-ECT, and CIRC-BCT.

(4) Empirical validation. Section 5.2 discusses the experimental results. CIRC-PHT achieves perfect chirality-aware retrieval on 4 molecular datasets using only three directions and a single circle. On a long-range task, it is unaffected by the separation of chiral groups, while an $\mathrm{SE}(3)$-equivariant transformer fails beyond its cutoff radius. In zero-shot stereoisomer separation, it outperforms equivariant and pretrained baselines. Furthermore, every baseline improves when combined with CIRC-PHT, highlighting its value as a complementary geometric descriptor.

2 BACKGROUND AND PRELIMINARIES

We write $\mathcal{S}$ for a class of shapes embedded in $\mathbb{R}^n$, and $X \in \mathcal{S}$ for an individual shape. Examples include finite point clouds and embedded simplicial complexes. A geometric graph is obtained by equipping the vertices of a graph $G = (V, E)$ with coordinates in $\mathbb{R}^n$ and edges between coordinates.

We use $O(n)$ to denote the orthogonal group for orthogonal matrices on $\mathbb{R}^n$ and $SO(n) \subseteq O(n)$ for the subgroup of special orthogonal matrices. A matrix in $\mathrm{SO}(n)$ is a rotation, while a matrix in $\mathrm{O}(n) \setminus \mathrm{SO}(n)$ is orientation-reversing and contributes to chirality. Much of our discussions will extend to a general class of shapes defined as follows.

Definition 2.1. A class of **shapes** is a collection $\mathcal{S}$ of non-empty compact subsets of $\mathbb{R}^n$ such that whenever $X \in \mathcal{S}$, $A \in O(n)$ and $v \in \mathbb{R}^n$, the shape $A \cdot X + v$ is in $\mathcal{S}$.

Suppose $\mathcal{S}$ is the collection of embedded finite graphs (call this $\mathcal{S}^{\mathrm{graph}}$) or embedded finite simplicial complexes (call this $\mathcal{S}^{\Delta}$). A shape X here has vertices $x_1, \ldots, x_N \in \mathbb{R}^n$, and we can center X by subtracting the average vector of all vertices on X. When $\mathcal{S}$ is a class of shapes in general, it is less obvious if we could always do this. In Appendix D, we show that this is always possible.

Thus, unless stated otherwise, we can center a shape $X \in \mathcal{S}$ to the origin in a pre-processing step. This will settle the translation-invariant part of the story, and we focus on building rotation-invariant (invariant to $SO(n)$) and chirality-sensitive (sensitive to $O(n) \setminus SO(n)$) architectures.

Topological Descriptors. The topological descriptors we use here are homology, Betti numbers, and the Euler characteristic. We briefly mention them and refer to Hatcher (2002) for more details.

Definition 2.2. Given any shape $Y \in \mathcal{S}^{\Delta}$ and $k \geq 0$, its **k-th homology** $H_k(Y)$ is a $\mathbb{Z}/2$-vector space with a basis given by, roughly, the collection of k-dimensional holes in Y. The dimension of $H_k(Y)$, denoted $\beta_k(Y)$ is called the **k-th Betti number** of Y. The **Euler characteristic** of Y is the alternating sum $\chi(Y) \coloneqq \sum_{k=0}^{\dim Y} (-1)^k (\#\ k\text{-simplices in } Y)$.

Example 2.3. When Y is a graph, $\beta_0(Y)$ is the number of connected components in Y, $\beta_1(Y)$ is the number of independent cycles in Y, $\beta_k(Y) = 0$ for $k \geq 2$, and $\chi(Y) = |V(Y)| - |E(Y)|$.

A guiding principle from TDA is to examine a shape not at a fixed snapshot but rather how topological features evolve across time. This evolution of shape is described by a filtration.

Definition 2.4. A filtration of X is a family $\{X_t\}_{t \in \mathbb{R}}$ such that $X_s \subseteq X_t \subseteq X$ for all $s \leq t \in \mathbb{R}$.

Example 2.5. Let $\nu \in S^{n-1} = \{x \in \mathbb{R}^n \mid ||x|| = 1\}$ be a unit vector. The **directional filtration** of X in the direction ν is the filtration $X_t^{\nu} \coloneqq \{x \in X \mid \nu \cdot x \leq t\}$ for all $t \in \mathbb{R}$.

Given these filtrations, we can examine the evolution of topological descriptions across time.

Definition 2.6. Given a shape $X \in \mathcal{S}^{\Delta}$ and a filtration $\{X_t\}_{t \in \mathbb{R}}$.

1. The **Betti curve** with respect to $\{X_t\}_{t \in \mathbb{R}}$ is the function $t \mapsto \{\beta_k(X_t)\}_{k=0,\ldots,\dim X}$.
2. The **Euler characteristic curve** with respect to $\{X_t\}_{t \in \mathbb{R}}$ is the function $t \mapsto \chi(X_t)$.
3. The **persistent homology** (PH) with respect to $\{X_t\}_{t \in \mathbb{R}}$ is, roughly, a multi-set $\mathrm{PH}_k(\{X_t\}_{t \in \mathbb{R}})$ for $k = 0, \ldots, \dim X$ that records the appearance (called birth) and disappearance (called death)

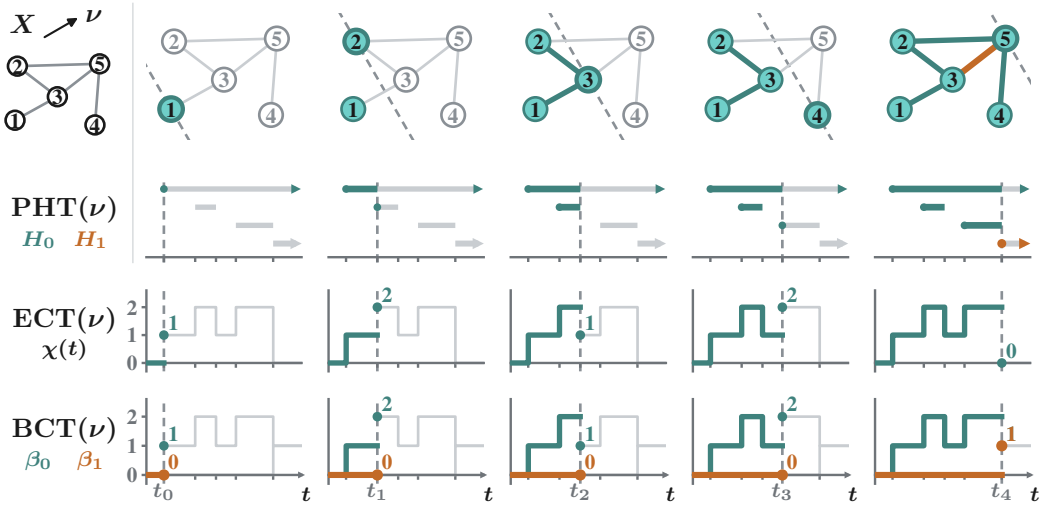

Figure 4: **Directional transforms.** For a fixed direction $\nu \in S^{n-1}$, the function $x \mapsto \nu \cdot x$ sweeps the embedded shape X and induces the sublevel-set filtration $\{x \in X \mid \nu \cdot x \leq t\}_{t\in\mathbb{R}}$. The figure displays the output of the directional transforms PHT, ECT, and BCT on X in the direction ν.

of k-dimensional holes as t increases. An element $(b, d) \in \mathrm{Dgm}_k(\{X_t\}_{t\in\mathbb{R}})$ corresponds to a feature that appears in $H_k(X_b)$ and disappears at $H_k(X_d)$.

We refer to Dey & Wang (2022) and Appendix A of Ballester & Rieck (2025) for more details on PH. For graphs, PH tracks the appearance of permanent cycles and how components merge.

Directional Transforms. We would like to apply Definition 2.6 to directional filtrations arising from all vectors $\nu \in S^{n-1}$, which lead to **directional transforms**.

Definition 2.7 (Directional transform)**.** A **directional transform** Φ is a function that assigns each shape $X \in \mathcal{S}$ a collection of data $\Phi(X)(\nu)$ for each $\nu \in S^{n-1}$. It is of the form

$$\Phi : \mathcal{S} \to \{\text{Functions on } S^{n-1}\}.$$

Example 2.8. Let $\mathcal{S} = \mathcal{S}^{\Delta}$ or $\mathcal{S}^{\text{graph}}$ and $X \in \mathcal{S}$. The following are directional transforms.

1. The **persistent homology transform** (PHT) is the function $\mathrm{PHT}(X)(\nu) := \mathrm{PH}(\{X_t^\nu\}_{t\in\mathbb{R}})$.
2. The **Euler characteristics transform** (ECT) is the function $\mathrm{ECT}(X)(\nu) := \{\chi(X_t^\nu)\}_{t\in\mathbb{R}}$.
3. The **Betti curve transform** (BCT) is the function $\mathrm{BCT}(X)(\nu) :=$ Betti curve of $\{X_t^\nu\}_{t\in\mathbb{R}}$.

See Figure 4 for an example of these transforms.

Definition 2.9. Let H be a subgroup of $O(n)$. We say a directional transform Φ is H-equivariant if for all $h \in H$ and $\nu \in S^{n-1}$, $\Phi(h \cdot X)(\nu) = \Phi(X)(h^{-1} \cdot \nu)$.

Example 2.10. Any directional transform Φ whose value on (X, ν) depends only on the filtration $\{X_t^\nu\}_{t\in\mathbb{R}}$ will be $O(n)$-equivariant. This includes the PHT, ECT, and BCT.

Our goal is to convert such an equivariant directional transform into a design that is invariant to a prescribed subgroup H and, ideally, sensitive to elements outside of H.

Definition 2.11 (H-invariantization)**.** Let H be a subgroup of $O(n)$ and Φ be a directional transform. An H**-invariantization** of Φ is a map $\mathcal{I}_H[\Phi] : \mathcal{S} \to \mathcal{Y}$ for some output space $\mathcal{Y}$, such that

$$\mathcal{I}_H[\Phi](h \cdot X) = \mathcal{I}_H[\Phi](X) \qquad \text{for all } X \in \mathcal{S},\ h \in H.$$

We mainly focus on the case for $\mathrm{SO}(n)$, but much of the discussions will be applicable to H. Combining $SO(n)$-invariance with the centering discussed earlier, this will achieve $\mathrm{SE}(n)$-invariance.

3 Naive Invariantization Kills Chirality

We first consider two natural ways of obtaining invariantization of equivariant directional transforms. The first is to average over $\mathrm{SO}(n)$ (as in Puny et al. (2022); Ma et al. (2024) and more); the

second is to select canonical directions (following Kaba et al. (2023)). Both constructions produce $\mathrm{SO}(n)$-invariance, but can remove information needed to distinguish reflections.

3.1 Group Averaging

Observe that PHT, ECT, and BCT are all topological descriptors valued on the sublevel set induced by the filtration function $f_\nu(x) \coloneqq \nu \cdot x, \nu \in S^{n-1}$. We may attempt to turn f_ν into a $\mathrm{SO}(n)$-invariant function, and the descriptors along this new function would be $\mathrm{SO}(n)$. One way to do this is via group-averaging by averaging f_ν over $\mathrm{SO}(n)$ before building the directional transform.

Definition 3.1 (Averaging Method)**.** For $\nu \in S^{n-1}$, we define the averaging of f_ν as the function $\mathrm{avg}(f_\nu)(x) \coloneqq \int_{A \in \mathrm{SO}(n)} f_\nu(Ax)\, dA$. Accordingly, we write $\mathcal{I}^{\mathrm{avg}}_{SO(n)}[\Phi]$ for $\Phi \in \{\mathrm{BCT}, \mathrm{ECT}, \mathrm{PHT}\}$ (or, more generally, any Φ that is induced by f_ν on each ν) for the invariantization obtained by replacing each f_ν by $\mathrm{avg}(f_\nu)$ before applying the descriptor to the filtration.

The average has the desired $\mathrm{SO}(n)$-invariance. However, it also has a major drawback.

Proposition 3.2. The function $\mathrm{avg}(f_\nu)$ is $\mathrm{SO}(n)$-invariant, but $\mathrm{avg}(f_\nu)(x) = 0$ for any $x \in \mathbb{R}^n$.

For the PHT, ECT, and BCT, this means that averaging removes not only dependence on rotation, but also the directional signal needed to retain chirality.

3.2 Sorting Canonicalization

Another possible invariantization is to use the directional transform itself to choose a small collection of canonical directions. For each direction $\nu \in S^{n-1}$, we first compute the transform output $\Phi(X)(\nu)$. A scoring function s assigns a scalar score to this output, and the directions are then ranked according to these scores. We retain the k directions with the smallest scores.

Definition 3.3 (Canonicalization)**.** Let Φ be a directional transform, $X \in \mathcal{S}$, and s be a scoring function on the outputs $\{\Phi(X)(\nu)\}_{\nu \in S^{n-1}}$ such that the smallest k vectors $w_1, \ldots, w_k$ under s is well-defined. Then we define $\mathrm{Best}^{\Phi,s}_k(X) \coloneqq \{w_1, \ldots, w_k\}$.

Note that $\mathrm{Best}^{\Phi,s}_k(X)$ does not use a fixed set of laboratory-frame directions. The selected directions depend on X through the values of its directional transform. However, the following proposition shows that canonicalization cannot turn ECT, PHT, or BCT into chirality-sensitive transforms.

Proposition 3.4. Let H be a subgroup of $\mathrm{O}(n)$. Suppose Φ is H-equivariant, then $\mathrm{Best}^{\Phi,s}_k(-)$ is H-equivariant. However, the $X \mapsto \Phi(X)(\mathrm{Best}^{\Phi,s}_k(X))$ is H-invariant with respect to X.

In particular, for any $O(n)$-equivariant directional transforms (e.g., PHT, ECT, and BCT), the $SO(n)$-invariantization $\mathcal{I}^{\mathrm{can}}_{k,s,\mathrm{SO}(n)}[\Phi](X) \coloneqq \Phi(X)(\mathrm{Best}^{\Phi,s}_k(X))$ would also be $O(n)$-invariant.

4 Circular Topological Transforms (the Circle Method)

Section 3 shows that naive methods to invariantize will make directional transformations insensitive to reflections. Thus, we need to delicately work with our geometric setting to design H-invariantizations that are chiral sensitive. In this section, we will introduce a novel method - the **circular topological transform** (or the **circle method**) - to resolve this.

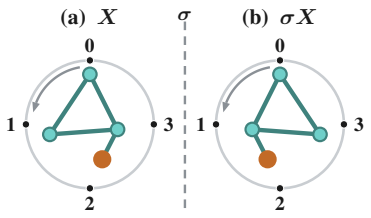

Figure 5: A shape X (a) and its mirror σX (b) in the frame described.

Motivating Example. To motivate the circle method in higher dimensions, we first give some intuition in 2D. Consider the displayed Figure 5 of a shape X and its mirror reflection σX, and 4 points placed evenly across the circle. Consider the ordered directional scan from $0 \to 1 \to 2 \to 3$. As we rotate X, the frame rotates with the shape and the same ordered scan is recovered in frame coordinates. Under reflection, the orientation of the shape relative to the proper frame is reversed. Thus, this design is invariant under rotation but detects reflections. This is a special instance of Theorem 4.4 and Theorem 4.6 later.

Let us now see how to generalize this to $\mathbb{R}^n$ for $n \geq 2$. Throughout this section, H is a subgroup of $SO(n)$. We first describe some preliminary concepts. For two shapes that differ from an action of H, we would like a way to choose an H-equivariant frame to align them together, in the form of a selector Fr. Instead of a single circle, we now have also multiple circles to pick from, which comes in the form of a discretization parameter $\mathbb{d}$.

Definition 4.1. We assemble them formally into the following data.

1. A **frame function** is a function Fr that takes in a shape $X \in \mathcal{S}$ and produces a list $(v_1, ..., v_{n-1})$ of **orthonormal vectors**. It is H-equivariant if $\mathrm{Fr}(h \cdot X) = h \cdot \mathrm{Fr}(X)$ for all $h \in H$.

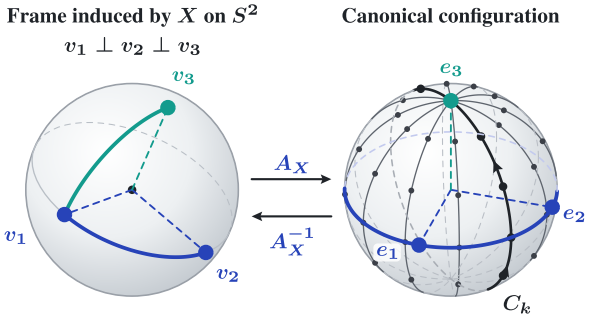

Figure 6: **The circle method in $\mathbb{R}^3$.** A shape X gives a frame $\mathrm{Fr}(X) = (v_1, v_2)$ on S^2, completed by v_3. We find the unique $A_X \in SO(3)$ with $A_X v_i = e_i$ and map to the canonical configuration, where the oriented circles C_k and directions $p_{k,j}$ are fixed. The ordered scan $w_{k,j}$ are given by mapping the $p_{k,j}$'s back along A_X^{-1}.

2. A **discretization parameter** is a list of ordered lists $\mathbb{d} := \left\{ \left((p_{k,j})_{j=1}^{m_k}, C_k \right)_{k=1}^{K} \right\}$, where K and m_k are positive integers, C_k is an **oriented circle** embedded in S^{n-1}, and the $p_{k,j}$'s are points on C_k in the order specified by the orientation of the circle.

Definition 4.2 (The Circle Method)**.** Let $X \in \mathcal{S}$ be a shape, Φ a directional transform on $\mathcal{S}$, Fr be a frame function, and $\mathbb{d}$ be a discretization parameter. The **circle method** is a transform $\mathfrak{C}$ that takes in $(X, \Phi, \mathrm{Fr}, \mathbb{d})$ and produces:

$$\mathfrak{C}(X, \Phi, \mathrm{Fr}, \mathbb{d}) := \left((\Phi(X)(w_{k,j}))_{j=1}^{m_k} \right)_{k=1}^{K}.$$

Here, for each fixed k, we construct the vectors $w_{k,j}$ as:

1. Write $\mathrm{Fr}(X) = (v_1, \ldots, v_{n-1})$ as the list of orthonormal vectors. There are exactly two unit vectors that are orthogonal to $v_1, ..., v_{n-1}$, and the two vectors are antipodal. We choose v_n as the unique vector so that $A_X := [v_1, ..., v_n]^{-1}$ has determinant 1.
2. From the discretization parameter $\mathbb{d}$, we pick $p_{k,1}, p_{k,2}, \ldots, p_{k,m_k}$ following a consistent cyclical ordering around the circle C_k. We obtain the vectors $w_{k,1}, \ldots, w_{k,m_k}$ as

$$[w_{k,1}, w_{k,2}, \ldots, w_{k,m_k}] = A_X^{-1} [p_{k,1}, p_{k,2}, \ldots, p_{k,m_k}].$$

Figure 6 displays the circle method in 3D. For a directional transform Φ, we call the method obtained by the circle method its corresponding **circular topological transforms** (CIRCUiT^2).

Definition 4.3. When $\Phi = \mathrm{PHT}, \mathrm{ECT}, \mathrm{BCT}$, we call the resulting circular topological transform Circ-PHT, Circ-ECT and Circ-BCT respectively.

4.1 Rotational Invariance and Reflection Detection

The first property we wish to show is that the circle method is H-invariant with respect to suitable conditions on the input.

Theorem 4.4. Let H be a subgroup of $\mathrm{SO}(n)$. Suppose Φ is **H-equivariant** and Fr is **H-equivariant**. The circular topological transform is then **H-invariant**, that is

$$\mathfrak{C}(X, \Phi, \mathrm{Fr}, \mathbb{d}) = \mathfrak{C}(h \cdot X, \Phi, \mathrm{Fr}, \mathbb{d}), \quad \text{for all } h \in H.$$

The next result we wish to discuss is to what extent is the circle method reflection-sensitive.

Definition 4.5. Let $\mathbb{H}$ be a hyperplane crossing the origin in $\mathbb{R}^n$. We write $\mathrm{flip}_{\mathbb{H}}$ to denote the reflection across the hyperplane $\mathbb{H}$.

Theorem 4.6. Let Fr be $O(n)$-equivariant, Φ be $O(n)$-equivariant, C be an oriented circle in $\mathbb{d}$, $\mathbb{H}$ be a hyperplane crossing the origin. Suppose $w_1, ..., w_k$ is the list of vectors chosen for X in the corresponding circle to C.

$$\mathfrak{C}(X, \Phi, \mathrm{Fr}, \mathbb{d})(C) = (\Phi(X)(w_1), ..., \Phi(X)(w_k))\,, \text{ and}$$

$$\mathfrak{C}(\mathrm{flip}_{\mathbb{H}} X, \Phi, \mathrm{Fr}, \mathbb{d})(C) = \left(\Phi(X)(\mathrm{flip}_{\mathrm{span}(\mathrm{Fr}(X))}(w_1)), ..., \Phi(X)(\mathrm{flip}_{\mathrm{span}(\mathrm{Fr}(X))}(w_k))\right).$$

Note that $H \subseteq SO(n)$ in Theorem 4.4, so Theorem 4.6 does not contradict Theorem 4.4.

Remark 4.7. Theorem 4.6 only tells us that the output lists can be potentially different, but the actual chiral-sensitivity depends on the transform Φ. If $\Phi(X)$ is defined to be the zero function, clearly it cannot tell apart a shape and its reflection.

In practice, we will take $\Phi = \mathrm{PHT}, \mathrm{ECT}, \mathrm{BCT}$, and the upshot of this section is that we can turn them into $\mathrm{SO}(n)$-invariant and reflection-sensitive transforms with an $O(n)$-equivariant frame.

4.2 Completeness

In this section, we would like to show that for suitable directional transforms (e.g., PHT, ECT, BCT), it is possible to establish our aforementioned completeness result on broad classes of shapes.

Definition 4.8. A directional transform Φ is **admissible** if it is $O(n)$-equivariant and satisfies:

1. (Continuity) For each $Z \in \mathcal{S}$, $\Phi(Z) : S^{n-1} \to T$ is a continuous function and T is Hausdorff.
2. (Injectivity) $\Phi(Z_1) \neq \Phi(Z_2)$ for distinct $Z_1, Z_2 \in \mathcal{S}$.

Example 4.9. When $\mathcal{S} \subseteq \mathcal{S}^{\Delta}$, the PHT, BCT, and ECT are admissible. Continuity is from Lemma 4.8, Corollary 4.12 of Curry et al. (2022). Injectivity is from Theorem 5 of Ghrist et al. (2018) or Theorem 3.5 of Curry et al. (2022), with the case $n = 2, 3$ earlier in Turner et al. (2014).

Definition 4.10. Let $\{\mathbb{d}_n\}_{n=1,2,...}$ be a sequence of discretization parameters. We say the sequence is **equidistributed** if the $\mathbb{d}_n$'s forms an equidistributed sequence in the sense of measure theory (see Appendix C for more details). Examples include when the $\mathbb{d}_n$'s are uniformly distributed.

Theorem 4.11. Suppose Φ is admissible, and $\{\mathbb{d}_n\}_{n\geq 1}$ is an equidistributed sequence of discretization parameters. Let $X, Y \in \mathcal{S}$ be two shapes that are not rotations of each other. Suppose Fr is $\mathrm{SO}(n)$-equivariant on the subcollection $\{g \cdot X \mid g \in \mathrm{SO}(n)\} \cup \{g \cdot Y \mid g \in \mathrm{SO}(n)\} \subseteq \mathcal{S}$. There exists some $N > 0$ such that for all $m \geq N$,

$$\mathfrak{C}(X, \Phi, \mathrm{Fr}, \mathbb{d}_m) \neq \mathfrak{C}(Y, \Phi, \mathrm{Fr}, \mathbb{d}_m).$$

Remark 4.12. In Appendix E.1, we explain how we construct frame functions Fr in our experiments. $O(3)$-equivariance holds in exact arithmetic when both selection stages are defined and determine unique output directions. The resulting transforms do quite well for chirality-detections and empirically exhibit completeness as in Theorem 4.11 in practice (see Section 5).

Combining Theorem 4.11 with the pre-processing step of **translation invariance** in Section 2, we see that the circular topological transforms of admissible Φ can be made to be $\mathrm{SE}(n)$-complete.

5 Experiments

We evaluate Circ-PHT and its variants on molecular retrieval, long-range chirality, and stereoisomer separation. The experiments address four questions. **RQ1.** Do circular transforms distinguish mirror images while remaining invariant to proper rotations? **RQ2.** What is the minimum number of directions needed for an order-sensitive circular topological transform? **RQ3.** Can circular transforms capture chirality determined by geometrically distant substructures? **RQ4.** Do they provide information not already present in standard molecular representations? Full experimental details and additional results appear in Appendix F.

5.1 Experimental setup

Datasets. For RQ1 and RQ2, we use QM9-OR (Zhou et al., 2025) and the BZR, COX2, and DHFR datasets from TUDataset (Morris et al., 2020). For each molecule X, we construct its reflection QX and eight rotated poses of each hand. All methods receive the unlabelled graph and 3D coordinates. Figure 7 shows a molecule from the BZR dataset in this setting. For RQ3, we use the noise-free Chi-Geometry benchmark (Weaver & Pasini, 2025), with the substructures determining the R/S label placed at graph distances $d = 1, \ldots, 9$ from the labelled centre. For RQ4, we use the 3DCS stereoisomer benchmark (Wang et al., 2026).

Table 1: **Handedness retrieval** (mAP; 88 directions). Best performer **bolded**. All rows: Table 5.

Method	QM9-OR	BZR	COX2	DHFR
Learned				
EGNN	0.557	0.548	0.545	0.552
GotenNet	0.557	0.546	0.558	0.558
EquiformerV2	**1.000**	0.998	**1.000**	0.999
ChiENN	0.952	0.578	0.467	0.538
PersLay	0.555	0.560	0.559	0.555
DECT	0.547	0.503	0.486	0.541
Unlearned				
VR-PH	0.554	0.554	0.558	0.553
Lab-PHT	0.056	0.046	0.019	0.033
Avg-PHT	0.004	0.012	0.009	0.006
Canon-PHT	0.551	0.552	0.552	0.556
CIRC-PHT	**1.000**	**1.000**	**1.000**	**1.000**
CIRC-ECT	**1.000**	**1.000**	**1.000**	**1.000**
CIRC-BCT	**1.000**	**1.000**	**1.000**	**1.000**

Baselines. Avg-PHT and Canon-PHT are the two invariantization schemes from Section 3: Avg-PHT averages the directional height function over $SO(3)$, while Canon-PHT evaluates PHT at $O(3)$-equivariant canonical directions. The geometric baselines are EGNN (Satorras et al., 2021), GotenNet (Aykent & Xia, 2025), EquiformerV2 (Liao et al., 2024), and ChiENN (Gaiński et al., 2023). Topological baselines include Vietoris–Rips PH (Zomorodian & Carlsson, 2005), laboratory-frame PHT/ECT/BCT (Turner et al., 2014), PersLay (Carrière et al., 2020), DECT (Roell & Rieck, 2024), and Neural k-Forms (Maggs et al., 2024). For RQ4, we additionally compare with MACE (Batatia et al., 2022), E3FP (Axen et al., 2017), and UniMol (Zhou et al., 2023). Our unlearned methods are CIRC-PHT, CIRC-ECT, and CIRC-BCT. Additional baselines are reported in Appendix F.

Evaluation protocol. For retrieval, each pose is queried against all remaining test poses, with the seven other rotations of the same hand treated as relevant. We report mean average precision (mAP) on test fold 0. A rotation-invariant, reflection-blind descriptor ties these seven poses with the eight poses of the reflected hand, giving expected mAP 0.555. For long-range chirality, we report balanced handedness accuracy (chance 0.5) using frozen circular descriptors with a logistic-regression readout. For 3DCS, we report zero-shot ES-AUC, the ROC AUC of pairwise embedding distances for conformers from different versus the same stereoisomer.

5.2 RESULTS

Invariance and chirality (RQ1). VR-PH, EGNN, GotenNet, PersLay, DECT, and Canon-PHT remain near the reflection-blind value of 0.555 (0.486–0.560 in Table 1). The first five are reflection-invariant, while Canon-PHT loses handedness by Proposition 3.4. Avg-PHT falls to 0.004–0.012, consistent with the collapse of the averaged height function in Proposition 3.2. Lab-PHT is also poor (0.019–0.056), since its fixed directions retain pose dependence. Among the learned baselines, EquiformerV2 separates the two hands consistently across all 4 datasets. ChiENN is strong on QM9-OR but not on the other three datasets, where the unlabelled bond graph less often distinguishes the neighbour ordering needed by its chiral mechanism (Appendix F.1). CIRC-PHT, CIRC-ECT, and CIRC-BCT obtain mAP 1.000 on all four datasets without training.

Table 2: **Directions needed** (mAP; one circle, D directions). Best performer is **bolded**. CIRC-BCT is within 0.002 of CIRC-ECT. Full sweep: Table 6.

		Scan directions D			
Data	**Read-out**	2	3	4	8
QM9-OR	CIRC-PHT	0.552	**1.000**	**1.000**	**1.000**
	CIRC-ECT	0.555	**1.000**	0.994	**1.000**
BZR	CIRC-PHT	0.557	**1.000**	**1.000**	**1.000**
	CIRC-ECT	0.555	0.994	0.975	**1.000**
COX2	CIRC-PHT	0.552	**1.000**	**1.000**	**1.000**
	CIRC-ECT	0.539	**1.000**	**1.000**	**1.000**
DHFR	CIRC-PHT	0.551	**1.000**	**1.000**	**1.000**
	CIRC-ECT	0.555	0.994	0.991	0.994

Minimum scan directions (RQ2). Table 2 varies the number of directions on a single circle. At $D = 2$, reversal leaves the scan unchanged and the circular descriptors remain reflection-blind. At $D = 3$, CIRC-PHT reaches mAP 1.000 on all four datasets. CIRC-ECT and CIRC-BCT require more directions on BZR and DHFR: BZR first reaches 1.000 at $D = 8$ and is exact at every tested resolution from $D = 16$ on (0.994 at $D = 12$), and DHFR first reaches 1.000 at $D = 32$ (Table 6). The three methods use the same frame and scan. Euler and Betti curves retain feature counts at each filtration value, while persistence also records birth–death pairs. At four directions and the same output, persistence landscapes (Bubenik, 2015) obtain 1.000 on BZR and DHFR, compared with

Table 3: **Long-range chirality.** Balanced handedness accuracy at bond distance d (chance 0.5). CIRC descriptors frozen; CIRC-ECT and CIRC-BCT give identical accuracies. EquiformerV2 has $L = 4$; R is the cutoff of a rebuilt graph. Best value per column **bolded**. All rows: Table 8.

Distance d (bonds)	1	2	3	4	5	6	7	8	9
CIRC-PHT	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
CIRC-ECT, CIRC-BCT	0.861	0.793	**1.000**	0.682	0.863	0.822	**1.000**	0.904	0.845
CIRC-PHT, direction order pooled	0.509	0.494	0.496	0.501	0.504	0.486	0.504	0.509	0.492
EquiformerV2, bond graph	**1.000**	**1.000**	**1.000**	0.498	0.495	0.500	0.511	0.499	0.498
EquiformerV2, $R = 1.2$ Å	**1.000**	**1.000**	**1.000**	0.498	0.489	0.498	0.500	0.496	0.496
EquiformerV2, $R = 2.5$ Å	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	0.496	0.507	0.487
ChiENN	0.507	0.502	0.482	0.509	0.484	0.509	0.489	0.522	0.496

0.975 and 0.991 for Betti curves (Table 7). Betti and Euler curves differ by at most 0.002 in this experiment.

Long-range chirality (RQ3). Table 3 varies the bond distance d; atomic species are included through the fixed radial lift described in Appendix F.3. CIRC-PHT reaches accuracy 1.000 at every distance, while CIRC-ECT and CIRC-BCT range from 0.682 to 1.000 with 88 directions and from 0.923 to 1.000 with 704 (Table 8); the accuracy of Euler and Betti curves increase as the number of directions increase (Appendix F.3). "Direction order pooled" in CIRC-PHT is an averaging operation that removes the order of the scans, which reduces performance.

EquiformerV2 on the bond graph is 1.0 through $d = 3$ and near chance from $d = 4$ onward for every tested depth $L \in 3, 4, 5, 7$ (Table 8). Rebuilding the graph with radius $R = 1.2$ Å reproduces the bond-graph behaviour (exact through $d = 3$), while $R = 2.5$ Å is exact until $d = 6$, at the cost of 8–9× more edges than the bond graph; $R = 5$ Å is exact at every tested distance with 9.5–14× more edges (Table 8). This dependence on graph connectivity is consistent with a finite receptive field. The directional filtration used by the circular descriptors is global and does not require a radius parameter. ChiENN remains with a low performance because the relevant neighbour cycles contain 2 elements, whose circular order is unchanged by reversal.

Table 4: **3DCS**, zero-shot ES-AUC.

Method	**ES-AUC**
CIRC *descriptors*	
CIRC-PHT, 88 dirs.	0.707
CIRC-PHT, 4 dirs.	0.706
CIRC-BCT	0.700
CIRC-ECT	0.698
Molecular representations	
MACE	0.467
E3FP	0.580
UniMol	0.622
Combined with CIRC-PHT, 4 *dirs.*	
MACE	0.704
E3FP	0.711
UniMol	**0.714**

Complementarity on 3DCS (RQ4). On 3DCS, CIRC-PHT reaches zero-shot ES-AUC 0.707 with 88 directions and 0.706 with four, compared with 0.467 for MACE, 0.580 for E3FP, and 0.622 for UniMol (Table 4). Concatenating the four-direction CIRC-PHT descriptor gives 0.704, 0.711, and 0.714, respectively. The E3FP and UniMol combinations exceed either constituent representation alone by 0.005 and 0.008. These gains show that the circular descriptor contains information not already captured by those representations. Further details are given in Appendix F.4.

6 CONCLUSION

We introduced circular topological transforms for constructing $\mathrm{SE}(n)$-invariant, chirality-sensitive representations from directional topological transforms. A shape-dependent equivariant frame removes proper rotations, while reflection reverses the ordered circular scan; for continuous injective directional transforms, equidistributed refinements separate distinct rotational orbits (complete descriptor) on broad shape classes. In the molecular experiments, CIRC-PHT achieves perfect handedness retrieval with three directions and retains the chirality signal across long graph distances without changing a neighbourhood radius. On 3DCS, adding CIRC-PHT also improves the tested molecular representations. Our experiments use fixed circles and fixed scan directions, so we have not studied whether learning the number, placement, or resolution of the circles could improve accuracy or reduce the number of directional evaluations. We also evaluate only three-dimensional molecular graphs. Although the construction is defined in general dimension, its empirical behavior in higher-dimensional geometric data remain unexplored.

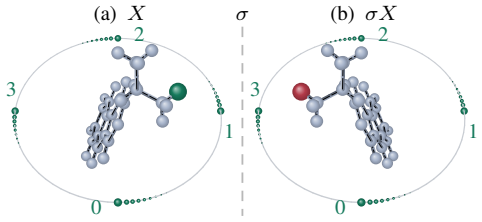

Figure 7: A chiral BZR ligand and its mirror image. The green dots are the directions.

Ethics Statement

This paper presents work whose goal is to make directional transforms arising from topological data analysis into invariant/equivariant architectures. Our main methodology is directly applicable to biology and chemistry, but they also applies to many higher dimensional data arising in settings outside of these areas. There may be potential societal consequences of our work, none of which we feel particularly necessary to mention here.

Reproducibility Statement

All data used in this paper is publicly available and can be easily reproduced, where we give more details in Appendix F. The coding used for experiments in this paper can be found available at the following anonymized repository: https://anonymous.4open.science/r/anonymousInvPHT-5E16/. Appendix B contains the full statements and proofs of all the propositions and theorems.

AI Use Statement

We used AI/LLM tools mainly to assist with text editing, coding for experiments and plots, and literature searches to identify related work and assess the novelty and positioning of our contributions. The paper's main methodology was developed by the authors without AI assistance. Most proofs were initially written by the authors, and LLMs provided limited assistance with proofs of some lemmas and were also used to help check proofs and improve their exposition. All proofs, including those developed with LLM assistance, were checked and verified by the authors. The authors take full responsibility for the content and correctness of the final manuscript.

References

Uri Alon and Eran Yahav. On the bottleneck of graph neural networks and its practical implications. *arXiv preprint arXiv:2006.05205*, 2020.

Dan Amir. *Characterizations of Inner Product Spaces*, volume 20 of *Operator Theory: Advances and Applications*. Birkhäuser Basel, 1986. doi: 10.1007/978-3-0348-5487-0.

Kenneth Atz, Francesca Grisoni, and Gisbert Schneider. Geometric deep learning on molecular representations. *Nature Machine Intelligence*, 3(12):1023–1032, 2021.

Seth D. Axen, Xi-Ping Huang, Elena L. Cáceres, Leo Gendelev, Bryan L. Roth, and Michael J. Keiser. A simple representation of three-dimensional molecular structure. *Journal of Medicinal Chemistry*, 60(17):7393–7409, 2017. doi: 10.1021/acs.jmedchem.7b00696.

Sarp Aykent and Tian Xia. GotenNet: Rethinking efficient 3d equivariant graph neural networks. In *International Conference on Learning Representations*, 2025.

Rubén Ballester and Bastian Rieck. On the expressivity of persistent homology in graph learning. In Guy Wolf and Smita Krishnaswamy (eds.), *Proceedings of the Third Learning on Graphs Conference*, volume 269 of *Proceedings of Machine Learning Research*, pp. 42:1–42:31. PMLR, 26–29 Nov 2025. URL https://proceedings.mlr.press/v269/ballester25a.html.

Levent Batakci, Abigail Branson, Bryan Castillo, Candace Todd, Erin Wolf Chambers, and Elizabeth Munch. Comparing embedded graphs using average branching distance. *Involve, a Journal of Mathematics*, 16(3):365–388, August 2023. ISSN 1944-4176. doi: 10.2140/involve.2023.16.365. URL http://dx.doi.org/10.2140/involve.2023.16.365.

Ilyes Batatia, David Peter Kovacs, Gregor N. C. Simm, Christoph Ortner, and Gabor Csanyi. MACE: Higher order equivariant message passing neural networks for fast and accurate force fields. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL https://openreview.net/forum?id=YPpSngE-ZU.

Robin Lynne Belton, Brittany Terese Fasy, Rostik Mertz, Samuel Micka, David L. Millman, Daniel Salinas, Anna Schenfisch, Jordan Schupbach, and Lucia Williams. Learning simplicial complexes from persistence diagrams, 2018. URL https://arxiv.org/abs/1805.10716.

Nicholas Bermingham, Vanessa Robins, and Katharine Turner. Planar symmetry detection and quantification using the extended persistent homology transform. In *2023 Topological Data Analysis and Visualization (TopoInVis)*, pp. 1–9, 2023. doi: 10.1109/TopoInVis60193.2023.00007.

Patrick Billingsley. *Convergence of Probability Measures*. John Wiley & Sons, Ltd, 1999. ISBN 9780470316962. doi: https://doi.org/10.1002/9780470316962.

Theodor Bröcker and Tammo tom Dieck. *Representations of Compact Lie Groups*, volume 98 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1985. ISBN 978-3-662-12918-0. doi: 10.1007/978-3-662-12918-0.

Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.

Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

Peter Bubenik. Statistical topological data analysis using persistence landscapes. *Journal of Machine Learning Research*, 16(3):77–102, 2015. URL http://jmlr.org/papers/v16/bubenik15a.html.

Robert S Cahn, Christopher Ingold, and Vladimir Prelog. Specification of molecular chirality. *Angewandte Chemie International Edition in English*, 5(4):385–415, 1966.

Mathieu Carrière, Frédéric Chazal, Yuichi Ike, Théo Lacombe, Martin Royer, and Yuhei Umeda. Perslay: A neural network layer for persistence diagrams and new graph topological signatures. In *International Conference on Artificial Intelligence and Statistics*, pp. 2786–2796. PMLR, 2020.

Erin Chambers, Shankha Shubhra Mukherjee, and Katharine Turner. The Reeb Transform, June 2025. URL https://socg25.github.io/yrf.html. Extended abstract at the 41st International Symposium on Computational Geometry: Young Researchers Forum (CG:YRF 2025); Non-archival Contribution.

Erin Wolf Chambers, Elizabeth Munch, Sarah Percival, and Xinyi Wang. A distance for geometric graphs via the labeled merge tree interleaving distance, 2024. URL https://arxiv.org/abs/2407.09442.

Zhengdao Chen, Lei Chen, Soledad Villar, and Joan Bruna. Can graph neural networks count substructures? In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 10383–10395. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/75877cb75154206c4e65e76b88a12712-Paper.pdf.

Lorin Crawford, Anthea Monod, Andrew X. Chen, Sayan Mukherjee, and Raúl Rabadán. Predicting clinical outcomes in glioblastoma: An application of topological and functional data analysis. *Journal of the American Statistical Association*, 115(531):1139–1150, 2020. doi: 10.1080/01621459.2019.1671198. URL https://doi.org/10.1080/01621459.2019.1671198.

Justin Curry, Sayan Mukherjee, and Katharine Turner. How many directions determine a shape and other sufficiency results for two topological transforms. *Transactions of the American Mathematical Society, Series B*, 9:1006–1043, 2022. doi: 10.1090/btran/122.

Tamal Krishna Dey and Yusu Wang. *Computational Topology for Data Analysis*. Cambridge University Press, 2022.

Alexandru Dumitrescu, Dani Korpela, Markus Heinonen, Yogesh Verma, Valerii Iakovlev, Vikas Garg, and Harri Lähdesmäki. E(3)-equivariant models cannot learn chirality: Field-based molecular generation. In *International Conference on Learning Representations*, 2025.

Herbert Edelsbrunner and Dmitriy Morozov. Persistent homology: Theory and practice. In *Proceedings of the European Congress of Mathematics*, 2012.

Ernest L Eliel and Samuel H Wilen. *Stereochemistry of organic compounds*. John Wiley & Sons, 1994.

Brittany Terese Fasy, Samuel Micka, David L. Millman, Anna Schenfisch, and Lucia Williams. Challenges in reconstructing shapes from euler characteristic curves, 2018. URL https://arxiv.org/abs/1811.11337.

Brittany Terese Fasy, Samuel Micka, David L. Millman, Anna Schenfisch, and Lucia Williams. A Faithful Discretization of Verbose Directional Transforms. *Discrete & Computational Geometry*, 75(3):904–949, November 2025. ISSN 1432-0444. doi: 10.1007/s00454-025-00791-w. URL http://dx.doi.org/10.1007/s00454-025-00791-w.

Gerald B. Folland. *Real Analysis: Modern Techniques and Their Applications*. Pure and Applied Mathematics. John Wiley & Sons, New York, 2nd edition, 1999.

Jesús Gacías Franco. The Euler Characteristic Transform from a Convex Geometric Perspective, 2026. URL https://arxiv.org/abs/2607.28021.

Piotr Gaiński, Michał Koziarski, Jacek Tabor, and Marek Śmieja. Chienn: embracing molecular chirality with graph neural networks. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 36–52. Springer, 2023.

Vikas Garg, Stefanie Jegelka, and Tommi Jaakkola. Generalization and representational limits of graph neural networks. In *International conference on machine learning*, pp. 3419–3430. PMLR, 2020.

Johannes Gasteiger, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs. *arXiv preprint arXiv:2003.03123*, 2020.

Johannes Gasteiger, Florian Becker, and Stephan Günnemann. GemNet: Universal directional graph neural networks for molecules. In *Advances in Neural Information Processing Systems*, 2021.

Robert Ghrist, Rachel Levanger, and Huy Mai. Persistent homology and euler integral transforms. *Journal of Applied and Computational Topology*, 2(1):55–60, Oct 2018. ISSN 2367-1734. doi: 10.1007/s41468-018-0017-1. URL https://doi.org/10.1007/s41468-018-0017-1.

Allen Hatcher. *Algebraic Topology*. New York : Cambridge University Press, 2002.

Max Horn, Edward De Brouwer, Michael Moor, Yves Moreau, Bastian Rieck, and Karsten Borgwardt. Topological graph neural networks. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=oxxUMeFwEHd.

Johanna Immonen, Amauri H. Souza, and Vikas Garg. Going beyond persistent homology using persistent homology. In *Advances in Neural Information Processing Systems*, 2023.

Mattie Ji, Amauri H Souza, and Vikas K Garg. Graph persistence goes spectral. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=wU8IKGLpbi.

Chaitanya K Joshi, Cristian Bodnar, Simon V Mathis, Taco Cohen, and Pietro Lio. On the expressive power of geometric graph neural networks. In *International conference on machine learning*, pp. 15330–15355. PMLR, 2023.

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.

Sékou-Oumar Kaba, Arnab Kumar Mondal, Yan Zhang, Yoshua Bengio, and Siamak Ravanbakhsh. Equivariance with learned canonicalization functions. In *International Conference on Machine Learning*, pp. 15546–15566. PMLR, 2023.

Achim Klenke. *Probability Theory: A Comprehensive Course*. Universitext. Springer, London, 1 edition, 2008. ISBN 978-1-84800-048-3. doi: 10.1007/978-1-84800-048-3.

Yunyang Li, Yusong Wang, Lin Huang, Han Yang, Xinran Wei, Jia Zhang, Tong Wang, Zun Wang, Bin Shao, and Tie-Yan Liu. Long-short-range message-passing: A physics-informed framework to capture non-local interaction for scalable molecular dynamics simulation. In *International Conference on Learning Representations*, volume 2024, pp. 37781–37820, 2024.

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.

George Ma, Yifei Wang, Derek Lim, Stefanie Jegelka, and Yisen Wang. A canonicalization perspective on invariant and equivariant learning. *Advances in Neural Information Processing Systems*, 37:60936–60979, 2024.

Kelly Maggs, Celia Hacker, and Bastian Rieck. Simplicial representation learning with neural k-forms. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=Djw0XhjHZb.

Lewis Marsh, Felix Y Zhou, Xiao Qin, Xin Lu, Helen M Byrne, and Heather A Harrington. Detecting temporal shape changes with the Euler characteristic transform. *Transactions of Mathematics and its Applications*, 8(2):tnae002, 2024.

Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *2017 IEEE conference on computer vision and pattern recognition (CVPR)*, pp. 5425–5434. IEEE, 2017.

Christopher Morris, Nils M. Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion Neumann. Tudataset: A collection of benchmark datasets for learning with graphs, 2020. URL https://arxiv.org/abs/2007.08663.

Elizabeth Munch. An invitation to the Euler characteristic transform. *The American Mathematical Monthly*, 132(1):15–25, 2025.

Lien Ai Nguyen, Hua He, and Chuong Pham-Huy. Chiral drugs: an overview. *International journal of biomedical science: IJBS*, 2(2):85, 2006.

Taejin Paik. Invariant representations of embedded simplicial complexes, 2023. URL https://arxiv.org/abs/2302.13565.

Omri Puny, Matan Atzmon, Heli Ben-Hamu, Ishan Misra, Aditya Grover, Edward J. Smith, and Yaron Lipman. Frame averaging for invariant and equivariant network design. 2022. 10th International Conference on Learning Representations, ICLR 2022 ; Conference date: 25-04-2022 Through 29-04-2022.

Ernst Roell and Bastian Rieck. Differentiable Euler characteristic transforms for shape classification. In *International Conference on Learning Representations*, volume 2024, pp. 44481–44492, 2024.

Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pp. 8459–8468. Pmlr, 2020.

Víctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E(n) equivariant graph neural networks. In *Proceedings of the 38th International Conference on Machine Learning*, 2021.

Anna Schenfisch, Brittany Terese Fasy, Samuel Micka, David L. Millman, and Lucia Williams. Efficient graph reconstruction and representation using augmented persistence diagrams. In Yeganeh Bahoo and Konstantinos Georgiou (eds.), *Proceedings of the 34th Canadian Conference on Computational Geometry, CCCG 2022, Toronto Metropolitan University, Toronto, Ontario, Canada, August 25-27, 2022*, pp. 284–292, 2022.

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*, 2021.

Katharine Turner, Sayan Mukherjee, and Doug M. Boyer. Persistent homology transform for modeling shapes and surfaces. *Information and Inference: A Journal of the IMA*, 2014.

Katharine Turner, Vanessa Robins, and James Morgan. The extended persistent homology transform of manifolds with boundary. *Journal of Applied and Computational Topology*, 8(7):2111–2154, Nov 2024. ISSN 2367-1734. doi: 10.1007/s41468-024-00175-8. URL https://doi.org/10.1007/s41468-024-00175-8.

Yogesh Verma, Amauri H. Souza, and Vikas Garg. Topological neural networks go persistent, equivariant, and continuous. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.

Julius von Rohrscheidt and Bastian Rieck. Diss-l-ECT: Dissecting graph data with local Euler characteristic transforms. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 61790–61809. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/von-rohrscheidt25a.html.

Weikang Wang, Tobias Weißberg, Nafie El Amrani, and Florian Bernard. χ: Symmetry understanding of 3d shapes via chirality disentanglement. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 28292–28302. IEEE, 2025.

Xi Wang, Yang Zhang, Yingjia Zhang, Yejia Cai, and Shenji Wan. 3DCS: Datasets and benchmark for evaluating conformational sensitivity in molecular representations. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=JAb0y8lkqL.

Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (tog)*, 38(5): 1–12, 2019.

Rylie Weaver and Massamiliano Lupo Pasini. Chi-geometry: A library for benchmarking chirality prediction of gnns, 2025. URL https://arxiv.org/abs/2508.09097.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.

Haochen Yang, Vadim Lebovici, Andreas Tarcevski, Liliana Tchernev, Saulius Zuklys, Georg A. Holländer, Helen M. Byrne, and Heather A. Harrington. Topological shape transform for thymus structures, 2026. URL https://arxiv.org/abs/2602.18889.

Manzil Zaheer, Satwik Kottur, Siamak Ravanbhakhsh, Barnabás Póczos, Ruslan Salakhutdinov, and Alexander J Smola. Deep sets. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pp. 3394–3404, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.

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 *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=6K2RM6wVqKu.

Yilin Zhou, Haoran Zhu, Yijie Yuan, Ziyu Song, and Brendan C. Mort. Machine learning classification of chirality and optical rotation using a simple one-hot encoded Cartesian coordinate molecular representation. *Journal of Chemical Information and Modeling*, 2025. doi: 10.1021/acs.jcim.4c02374.

Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 2005. doi: 10.1007/s00454-004-1146-y.

Circular Topological Transforms: Geometric Chirality *with* Completeness (Appendix)

A Related Work

Directional topological transforms and Invariance. We continue the discussion from the introduction here to compare our method with existing literatures.

1. The work by Curry et al. (2022) gave a construction based on the ECT that is $O(d)$-invariant and injective on the embedded simplicial complexes, but it does not accomplish what we want here as it cannot detect reflections.
2. The work by Marsh et al. (2024) introduces DETECT, which builds on Curry et al. (2022) that integrates out the dependence on directions for rotationally invariant descriptors. However, from the construction, their DETECT is also reflection invariant, and they do not establish any completeness/finiteness guarantees of the DETECT method.
3. The work by Yang et al. (2026) introduced SampEuler to approximate the construction in Curry et al. (2022), so their construction is also $O(d)$-invariant and do not address reflection detection.
4. The work by Franco (2026) produces variants of the ECT and PHT with the concept of Steiner points in convex geometry, their resulting descriptors are isometry invariants (and hence cannot detection reflections).
5. The work by Paik (2023) combines the ECT and PHT with spherical neural processing and includes a reflection-sensitive $\mathrm{SE}(3)$-invariant but $\mathrm{E}(3)$-sensitive operator. Their main construction was done using the spherical Fourier transform and could be generalized to higher dimensions. However, Paik (2023) did not establish any completeness guarantees or finiteness guarantees on the directions needed.
6. The work by von Rohrscheidt & Rieck (2025) defines an ECT-based distance on shapes and prove that it is a metric on the rotational orbits. Precisely, it is defined as

$$d_{\mathrm{ECT}}(X, Y) \coloneqq \inf_{\rho \in SO(n)} ||\,\mathrm{ECT}(X) - \mathrm{ECT}(\rho Y)||_{\infty}.$$

   However it appears that to compute this metric, one needs to rotate Y across many different samples on $\rho \in SO(n)$. This requires an optimisation over rotations per comparison and yields a distance, instead of a descriptor. In contrast, our canonical frame assigns each shape a single invariant representation, with completeness proven for a finite scan. von Rohrscheidt & Rieck (2025) did not investigate further

Our contribution is a general completeness-preserving construction: CIRCUiT^2 converts any admissible directional transforms into $\mathrm{SE}(n)$-complete representations through our circular construction. On shape classes admitting equivariant frames, the family of circular transforms is complete modulo proper rigid motions. For each pair in distinct orbits, equidistributed finite scans separate the pair at every sufficiently fine stage.

It is also important to note that Paik (2023); von Rohrscheidt & Rieck (2025) appears to mainly work in the setting of embedded simplicial complexes or semi-algebraic sets, whereas our main theorems work in the setting of arbitrary shapes as in Definition 2.1 (a collection of non-empty compact subsets of $\mathbb{R}^n$ closed under Euclidean isometries).

We further establish finiteness guarantees under explicit hypotheses. Our methodology is also discrete and combinatorial in nature.

Other Directional Transforms in TDA. Although we have focused on the persistent homology transform, the Euler characteristic transform, and the Betti curve transform as our main sources of examples in the main text (Example 2.8), there are plenty of other topological descriptors and their corresponding transforms that have been studied in the TDA literature. To name a few,

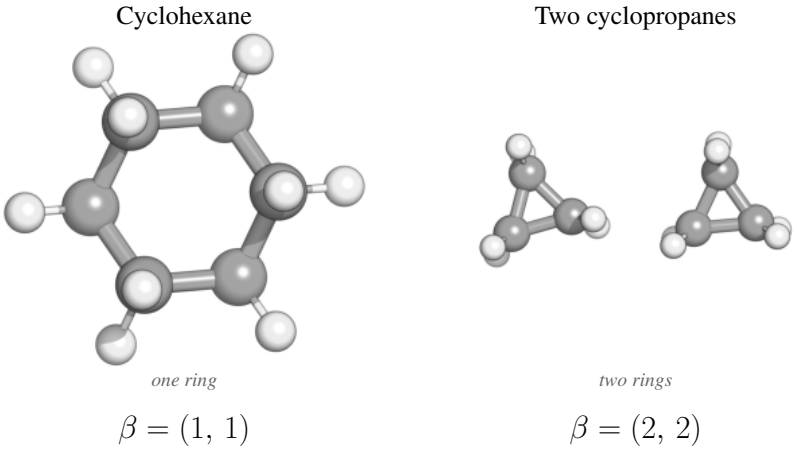

Figure 8: Cyclohexane and two disconnected cyclopropanes have the same number of vertices and the same degree-2 local neighborhoods. With identical node features, an untrained GNN assigns the two graphs the same embedding, while their Betti numbers $\beta = (\text{components}, \text{loops})$ are $(1, 1)$ and $(2, 2)$.

1. On the algebraic side, the extended persistent homology transform (Turner et al., 2024; Bermingham et al., 2023) studies the extended persistence arising from directional filtrations.
2. On more geometric sides, directional transforms producing (labeled) merge trees have been studied in Batakci et al. (2023) and Chambers et al. (2024), and direction transforms producing Reeb graphs have been studied in Chambers et al. (2025).
3. On the more analytic side, Crawford et al. (2020) proposed the smooth Euler characteristic transform which augments the ECT with more regularity to study clinical outcomes in glioblastoma.

Recently,Fasy et al. (2025) had proposed a verbose version of arbitrary directional transforms and the faithful discretizations for the verbose PHT/ECT/BCT.

There are also interests in practically reconstructing the original shape based on the given directional transform. Curry et al. (2022) gave an explicit bound on the (finite) number of directions ECT needs to distinguish two embedded simplicial complexes when the class of shapes $\mathcal{S}$ has a uniform bound on the angle at its vertices. Practical reconstructions of graphs and simplicial complexes from the PHT are in Belton et al. (2018); Schenfisch et al. (2022), while reconstructions from the ECT are considered in Fasy et al. (2018).

TDA in Graph ML In Figure 8 we show a pair of graphs that 1-WL GNNs Xu et al. (2019) cannot distinguish. This is one of the main motivations to incorporate TDA into ML pipelines, even more because a huge literature in GNNs show that these topological blindspots of these models are a intrisic part of how they work (Chen et al., 2020; Garg et al., 2020). Topological deep learning comes from this gap and numerous recent works develop better ways to incorporate TDA and thus PH-based descriptors into neural networks (Carrière et al., 2020; Horn et al., 2022; Immonen et al., 2023). A nice extension of our work would be to do our entire pipeline learnable end-to-end, that is, the number of sampled directions and circles could be learned based on the downstream task. Recent works have also tried to incorporating equivariance and TDA together in ML pipelines (Verma et al., 2024)).

B PROOFS

B.1 PROOFS IN SECTION 3.1

The proof of Proposition 3.2 requires the following lemma.

Lemma B.1 (Equation (4.1) in Chapter II of Bröcker & tom Dieck (1985))**.** Let $\rho : G \to \mathrm{GL}(V)$ be a matrix representation and dg be the normalized Haar measure then the matrix

$$\int_G \rho(g) dg$$

is the projection matrix onto the subspace $V^G = \{v \in V : \rho(g)v = v, \text{ for all } g \in G\}$.

Proof of Proposition 3.2. To show $\mathrm{SO}(n)$-invariance, let $B \in \mathrm{SO}(n)$, then observe that we can write

$$\mathrm{avg}(f_v)(B \cdot x) = \int_{A \in \mathrm{SO}(n)} f_v(ABx) dA = \int_{C = AB \in \mathrm{SO}(n)} f_v(Cx) dC = \mathrm{avg}(f_v)(x).$$

To show $\mathrm{avg}(f_v)(x) = 0$, we can write

$$\int_{g \in \mathrm{SO}(n)} \nu \cdot (gx) dg = \nu \cdot ((\int_{g \in \mathrm{SO}(n)} \mathrm{g}\, dg)x) = \nu \cdot (0x) = 0.$$

Here, $\int_{g \in \mathrm{SO}(n)} \mathrm{g}\, dg$ is the zero matrix. This follows from Lemma B.1 by choosing $\rho : \mathrm{SO}(n) \hookrightarrow \mathrm{GL}_n(\mathbb{R})$ as the natural inclusion map, in which the subspace of fixed points is trivial. □

B.2 PROOFS IN SECTION 3.2

Proof of Proposition 3.4. Let $h \in H$ and consider the shapes X and $h \cdot X$. Since Φ is H-equivariant, we see there is an equality of sets

$$\{\Phi(h \cdot X)(\nu)\}_{\nu \in S^{n-1}} = \{\Phi(X)(h^{-1} \cdot \nu)\}_{\nu \in S^{n-1}} = \{\Phi(X)(w)\}_{w = h^{-1} \cdot \nu \in S^{n-1}}.$$

It follows that if $\mathrm{Best}_k^{\Phi,s}(X) = \{w_1, ..., w_k\}$, then $\mathrm{Best}_k^{\Phi,s}(h \cdot X)$ is a list of vectors $w_1', ..., w_k'$ such that $h^{-1} \cdot w_i' = w_i$ and hence $w_i' = h \cdot w_i$. Thus, the output $\mathrm{Best}_k^{\Phi,s}(-)$ is H-equivariant. Since Φ is H-equivariant, it follows that $\Phi(-)(\mathrm{Best}_k^{\Phi,s}(-))$ is H-invariant. □

B.3 PROOFS IN SECTION 4.1

The goal of this part is to prove Theorem 4.4 and Theorem 4.6. We will first prove Theorem 4.4.

Proof of Theorem 4.4. For simplicity we drop the subscript for A_X. Write $(v_1, ..., v_{n-1})$ as $\mathrm{Fr}(X)$. Let $v_n \in S^{n-1}$ be the unique vector such that $[v_1, ..., v_{n-1}, v_n]$ is a special orthogonal matrix. The left hand side creates the matrix A

$$A = [v_1, ..., v_n]^{-1}.$$

Since $h \in H \subseteq \mathrm{SO}(n)$ and Fr is H-equivariant, the right hand side creates the matrix

$$A' = [hv_1, ..., hv_n]^{-1} = (h \cdot [v_1, ..., v_n])^{-1} = [v_1, ..., v_n]^{-1} h^{-1}.$$

The composition $(A')^{-1} \circ A$ is then

$$h \cdot [v_1, ..., v_n] \circ [v_1, ..., v_n]^{-1} = h.$$

This shows that any vector v used in the discretization scheme of the circle method on the left corresponds to the vector $h \cdot v$ in the discretization scheme of the circle method on the right. The bijection also preserves the order of the list, and hence $\mathfrak{C}(X, \Phi, \mathrm{Fr}, \mathfrak{d}) = \mathfrak{C}(h \cdot X, \Phi, \mathrm{Fr}, \mathfrak{d})$ as Φ is H-equivariant. □

In the remainder of this part, we will work to prove Theorem 4.6. Let us first set-up some constructions.

Let $\mathbb{H}$ be a hyperplane crossing the origin. Recall that we use $\text{flip}_{\mathbb{H}}$ to denote reflection across $\mathbb{H}$. Note, as a matrix, $\text{flip}_{\mathbb{H}}$ may be written as

$$\text{flip}_{\mathbb{H}} = RFR^{-1}, \text{ where } R \in \text{SO}(n), F = \begin{pmatrix} 1 & 0 & ... & 0 \\ 0 & 1 & ... & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & ... & -1 \end{pmatrix}.$$

It follows directly from the definition of $O(n)$-equivariance for directional transforms that we have the following.

Lemma B.2. Let Φ be $O(n)$-equivariant, then $\Phi(X)(A_X^{-1}\omega) = \Phi(A_X X)(\omega)$.

We now establish the following lemma.

Lemma B.3. Let $X \in \mathcal{S}$, Fr be $O(n)$-equivariant, and $\mathbb{H}$ be a hyperplane crossing the origin. Write $Y = \text{flip}_{\mathbb{H}} X$, then

$$A_Y^{-1} \circ A_X \text{ is reflect over } \text{span}(v_1, ..., v_{n-1}) \text{ and then reflect over } \mathbb{H}.$$

Proof. Consider the shape X and $Y := \text{flip}_{\mathbb{H}} X = RFR^{-1}X$. Since Fr is $O(n)$-equivariant, we know that

$$\text{Fr}(Y) = RFR^{-1}\,\text{Fr}(X).$$

Let A_X and A_Y be the matrices associated to X and Y respectively from the circle method.

Write

$$A_X = [v_1, ..., v_{n-1}, v_n]^{-1}.$$

Then it follows that

$$A_Y = [RFR^{-1}v_1, ..., RFR^{-1}v_{n-1}, -RFR^{-1}v_n]^{-1}.$$

The composition $A_Y^{-1} \circ A_X$ is then

$$RFR^{-1} \circ [v_1, ..., v_{n-1}, -v_n] \circ [v_1, ..., v_{n-1}, v_n]^{-1}.$$

The matrix RFR^{-1} is reflection across $\mathbb{H}$ and the matrix $v_1, ..., v_{n-1}, -v_n] \circ [v_1, ..., v_{n-1}, v_n]^{-1}$ is reflection across $\text{span}(v_1, ..., v_{n-1})$. □

Remark B.4. However this expression $A_Y^{-1} \circ A_X$ can never be equal to RFR^{-1}. Indeed suppose for contradiction it is, then since RFR^{-1} is invertible, this would imply that $[v_1, ..., v_{n-1}, -v_n] \circ [v_1, ..., v_{n-1}, v_n]^{-1} = I$, so

$$[v_1, ..., v_{n-1}, -v_n] = ([v_1, ..., v_{n-1}, v_n])^{-1})^{-1} = [v_1, ..., v_{n-1}, v_n] \implies v_n = 0,$$

which contradicts to the fact that $[v_1, ..., v_n]$ is invertible.

We may now prove Theorem 4.6 as follows.

Proof of Theorem 4.6. Write $\text{Fr}(X) = (v_1, ..., v_{n-1})$. For each vector v looking at X, we are looking at $\text{flip}_{\mathbb{H}} \circ \text{flip}_{\text{span}(v_1,...,v_{n-1})}(v)$ by Lemma B.3 on $\text{flip}_{\mathbb{H}} X$. The latter is equivalent to looking at $\text{flip}_{\text{span}(v_1,...,v_{n-1})}(v)$ at X, by canceling out $\text{flip}_{\mathbb{H}}$ on both sides. □

B.4 Proofs in Section 4.2

In this part, we will prove Theorem 4.11. This will involve some concepts from measure theory (specifically the notion of an equidistributed sequence and Corollary C.11), which we discuss in Appendix C. For our purposes, we explain equidistribution again below more formally.

Definition B.5. Let $\mathbb{d}$ be a discretization parameter. We use $\texttt{set}(\mathbb{d})$ to denote the underlying set of points given by all the lists in $\mathbb{d}$. More formally, $\texttt{set}(\mathbb{d}) := \bigcup_{k,j}\{p_{k,j}\}$.

Definition B.6. Let $\mathbb{d}_n, n = 1, 2, ...$ be a sequence of discretization parameters. We say $\{\mathbb{d}_n\}_{n\geq 1}$ is **equidistributed** if the sampling $\texttt{set}(\mathbb{d}_n)$ is equidistributed, in the sense of Definition C.9, with respect to the normalized spherical measure on S^{n-1}.

We also record a well-known fact in point-set topology. For completeness, we write down a proof.

Lemma B.7. Let $f, g : X \to Y$ be continuous functions between topological spaces such that Y is Hausdorff. If there is a point $x \in X$ such that $f(x) \neq g(x)$, then there is an open neighborhood $U \subseteq X$ containing x such that $f(x') \neq g(x')$ for all $x' \in U$.

Proof. Let $f(x), g(x) \in Y$ be distinct points. Since Y is Hausdorff, there exists open subsets $V_1, V_2 \subseteq Y$ containing $f(x), g(x)$ respectively such that $V_1 \cap V_2 = \emptyset$. Now consider the open set $U \coloneqq f^{-1}(V_1) \cap g^{-1}(V_2)$ in X, which contains the point x. For any $x' \in U$, we see that $f(x') \in V_1$ and $g(x') \in V_2$, and hence $f(x') \neq g(x')$. □

We will now prove Theorem 4.11 as follows. Morally, Theorem 4.11 is saying that admissible Φ is a complete invariant on the rotational orbits of shapes in $\mathcal{S}$.

Proof of Theorem 4.11. Let $A_X, A_Y \in \mathrm{SO}(n)$ be the rotational matrices associated to X and Y. By Theorem 4.4, we have that

$$\mathfrak{C}(X, \Phi, \mathrm{Fr}, \mathbb{d}_m) = \mathfrak{C}(A_X \cdot X, \Phi, \mathrm{Fr}, \mathbb{d}_m) \text{ and } \mathfrak{C}(Y, \Phi, \mathrm{Fr}, \mathbb{d}_m) = \mathfrak{C}(A_Y \cdot Y, \Phi, \mathrm{Fr}, \mathbb{d}_m).$$

It then suffices to prove that

$$\mathfrak{C}(A_X \cdot X, \Phi, \mathrm{Fr}, \mathbb{d}_m) \neq \mathfrak{C}(A_Y \cdot Y, \Phi, \mathrm{Fr}, \mathbb{d}_m)$$

for m sufficiently large. We first note that $\mathrm{Fr}(A_X \cdot X) = \mathrm{Fr}(A_Y \cdot Y)$ is equal to $(e_1, ..., e_{n-1})$ of $\mathbb{R}^n$, hence the directions sampled in the circle method for both $A_X \cdot X$ and $A_Y \cdot Y$ are the same. Since X and Y are in two different rotational orbits, we know that $A_X \cdot X \neq A_Y \cdot Y$. Since Φ is admissible, there is at least one vector $v \in S^{n-1}$ such that $\Phi(A_X \cdot X)(\nu) \neq \Phi(A_Y \cdot Y)(\nu)$. The admissibility of Φ and Lemma B.7 then implies that there is an open neighborhood $U \subseteq S^{n-1}$ containing ν such that,

$$\Phi(A_X \cdot X)(\omega) \neq \Phi(A_Y \cdot Y)(\omega), \text{ for all } \omega \in U.$$

Since U is non-empty and open, its value under the normalized spherical measure is positive. The theorem now follows from Corollary C.11, as the sequence $\{\mathbb{d}_n\}_{n\geq 1}$ is equidistributed. □

Remark B.8. The N chosen in Theorem 4.11 depends on the input X and Y. In general, when $\Phi = \mathrm{PHT}$, this N need not be uniformly bounded. The reasoning is similar to why the usual PHT also need not have a uniform bound on the number of directions needed to distinguish all shapes, as explained in Curry et al. (2022). However, if one impose certain constraints onto $\mathcal{S}$, then N can admit a uniform bound independent of the input shapes. The approach taken in Curry et al. (2022) is, roughly, to take $\mathcal{S}$ as geometrically embedded simplicial complexes such that (i) each embedded complex K is not "too flat" near any vertices and (ii) for any δ-ball of directions, no embedded complex K has "too many critical values". By imposing these certain uniform constraints on $\mathcal{S}$, they were able to obtain a uniform estimate on N. However, the directions they sampled on depended on the input shapes. Finding a uniform bound from a uniform list of directions independent of the input appears to be a challenge to the entire TDA community.

C Background on Measure Theory

In this section, we give a brief introduction to measure theory, introducing only the necessary concepts needed for the proofs and equidistribution. We refer the reader to Folland (1999) for a comprehensive coverage of the subject, and to Billingsley (1999); Klenke (2008) for probability measures specifically.

Definition C.1. Let $\mathfrak{X}$ be a topological space.

1. A σ-algebra Σ (thought of as **measurable subsets** of $\mathfrak{X}$) is a collection of subsets of $\mathfrak{X}$ such that $\mathfrak{X} \in \Sigma$ and Σ is closed under taking complements and countable unions.
2. A **measure** μ with respect to (X, Σ) is a function $\mu : \Sigma \to [0, +\infty]$ such that $\mu(\emptyset) = 0$, $\mu(E) \leq \mu(F)$ for $E \subseteq F \in \Sigma$, for any countable collection of pairwise disjoint subsets $\{E_k \in \Sigma\}_{k \geq 1}$, $\mu(\bigcup_{k \geq 1} E_k) = \sum_{k=1}^{\infty} \mu(E_k)$.
3. A measure μ is a **probability measure** if $\mu(\mathfrak{X}) = 1$.
4. A **measure space** is the space $(\mathfrak{X}, \Sigma, \mu)$ of a topological space $\mathfrak{X}$, a choice of σ-algebra Σ, and a compatible measure μ.
5. A measure space $(\mathfrak{X}, \Sigma, \mu)$ is a **probability measure space** if μ is a probability measure.
6. A function $f : (X, \Sigma_X) \to (Y, \Sigma_Y)$ is **measurable** if $f^{-1}(S) \in \Sigma_X$ for all $S \in \Sigma_Y$.

It is often reasonable to require all open sets in $\mathfrak{X}$ to be measurable, and this motivates the following definition.

Definition C.2. Let $\mathfrak{X}$ be a topological space. The **Borel σ-algebra** $\mathcal{B}_{\mathfrak{X}}$ is the smallest σ-algebra containing all open subsets of $\mathfrak{X}$. Continuous functions between measure spaces with Borel σ-algebras are measurable.

From now on, for simplicity we take Σ to be the Borel σ-algebra and omit the notation. A great benefit of measure theory is it allows us to formalize the theory of integrals in different settings.

The theory of integrals is built using what are called simple functions.

Definition C.3. Let $f : \mathfrak{X} \to \mathbb{R}$ be a measurable function. We say f is **simple** if it is of the form

$$f = \sum_{i=1}^{n} a_i \mathbb{1}_{A_i}, a_i \in \mathbb{R}, A_i \in \mathcal{B}_{\mathfrak{X}},$$

and $\mathbb{1}_{A_i}$ is the indicator function on A_i that sends $x \in \mathfrak{X}$ to 1 if $x \in A_i$ and 0 otherwise.

Definition C.4. Given a measure space $(\mathfrak{X}, \mathcal{B}_{\mathfrak{X}}, \mu)$, we define its theory of **integrals** as follows. Let $f : \mathfrak{X} \to \mathbb{R}$ be measurable.

1. Let $f = \sum_{i=1}^{n} a_i \mathbb{1}_{A_i}$ be simple, then
$$\int_{\mathfrak{X}} f d\mu \coloneqq \sum_{i=1}^{n} a_i \mu(A_i).$$
2. Suppose furthermore $f \geq 0$, we define
$$\int_{\mathfrak{X}} f d\mu \coloneqq \sup\{\int_{\mathfrak{X}} g d\mu : g : \mathfrak{X} \to \mathbb{R}_{\geq 0} \text{ is simple and } g \leq f\}.$$
3. For a general measurable function $f : \mathfrak{X} \to \mathbb{R}$, we write
$$f = f^+ - f^-$$
where $f^+ = \max(0, f)$ and $f^- = \max(0, -f)$. We then define
$$\int_{\mathfrak{X}} f d\mu \coloneqq \int_{\mathfrak{X}} f^+ d\mu - \int_{\mathfrak{X}} f^- d\mu.$$

Example C.5. Here are some examples of measures and their respective theory of integrals.

1. Let $x \in \mathfrak{X}$, the **Dirac measure** is a measure on $\mathfrak{X}$ defined by,
$$\delta_x(S) \coloneqq \begin{cases} 1, x \in S \\ 0, x \notin S \end{cases}.$$
This has the property that
$$\int_{\mathfrak{X}} f d\mu_x = f(x).$$
2. The standard **Lebesgue measure** on $\mathbb{R}^n$ (see Section 1.5 and 2.6 of Folland (1999)) recovers the usual theory of integrals.
3. The unit sphere $S^{n-1} \subseteq \mathbb{R}^n$ is equipped with the **spherical measure** (see Section 2.7 of Folland (1999)) and recovers the usual spherical integrals. We can normalize the spherical measure to obtain a probability measure.

Definition C.6. A sequence of probability measures μ_n for $n = 1, 2, ...$ **weakly converges** (or **converges weakly**) to a probability measure μ if for all bounded continuous functions $f : \mathfrak{X} \to \mathbb{R}$,
$$\lim_{n\to\infty} \int_{\mathfrak{X}} f d\mu_n = \int_{\mathfrak{X}} f d\mu.$$

Theorem C.7 (Portmanteau Theorem, see Theorem 2.1 of Billingsley (1999))**.** A sequence of probability measures $\{\mu_n\}_{n\geq 1}$ weakly converges to μ if and only if for any open set $U \subseteq \mathfrak{X}$,
$$\liminf_n \mu_n(U) \geq \mu(U).$$

Definition C.8. A **sampling** P of $(\mathfrak{X}, \mu)$ is a function
$$P : \{1, 2, \ ...\} \to \{\text{Non-Empty Finite Subset of } \mathfrak{X}\}.$$
We write P_n as $P(n) \subset \mathfrak{X}$.

Definition C.9. A sampling P of $(\mathfrak{X}, \Sigma, \mu)$ is **equidistributed** if the probability measures
$$\left\{\frac{1}{|P_n|}\sum_{x\in P_n} \mu_x\right\}_{n\geq 1} \text{ converges weakly to } \mu.$$

Example C.10. The examples of interests to this work are as follows. Let $\mathfrak{X} = S^{n-1}$ with the usual spherical measure. An equidistributed sampling P of $\mathfrak{X}$ could be a way to sample points uniformly on S^{n-1} that gets finer as n increases. This is implied by the strong law of large numbers.

As a corollary of the Portmanteau theorem, we have the following.

Corollary C.11. Let P be an equidistributed sampling of $(\mathfrak{X}, \mu)$ and U be an open set such that $\mu(U) > 0$, then $P_n \cap U \neq \emptyset$ for all n sufficiently large.

Proof. By Portmanteau theorem (Theorem C.7), we have that
$$\liminf_n \frac{1}{|P_n|}\sum_{x\in P_n} \mu_x(U) = \liminf_n \frac{|U \cap P_n|}{|P_n|} \geq \mu(U) > 0.$$
This means that the infimum of $\frac{|U\cap P_m|}{|P_m|}$ for $m \geq n$ is greater than 0 as $n \to \infty$. This implies that the numerator $|U \cap P_k|$ can only be 0 for finitely many k. □

D TRANSLATIONAL INVARIANCE OF GENERAL SHAPES

Let $\mathcal{S}$ be a collection of shapes (as in Definition 2.1). When $\mathcal{S}$ is the collection of embedded finite graphs or finite simplicial complexes, we described in Section 2 a way to center the all the shapes at the origin by taking the average of the vertices and subtract the average vector.

One might wonder if it would be possible to encode translation invariance for $\mathcal{S}$ in general. We formally encode this as follows:

Definition D.1. Let $\mathcal{S}$ be a collection of shapes in $\mathbb{R}^n$ (as in Definition 2.1). We say **translational invariance can be achieved** on $\mathcal{S}$ if there exists a function (called the **center function**)

$$\text{center} : \mathcal{S} \to \mathbb{R}^n$$

such that for all $X \in \mathcal{S}$, $A \in O(n)$, and $v \in \mathbb{R}^n$, we have

$$\text{center}(A \cdot X + v) = A \cdot \text{center}(X) + v.$$

Note that whenever such **center function** exists, we can always place X at the center via the isometry

$$X \mapsto X - \text{center}(X) \in \mathcal{S}.$$

For a shape X and its translation $X + v$, we have that

$$(X + v) - \text{center}(X + v) = X + v - \text{center}(X) - v = X - \text{center}(X),$$

so they are centered to the same shape.

Example D.2. The centers in Section 2 are valid center functions.

1. If $\mathcal{S}$ is the collection of embedded finite (non-empty) graphs in $\mathbb{R}^n$, then
$$\text{center}(G) \coloneqq \frac{1}{|V(G)|} \sum_{v \in V(G)} v$$
is a center function.
2. If $\mathcal{S}$ is the collection of embedded finite (non-empty) simplicial complexes in $\mathbb{R}^n$, then
$$\text{center}(K) \coloneqq \frac{1}{|V(K)|} \sum_{v \in V(K)} v$$
is a center function.

It is less obvious whether such center functions could exist for an arbitrary collection of shapes. An idea one might have is to use the integral definition of centroids, but this is not well-defined if the shape has zero volume (i.e., embedded graphs).

The purpose of this appendix section is to prove the following.

Proposition D.3. Translational invariance can be achieved on any collection of shapes defined as in Definition 2.1 (i.e., a collection of non-empty compact subsets closed under Euclidean isometries).

We prove Proposition D.3 using the concept of Chebyshev centers (see Amir (1986) for example). The conclusion should not be surprising to experts on inner product spaces.

Definition D.4. Let X be a compact non-empty subset of $\mathbb{R}^n$, the **Chebyshev center** of X is the center of the smallest closed ball in $\mathbb{R}^n$ that contains X. Formally, it is denoted:

$$\text{cc}(X) \coloneqq \arg\min_{y \in \mathbb{R}^n} \sup_{x \in X} ||x - y||^2.$$

The following lemma is well-known, and we include its proof for completeness.

Lemma D.5. Let X be a compact non-empty subset of $\mathbb{R}^n$. The Chebyshev center $\text{cc}(X)$ is well-defined and unique. Furthermore, for any $v \in \mathbb{R}^n$ and $A \in O(n)$,

$$\text{cc}(A \cdot X + v) = A \cdot \text{cc}(X) + v.$$

Proof. The well-definedness follows from X being compact in $\mathbb{R}^n$, as X is in particular closed and bounded. Now suppose y_1 and y_2 are two vectors that both minimize $\sup_{x\in X} ||x-y||^2$, with value L, and $y_1 \neq y_2$. Now consider the midpoint $m = \frac{y_1+y_2}{2}$, then for any $x \in X$, we have that

$$||x-m||^2 = ||x-(y_1+y_2)/2||^2 = \frac{1}{4}||2x-y_1-y_2||^2 = \frac{1}{4}||(x-y_1)+(x-y_2)||^2$$

$$= \frac{1}{4}(||x-y_1||^2 + 2\langle x-y_1, x-y_2\rangle + ||x-y_2||^2).$$

On the other hand, we know that

$$\frac{1}{4}||(x-y_1)-(x-y_2)||^2 = \frac{1}{4}||y_1-y_2||^2$$

and is also equal to

$$\frac{1}{4}(||x-y_1||^2 - 2\langle x-y_1, x-y_2\rangle + ||x-y_2||^2).$$

Thus, we have that

$$||x-m||^2 + \frac{1}{4}||y_1-y_2||^2 = \frac{1}{2}(||x-y_1||^2 + ||x-y_2||^2).$$

Taking the supremum over $x \in X$ shows that

$$\sup_{x\in X} ||x-m||^2 \leq \frac{1}{2}(L+L) - \frac{1}{4}||y_1-y_2||^2 = L - \frac{1}{4}||y_1-y_2||^2.$$

Since $y_1 \neq y_2$, we conclude that m is a smaller minimizer than y_1 or y_2, which is a contradiction. Thus, we see that $\mathrm{cc}(X)$ is unique.

Finally, the last statement follows from the observation that

$$\mathrm{cc}(A\cdot X+v) = \operatorname*{arg\,min}_{y\in\mathbb{R}^n} \sup_{x\in X} ||(A\cdot x+v)-y||^2 = \operatorname*{arg\,min}_{y=A\cdot w+\nu\in\mathbb{R}^n} \sup_{x\in X} ||(A\cdot x+v)-((A\cdot w)+v)||^2$$

$$= \operatorname*{arg\,min}_{y=A\cdot w+\nu\in\mathbb{R}^n} \sup_{x\in X} ||A\cdot x+v-(A\cdot w)-v||^2$$

$$= \operatorname*{arg\,min}_{y=A\cdot w+\nu\in\mathbb{R}^n} \sup_{x\in X} ||A(x-w)||^2$$

$$= \operatorname*{arg\,min}_{y=A\cdot w+\nu\in\mathbb{R}^n} \sup_{x\in X} ||x-w||^2$$

$$= A\cdot cc(X)+v.$$

□

Now we will prove Proposition D.3.

Proof of Proposition D.3. Let $\mathcal{S}^{\max}$ be the collection of all non-empty compact subsets in $\mathbb{R}^n$. Note that any collection of shapes $\mathcal{S}$ defined in Definition 2.1 must be a sub-collection of $\mathcal{S}^{\max}$. It suffices to prove a center exists for $\mathcal{S}^{\max}$, as we can restrict to $\mathcal{S}$ and $\mathcal{S}$ is closed under Euclidean isometries. Now we take cc, the Chebyshev center, as the candidate center function on $\mathcal{S}^{\max}$. The conclusion of Proposition D.3 now follows from Lemma D.5. □

E Implementation of the Circle Method

The circle method of Section 4 requires a frame function Fr and a discretization parameter $\mathbb{d}$. Here we describe the choices used for geometric molecular graphs in $\mathbb{R}^3$, with vertices $x_1, \dots, x_N$ at the atoms and edges E along the bonds. The construction is the same for Circ-PHT, Circ-ECT, and Circ-BCT; only the finite representation of the directional transform changes. The scan resolutions used in each experiment are reported there.

Directional transforms of a geometric graph. We center each shape at its centroid and rescale it so that the root-mean-square distance of its vertices from the origin is one. For a direction ν, the vertex x_i enters the filtration at height $\nu \cdot x_i$, and an edge enters once both endpoints have appeared. We compute persistence by sorting the filtration values and running a union–find pass, with cost $O(N \log N + |E|, \alpha(N))$ per direction.

E.1 Constructing the frame

We take as candidates the directions from the centroid to the vertices,

$$\varphi(X) \coloneqq \left\{ \frac{x_i}{\|x_i\|} : x_i \neq 0 \right\}.$$

Since $\varphi(g \cdot X) = g \cdot \varphi(X)$ for every $g \in O(3)$, the candidate set moves equivariantly with the shape.

Let $\psi_X(v) \in \mathbb{R}^d$ be a fixed finite summary of $\mathrm{PHT}(X)(v)$, given by persistence landscapes of the H_0 and H_1 diagrams sampled on a fixed grid. We draw $W_a, W_b \in \mathbb{R}^d$ once from a standard normal distribution and keep them fixed for all shapes. The first direction is $v_1 \coloneqq \arg\max_{v \in \varphi(X)} \langle W_a, \psi_X(v) \rangle$.

We use $\nparallel$ to denote not parallel. For the second direction, each $w \in \varphi(X)$ with $w \nparallel v_1$ is orthogonalized against v_1,

$$u_w \coloneqq \frac{w - \langle w, v_1 \rangle v_1}{\|w - \langle w, v_1 \rangle v_1\|}.$$

We then take $w_* \coloneqq \arg\max_{w \in \varphi(X), , w \nparallel v_1} \langle W_b, \psi_X(u_w) \rangle$, set $v_2 \coloneqq u_{w_*}$, and define $\mathrm{Fr}(X) \coloneqq (v_1, v_2)$.

The motivation for the random linear scores is that scalar summaries such as total persistence can assign the same value to different directional diagrams. For a finite collection of distinct vectors $\psi_X(v)$, a continuously sampled weight vector gives distinct scores with probability one. The weights are sampled once and are not trained.

Computing v_1 uses at most one directional PHT per candidate and computing v_2 at most one more, for approximately $2N$ directional computations.

E.2 Placing the circles

In the canonical coordinates of Definition 4.2, we use K circles of constant first coordinate,

$$C_k \coloneqq \{ z_k e_1 + \rho_k (\cos\theta \, e_2 + \sin\theta \, e_3) : \theta \in [0, 2\pi) \},$$

where $z_k \coloneqq -1 + (2k+1)/K$ and $\rho_k \coloneqq \sqrt{1 - z_k^2}$, for $k = 0, \dots, K-1$. Each circle is oriented by increasing θ and sampled at m equally spaced angles $\theta_j = 2\pi j / m$.

In these coordinates, the reflection in Theorem 4.6 is $\sigma_3 = \mathrm{diag}(1, 1, -1)$. On each C_k, it sends θ to $-\theta$, and hence $p_{k,j}$ to $p_{k,-j \bmod m}$. Thus the circle is preserved and its cyclic order is reversed.

The points z_k are the midpoints of K equal intervals in $[-1, 1]$, giving equal-area bands on the sphere. As K and m increase, the sampled points form an equidistributed refinement of S^2, as

required by Theorem 4.11. For $K = 1$, this is the great circle orthogonal to e_1; three sampled points already give a non-trivial reversal, as discussed in Appendix F.2.

E.3 From the list of transforms to a vector

The circle method returns the ordered values $\Phi(X)(w_{k,j})$. For Circ-PHT, we use the same persistence-landscape summary ψ as above. For Circ-ECT, we sample the Euler characteristic curve on a fixed grid, and for Circ-BCT, the two Betti curves on the same grid. The resulting vectors are concatenated in scan order, with no pooling across directions.

If the scan uses $D = Km$ directions, the full construction requires about $2N + D$ directional computations. The total cost is $O((2N+D)(N \log N + |E|, \alpha(N)))$, or $O(N^2 \log N + DN \log N)$ for molecular graphs with $|E| = O(N)$.

For Circ-PHT, Circ-ECT, and Circ-BCT, the weights W_a, W_b, the directional summaries, and the discretization parameters are fixed across all shapes and are not trained. Circ-PHT-L uses the same frame and circular scan, but replaces the fixed persistence-landscape readout with a learned DeepSets (Zaheer et al., 2017) readout, which can be trained when labels from a downstream task are available.

F Experimental Setup

F.1 Molecular retrieval

Datasets and filtering. We use QM9-OR together with the BZR, COX2, and DHFR molecular graph datasets. The retrieval experiment is restricted to geometrically chiral base geometries. Since Table 1 withholds atom identities, a geometry $X = (G, x)$ with unlabelled bond graph G counts as achiral whenever its mirror image can be superimposed on it by a proper rigid motion after relabelling the vertices by a graph automorphism (i.e., a bijective graph homomorphism from a graph to itself). We therefore compute

$$\varepsilon(X) := \min_{\pi \in \mathrm{Aut}(G)} \min_{R \in SO(3)} \frac{\mathrm{RMSD}\big(x,\ R\,Q\,\pi x\big)}{\mathrm{rms}(x)}, \qquad Q = \mathrm{diag}(1, 1, -1),$$

where both point sets are centred, πx permutes the vertex coordinates by π, the root-mean-square deviation is taken over the fixed vertex correspondence, $\mathrm{rms}(x)$ is the root-mean-square distance of the vertices from the centroid, and the inner minimum is the Kabsch alignment restricted to proper rotations ($\det R = +1$), so that a reflection can never perform the alignment. A base geometry is kept if $\varepsilon(X) > \tau$ with $\tau = 10^{-3}$; the search over automorphisms stops as soon as one drives the residual below τ. After filtering, the experiments retain 355 of the 405 BZR, 467 of the 467 COX2, and 750 of the 756 DHFR base molecules. For QM9-OR, the candidates are the molecules with a stereocentre and at least nine atoms, scanned in a fixed random order until 1,200 pass the filter; these 1,200 are used for the full baseline comparison in Table 1.

Filtering and dataset splitting are performed at the level of the base molecule, before reflection or rotation augmentation. We use five folds grouped by base molecule. Consequently, no reflected copy or rotated pose of a molecule assigned to the test fold occurs in the corresponding training or validation folds. Table 1 reports fold 0. Table 5 reports mAP, P@1, and R@7 for every method on fold 0; Table 1 in the main text shows its mAP column for a subset of the rows.

Reflection and rotation augmentation. For each retained base geometry X, we construct a reflected geometry QX, where $Q \in O(3)$ has determinant -1. We then sample eight properly rotated poses of each hand using elements of $SO(3)$. Thus each base molecule contributes 16 poses: eight from X and eight from QX. The molecular adjacency graph is preserved by these transformations.

All methods in Table 1 receive the same geometry-only information: the unlabelled adjacency graph and Cartesian coordinates. Atom identities, atomic numbers, bond types or orders, formal charges, and stereochemical annotations are withheld. For the directional transforms, coordinates are translated to the common centering convention before the directional filtration is evaluated.

Learned baselines. GIN-Topo and GIN-XYZ use the GIN architecture (Xu et al., 2019). We additionally compare against EGNN (Satorras et al., 2021), PaiNN (Schütt et al., 2021), GemNet-Q (Gasteiger et al., 2021), GotenNet (Aykent & Xia, 2025), and EquiformerV2 (Liao et al., 2024). Neural k-Forms (Maggs et al., 2024), PersLay (Carrière et al., 2020), and DECT (Roell & Rieck, 2024) provide topological and geometric-topological comparisons, and ChiENN (Gaiński et al., 2023) is a chirality-oriented geometric architecture.

Retrieval protocol. Every pose is used once as a query against all remaining test poses. If $\mathcal{P}(q)$ denotes the seven remaining poses with the same base identity and hand as query q, then a gallery item is relevant if and only if it belongs to $\mathcal{P}(q)$. The eight poses of the opposite hand of the same molecule and all poses of all other molecules are non-relevant.

Hence each query has exactly seven relevant gallery items. This definition deliberately treats the reflected hand as non-relevant: the task asks whether a representation removes proper rotations without identifying the two reflected geometries.

Metrics. Mean average precision summarizes the complete ranking,

$$\mathrm{AP}(q) = \frac{1}{|\mathcal{P}(q)|} \sum_{r=1}^{|\mathcal{G}|-1} \mathrm{Prec}_q(r)\, \mathbf{1}\{\text{item at rank } r \in \mathcal{P}(q)\},$$

and

$$\mathrm{mAP} = \frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}} \mathrm{AP}(q).$$

Precision at rank one is

$$P@1 = \frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}} \mathbf{1}\{\text{top-ranked item is relevant}\}.$$

Finally,

$$R@7 = \frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}} \frac{|\mathcal{P}(q) \cap \mathrm{Top}_7(q)|}{7}.$$

Because every query has seven relevant items, $P@7 = R@7$ and we report only $R@7$.

Reflection-blind reference. Consider an idealized representation that is invariant to the sampled proper rotations and otherwise identifies the base molecule, but assigns the same descriptor to the two reflected hands. For a query, the seven relevant poses are then tied with the eight poses of the opposite hand. The tied block therefore contains $N = 15$ items, of which $R = 7$ are relevant.

Under a uniform ordering of this tied block,

$$\mathbb{E}[P@1] = \frac{7}{15}.$$

Likewise, the expected number of relevant items among the first seven ranks is $7(7/15)$, and hence

$$\mathbb{E}[R@7] = \frac{7}{15}.$$

For average precision, condition on rank r containing a relevant item. The expected number of relevant items among the first r ranks is

$$1 + (r-1)\frac{R-1}{N-1} = 1 + (r-1)\frac{6}{14}.$$

Since the probability that rank r itself is relevant is $7/15$,

$$\begin{aligned}
\mathbb{E}[\mathrm{AP}] &= \frac{1}{7} \sum_{r=1}^{15} \frac{7}{15} \frac{1 + (r-1)6/14}{r} \\
&= \frac{1}{15} \sum_{r=1}^{15} \left(\frac{3}{7} + \frac{4}{7r} \right) \\
&= \frac{3}{7} + \frac{4}{105} H_{15} \approx 0.55498,
\end{aligned}$$

where H_{15} is the fifteenth harmonic number. We report this value as 0.555.

The reference assumes that the representation retains enough information to identify the base molecule and loses only the distinction between its two hands. A coarser representation that also loses molecular identity falls below this value.

Directional baselines and circular descriptors. Lab-PHT, Lab-ECT, and Lab-BCT evaluate their directional transform at fixed laboratory-frame directions. They therefore retain pose dependence. Avg-PHT first averages the raw height function over $SO(3)$; this average is zero, so the directional filtration collapses. Canon-PHT instead chooses directions $O(3)$-equivariantly and evaluates the transform only at those directions. This removes proper-rotation dependence but also identifies reflected inputs.

The CIRC-* descriptors use the proper-frame construction and retain the order of the resulting directional scan. CIRC-PHT uses persistence diagrams followed by the fixed persistence readout used throughout the controlled experiments. CIRC-ECT and CIRC-BCT use the corresponding directional Euler and Betti curves. CIRC-PHT-L uses the same proper-frame PHT construction but replaces the fixed persistence readout with a learned directional readout. Tables 1 and 5 use the two-circle scan with 44 directions per circle, fixed in advance; the dependence on the resolution of a single circle is examined in Table 6.

ChiENN reads handedness through the circular ordering of the neighbours of an atom. Under the geometry-only retrieval protocol, however, all initial node features are constant. A graph-level reflection signal can therefore arise only when message passing makes the substituents around a tetrahedral centre distinguishable from the unlabelled graph structure.

In particular, if two substituents remain structurally identical, their reflection-sensitive contributions exchange under reflection and cancel under the final sum pooling. A tetrahedral centre with four graph-distinguishable substituents occurs in 1,148 of the 1,200 QM9-OR molecules, in 95 of 355 for BZR, in 4 of 467 for COX2, and in 86 of 750 for DHFR. This accounts for the strong dependence of the ChiENN retrieval result on the underlying molecular graph.

F.2 DIRECTION AND READOUT ABLATIONS

Minimum order-sensitive scan. For a single circular scan with D equally spaced directions, index the samples by $k = 0, \dots, D-1$. Reflection reverses the circular order by

$$k \longmapsto -k \pmod{D}.$$

For $D = 1$ this action is trivially the identity. For $D = 2$ the two indices are also fixed modulo 2. Hence a one- or two-direction circular scan cannot carry information through reversal of its ordering. The first non-trivial permutation occurs at $D = 3$.

Table 6 varies D from 1 to 88 on a single circle and also reports two-circle scans. For CIRC-PHT, $D = 1, 2$ lie at the reflection-blind level and $D = 3$ gives mAP 1.000 on all four datasets. CIRC-ECT and CIRC-BCT also increase sharply at $D = 3$, although a very coarse scan can additionally lose base-molecule information. They reach 1.000 on COX2 from $D = 3$, on BZR at $D = 8$ and from $D = 16$ on (with 0.994 at $D = 12$), and on DHFR from $D = 32$.

For QM9-OR, Table 6 uses the same 1,200-molecule selection as Table 1 (Appendix F.1); its fold 0 contains 240 base molecules and 3,840 reflected and rotated poses. A separate, earlier full-scale run applied the same chirality filter to all 85,239 eligible records, of which 80,811 passed; its fold 0 contains 16,163 base molecules and 258,608 poses. On that cohort, CIRC-PHT reaches mAP 1.000 at $D = 6$ (0.99999 at $D = 4$), while CIRC-ECT and CIRC-BCT obtain $D = 1$ mAP of 0.213 and 0.239, respectively, below the 0.555 reflection-blind reference: a one-direction scan also loses part of the base-molecule identity. These full-scale values are not part of Table 6.

Arrangement control. To distinguish the number of sampled directions from their ordering, we also evaluate four directions arranged as two independent two-point circles. Each two-point circle is invariant under reversal, so the combined four-direction descriptor is still reflection-blind despite using more directions than the minimal three-point scan. This control separates the effect of directional count from the order structure used by the circle method. Table 7 reports it next to the one-circle scan with the same four directions.

Table 5: **Geometry-only handedness retrieval** (mAP / P@1 / R@7; higher is better; fold 0). Each base molecule contributes two reflected hands and eight rotated poses per hand; for each query, the seven other same-hand poses are relevant. All methods receive only the unlabelled graph and 3-D coordinates. A rotation-invariant, reflection-blind representation ties the seven positives with the eight reflected poses, giving expected $\mathrm{mAP} = 0.555$ and $\mathrm{P@1} = \mathrm{R@7} = 7/15$ under uniform tie resolution. Lab-PHT/ECT/BCT use fixed laboratory directions and are pose-dependent. Canon-PHT uses $O(3)$-equivariant canonical directions and is reflection-blind. Avg-PHT uses the rotation-averaged height construction; the averaged height is constant, removing its directional geometric signal. CIRC-PHT/ECT/BCT instead use the proposed proper frame and retain reflection sensitivity. Every 1.000 is **bolded**.

	QM9-OR			BZR			COX2			DHFR		
Method	mAP	P@1	R@7	mAP	P@1	R@7	mAP	P@1	R@7	mAP	P@1	R@7
Unlearned												
VR-PH	0.554	0.465	0.466	0.554	0.468	0.463	0.558	0.472	0.468	0.553	0.471	0.464
Lab-PHT	0.056	0.138	0.058	0.046	0.113	0.048	0.019	0.037	0.020	0.033	0.084	0.036
Lab-ECT	0.043	0.104	0.042	0.042	0.097	0.040	0.016	0.033	0.013	0.029	0.072	0.030
Lab-BCT	0.051	0.130	0.052	0.045	0.111	0.042	0.017	0.031	0.014	0.034	0.083	0.035
Canon-PHT	0.551	0.458	0.462	0.552	0.462	0.464	0.552	0.461	0.463	0.556	0.462	0.468
Avg-PHT	0.004	0.002	0.002	0.012	0.003	0.006	0.009	0.002	0.004	0.006	0.003	0.003
Learned												
GIN-Topo	0.552	0.461	0.464	0.541	0.453	0.453	0.534	0.446	0.446	0.545	0.457	0.457
GIN-XYZ	0.549	0.464	0.462	0.529	0.451	0.444	0.494	0.467	0.403	0.541	0.440	0.456
EGNN	0.557	0.467	0.469	0.548	0.446	0.463	0.545	0.438	0.460	0.552	0.460	0.462
PaiNN	0.552	0.445	0.467	0.550	0.460	0.462	0.551	0.459	0.463	0.554	0.447	0.467
GemNet-Q	0.554	0.464	0.465	0.552	0.447	0.468	0.552	0.460	0.463	0.554	0.472	0.466
GotenNet	0.557	0.466	0.468	0.546	0.424	0.463	0.558	0.463	0.469	0.558	0.488	0.467
EquiformerV2	**1.000**	**1.000**	**1.000**	0.998	0.999	0.996	**1.000**	**1.000**	**1.000**	0.999	0.999	0.998
Neural k-Forms	0.534	0.462	0.450	0.333	0.369	0.272	0.214	0.253	0.178	0.422	0.417	0.352
PersLay	0.555	0.456	0.468	0.560	0.474	0.471	0.559	0.478	0.470	0.555	0.463	0.469
DECT	0.547	0.455	0.462	0.503	0.438	0.424	0.486	0.445	0.411	0.541	0.466	0.458
ChiENN	0.952	0.942	0.942	0.578	0.511	0.511	0.467	0.386	0.386	0.538	0.457	0.457
Ours												
CIRC-BCT	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
CIRC-ECT	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
CIRC-PHT	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
CIRC-PHT-L	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**

Readout resolution. We additionally vary the persistence-landscape depth, filtration-grid resolution, and per-direction readout while keeping the circular arrangement fixed (Table 7). These changes alter the number of scalar coordinates used to represent each directional transform, but do not change the reflection action on the ordered scan. Across the tested configurations, the transition between the reflection-blind $D \leq 2$ regime and the order-sensitive $D \geq 3$ regime is unchanged. The smallest exactly-perfect CIRC-PHT configuration measured in this ablation uses $D = 3$ with one persistence-landscape level, for a total descriptor size of 120 scalars.

F.3 LONG-RANGE CHIRALITY

Task. We use the noise-free long-range chirality task from Chi-Geometry (Weaver & Pasini, 2025). The task contains a labelled centre whose R/S class depends on species placed at a controlled graph distance d. We evaluate

$$d \in \{1, 2, \ldots, 9\}.$$

Increasing d therefore moves the information required to determine the label farther from the central site. Table 8 reports every method at every distance; Table 3 in the main text shows a subset of its rows.

Table 6: **Minimum order-sensitive scan, full sweep.** Retrieval mAP (fold 0) as the number of directions D on one circle is varied; the last column is the two-circle scan with 44 directions per circle used in Table 1. CIRC-* rows use the proposed frame; Lab-* rows read the same transform in fixed laboratory directions. Reflection of a single circular scan acts by $k \mapsto -k \bmod D$, the identity for $D \leq 2$. All four datasets use the test poses of Table 1 (fold 0; for QM9-OR the 1,200-molecule selection of Appendix F.1). Every 1.000 is **bolded**.

		Scan directions D											
Data	**Readout**	1	2	3	4	6	8	12	16	24	32	44	88
BZR	CIRC-PHT	0.553	0.557	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
	Lab-PHT	0.039	0.052	0.041	0.044	0.041	0.043	0.043	0.043	0.043	0.043	0.043	0.046
	CIRC-ECT	0.548	0.555	0.994	0.975	0.994	**1.000**	0.994	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
	Lab-ECT	0.024	0.034	0.031	0.031	0.035	0.035	0.037	0.037	0.038	0.038	0.038	0.042
	CIRC-BCT	0.548	0.555	0.994	0.975	0.994	**1.000**	0.994	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
	Lab-BCT	0.029	0.040	0.034	0.036	0.039	0.040	0.042	0.042	0.043	0.043	0.043	0.045
COX2	CIRC-PHT	0.555	0.552	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
	Lab-PHT	0.021	0.027	0.019	0.022	0.019	0.020	0.019	0.019	0.019	0.019	0.019	0.019
	CIRC-ECT	0.517	0.539	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
	Lab-ECT	0.014	0.017	0.014	0.015	0.015	0.015	0.015	0.015	0.015	0.015	0.015	0.016
	CIRC-BCT	0.517	0.539	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
	Lab-BCT	0.015	0.019	0.016	0.016	0.016	0.015	0.016	0.016	0.016	0.016	0.016	0.017
DHFR	CIRC-PHT	0.553	0.551	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
	Lab-PHT	0.026	0.040	0.025	0.032	0.031	0.030	0.030	0.030	0.030	0.030	0.030	0.033
	CIRC-ECT	0.555	0.555	0.994	0.991	0.994	0.994	0.994	0.997	0.997	**1.000**	**1.000**	**1.000**
	Lab-ECT	0.015	0.023	0.018	0.022	0.023	0.023	0.024	0.024	0.025	0.025	0.025	0.029
	CIRC-BCT	0.555	0.555	0.994	0.991	0.994	0.994	0.994	0.997	0.997	**1.000**	**1.000**	**1.000**
	Lab-BCT	0.019	0.030	0.023	0.027	0.029	0.028	0.029	0.029	0.029	0.030	0.030	0.034
QM9-OR	CIRC-PHT	0.555	0.552	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
	Lab-PHT	0.037	0.052	0.043	0.050	0.051	0.052	0.053	0.052	0.053	0.052	0.052	0.056
	CIRC-ECT	0.534	0.555	**1.000**	0.994	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
	Lab-ECT	0.023	0.029	0.027	0.031	0.032	0.034	0.035	0.037	0.038	0.039	0.039	0.043
	CIRC-BCT	0.534	0.555	**1.000**	0.994	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
	Lab-BCT	0.026	0.035	0.032	0.038	0.039	0.041	0.043	0.045	0.046	0.047	0.048	0.051

The noise-free geometry by itself does not determine the class: the species occupying the distant positions are part of the task. This experiment is therefore different from the geometry-only retrieval experiment of Section 5.2.

Species encoding for the circular descriptors. For the circular descriptors, species information is incorporated by a fixed radial lift. If $x_i \in \mathbb{R}^3$ is the coordinate of atom i, Z_i its atomic number, and $\bar{x}$ the molecular centroid, we use

$$x_i' = (x_i - \bar{x}) \left(1 + 0.2 \, \frac{Z_i}{118}\right).$$

The map is fixed before training and is used for every distance d. It introduces no learned parameter. The directional transforms are computed from the lifted geometry $X' = \{x_i'\}$.

Circular readouts. The main long-range experiment uses the same ordered 88-direction scan for CIRC-PHT, CIRC-ECT, and CIRC-BCT. The descriptors are computed once and then frozen.

Refining the ECT/BCT scan. The Euler and Betti curves count features at the 20 grid levels, so they are integer vectors, and the species lift moves an atom by at most 0.03–0.06 against a grid spacing of 0.27. On the achiral templates of this benchmark, a lifted molecule and its mirror image therefore often produce the same integer descriptor along every direction of the 88-direction scan, and such a molecule cannot be classified by any read-out of that descriptor. At $D = 88$, 19%–61% of the test molecules are tied with their mirror image in this way, except at $d = 3$ and $d = 7$, where the benchmark geometry is itself chiral and no ties occur. A finer scan gives more directions along which a lifted atom can cross a grid level. The four scans in Table 8 are different direction sets rather than nested refinements, so the tied fraction is not monotone in D, but it falls from 19%–61% at $D = 88$ to 0%–14% at $D = 704$, and the CIRC-ECT/CIRC-BCT accuracy rises from 0.682–0.904 to 0.923–1.000 on these seven distances, consistent with Theorem 4.11. The persistence-landscape read-out uses the exact birth and death heights rather than grid counts, so CIRC-PHT never ties and is exact at every $D \geq 3$; the gap between the rows is a property of the coarse read-out, not of the circle method.

Table 7: **Read-out and arrangement ablation** (mAP, fold 0). Three or four directions on one circle, the same four directions on two two-point circles (each two-point circle is fixed by reversal), and the 88-direction scan; "numbers" is the descriptor length. Landscapes with one level have the same length as Betti curves. Every 1.000 is **bolded**.

read-out	numbers	BZR	COX2	DHFR
3 directions, 1 circle × 3				
Euler curves	60	0.994	**1.000**	0.994
Betti curves	120	0.994	**1.000**	0.994
landscapes, K=1	120	**1.000**	**1.000**	**1.000**
landscapes, K=4, grid 5	120	**1.000**	**1.000**	**1.000**
landscapes, K=4, grid 20 (default)	480	**1.000**	**1.000**	**1.000**
landscapes, K=8	960	**1.000**	**1.000**	**1.000**
4 directions, 1 circle × 4				
Euler curves	80	0.975	**1.000**	0.991
Betti curves	160	0.975	**1.000**	0.991
landscapes, K=1	160	**1.000**	**1.000**	**1.000**
landscapes, K=4, grid 5	160	**1.000**	**1.000**	**1.000**
landscapes, K=4, grid 20 (default)	640	**1.000**	**1.000**	**1.000**
landscapes, K=8	1280	**1.000**	**1.000**	**1.000**
4 directions, 2 circles × 2 (reversal-fixed control)				
Euler curves	80	0.555	0.544	0.552
Betti curves	160	0.555	0.544	0.552
landscapes, K=1	160	0.555	0.556	0.553
landscapes, K=4, grid 5	160	0.554	0.557	0.554
landscapes, K=4, grid 20 (default)	640	0.553	0.559	0.554
landscapes, K=8	1280	0.557	0.556	0.555
88 directions, 2 circles × 44				
Euler curves	1760	**1.000**	**1.000**	**1.000**
Betti curves	3520	**1.000**	**1.000**	**1.000**
landscapes, K=1	3520	**1.000**	**1.000**	**1.000**
landscapes, K=4, grid 5	3520	**1.000**	**1.000**	**1.000**
landscapes, K=4, grid 20 (default)	14080	**1.000**	**1.000**	**1.000**
landscapes, K=8	28160	**1.000**	**1.000**	**1.000**

For the fixed CIRC-PHT, CIRC-ECT, and CIRC-BCT rows in Table 8, the R/S label is read using logistic regression fitted on the training split. CIRC-PHT-L instead uses the pre-declared two-layer learned readout on the same frozen directional representation. Hyperparameter selection is performed on the validation split; test labels are used only for the final evaluation.

The "direction order pooled" control starts from the same directional PHT blocks but averages over the direction index before classification. It therefore retains the per-direction ingredients while removing the ordering that changes under reflection. VR-PH is computed from reflection-invariant Euclidean distances and provides a second reflection-blind control.

EquiformerV2 on the supplied bond graph. For the EquiformerV2 rows on the supplied bond graph in Table 8, the network receives the benchmark's supplied bond graph. We verified the edge set used by the model against the supplied graph. In this setting, one message-passing hop corresponds to one bond, and after L layers information can propagate through at most L graph hops. Here L denotes model depth, i.e., the number of message-passing/Transformer blocks, not a spherical-harmonic degree.

On this graph, every tested depth

$$L \in \{3, 4, 5, 7\}$$

is perfect through $d = 3$ and approximately at chance from $d = 4$. Increasing depth over this range therefore does not move the measured transition.

Radius graphs. The radius rows in Table 8 change the graph itself. Rather than supplying the bond graph, we rebuild the neighbourhood graph from coordinates at radii

$$R \in \{1.2,\ 2.5,\ 5.0\}\ \text{Å}$$

Table 8: **Long-range chirality, all configurations.** Balanced handedness accuracy as graph distance d increases; chance is 0.5. The CIRC-* descriptors are frozen. L is the number of EquiformerV2 message-passing layers. The first four EquiformerV2 rows use the supplied bond graph; the radius rows rebuild the graph using the stated cutoff. The CIRC-ECT/CIRC-BCT rows with $D = 176$, 352, 704 refine the scan to $D/88$ latitude circles of 88 directions each (Appendix E); the other CIRC-* rows use the 88-direction scan of Table 3. Best value per column **bolded**.

Distance d (bonds)	1	2	3	4	5	6	7	8	9
CIRC-PHT	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
CIRC-PHT-L	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
CIRC-BCT	0.861	0.793	**1.000**	0.682	0.863	0.822	**1.000**	0.904	0.845
CIRC-ECT	0.861	0.793	**1.000**	0.682	0.863	0.822	**1.000**	0.904	0.845
CIRC-ECT, CIRC-BCT, $D = 176$	0.779	0.936	**1.000**	0.819	0.806	0.913	**1.000**	0.944	0.801
CIRC-ECT, CIRC-BCT, $D = 352$	0.977	0.917	**1.000**	0.943	0.958	0.914	**1.000**	0.955	0.959
CIRC-ECT, CIRC-BCT, $D = 704$	**1.000**	0.983	**1.000**	0.923	0.954	0.990	**1.000**	0.976	0.988
CIRC-PHT, direction order pooled	0.509	0.494	0.496	0.501	0.504	0.486	0.504	0.509	0.492
VR-PH	0.500	0.500	0.496	0.503	0.505	0.481	0.500	0.499	0.490
EquiformerV2, $L = 3$	**1.000**	**1.000**	**1.000**	0.500	0.508	0.500	0.523	0.495	0.504
EquiformerV2, $L = 4$	**1.000**	**1.000**	**1.000**	0.498	0.495	0.500	0.511	0.499	0.498
EquiformerV2, $L = 5$	**1.000**	**1.000**	**1.000**	0.500	0.512	0.503	0.503	0.496	0.487
EquiformerV2, $L = 7$	**1.000**	**1.000**	**1.000**	0.507	0.500	0.499	0.507	0.501	0.503
EquiformerV2, $L = 4$, $R = 1.2$ Å	**1.000**	**1.000**	**1.000**	0.498	0.489	0.498	0.500	0.496	0.496
EquiformerV2, $L = 4$, $R = 2.5$ Å	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	0.496	0.507	0.487
EquiformerV2, $L = 4$, $R = 5$ Å	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
ChiENN	0.507	0.502	0.482	0.509	0.484	0.509	0.489	0.522	0.496

Table 9: **EquiformerV2 capacity control on the supplied bond graph.** Balanced R/S accuracy. The graph is fixed; only network width, depth, and dropout are changed.

Configuration	Params.	$d = 2$	$d = 3$	$d = 4$	$d = 6$
128 ch., $L = 4$	8.4M	**1.000**	**1.000**	0.498	0.500
512 ch., $L = 4$	127.4M	**1.000**	0.488	0.500	0.500
512 ch., $L = 12$	376.9M	**1.000**	0.493	0.499	0.500
512 ch., $L = 12$, no dropout	376.9M	**1.000**	**1.000**	0.500	0.494

and train an $L = 4$ EquiformerV2 on the resulting graph. The $R = 1.2$ Å graph reproduces the bond-graph behaviour. At $R = 2.5$ Å, accuracy remains perfect through $d = 6$ and falls to chance at $d \geq 7$. At $R = 5$ Å, the model is perfect at every tested distance.

Table 1 already shows that the architecture can encode reflection-sensitive information; the radius experiment shows that changing which atoms can directly interact changes the distance over which the label can be recovered. The larger radius comes with a correspondingly denser graph: at $R = 2.5$ Å the rebuilt graph has $8.2\times$ the bond count at $d = 6$ and $9.4\times$ at $d = 9$; at $R = 5$ Å it has $9.5\times$ and $14\times$.

Width and depth control. We additionally enlarge EquiformerV2 while keeping the supplied bond graph fixed. Table 9 reports the measured capacity sweep. Increasing the model to 512 channels and up to 12 layers does not recover the $d \geq 4$ cases on the fixed bond graph.

ChiENN. ChiENN (Gaiński et al., 2023) is designed to read handedness from the circular ordering of the neighbours of an atom. The centres in this long-range task are three-armed. Consequently, every local neighbour cycle used by ChiENN contains at most two members. A two-element circular ordering is unchanged by reversal, so the mirrored inputs produce the same local ordering information for the model. We verified this directly at the input construction. ChiENN therefore remains at chance throughout this benchmark for a reason distinct from the finite-range effect measured for EquiformerV2.

F.4 3DCS EXTERNAL VALIDATION

Benchmark and metric. We use the chirality task from the 3DCS benchmark (Wang et al., 2026), containing 52,391 stereochemically validated conformers of chiral drug-like molecules. For a given parent molecule, the benchmark contains conformers belonging to different stereochemical configurations.

The evaluator compares distances between molecular embeddings. Pairs belonging to different stereoisomers should have larger representation distance than conformers belonging to the same stereoisomer. ES-AUC is the ROC AUC obtained from these pairwise distances; 0.5 corresponds to no separation and 1 to perfect separation. We use the released 3DCS evaluation procedure for all rows reported in Table 4.

Frozen circular descriptors. No parameter of CIRC-PHT, CIRC-ECT, or CIRC-BCT is fitted on 3DCS. The circular descriptors receive the molecular coordinates and unlabelled bond graph, but no atom identities, bond orders, stereochemical annotations, or pretrained chemical features.

We evaluate both the full 88-direction CIRC-PHT descriptor and the four-direction descriptor used in the fusion experiment. The corresponding zero-shot ES-AUC values are reported in Table 4.

Molecular baselines. We evaluate MACE (Batatia et al., 2022), E3FP (Axen et al., 2017), and UniMol (Zhou et al., 2023) through the same 3DCS evaluation procedure. The benchmark does not release the embeddings used for all of its learned baseline results, so Table 4 contains only representations for which we can compute the embeddings directly and evaluate them under the same protocol.

Combining molecular and circular representations. For the combined rows, we concatenate the molecular representation with the four-direction CIRC-PHT descriptor. A scalar relative weight on the CIRC-PHT block is selected from three pre-specified candidate values using validation performance only; the chosen value is then fixed for the final ES-AUC evaluation. The same candidate set and selection procedure are used for MACE, E3FP, and UniMol.

The purpose of this experiment is to test whether the circular geometric-topological descriptor contains information complementary to the molecular representation. MACE increases from 0.467 to 0.704, E3FP from 0.580 to 0.711, and UniMol from 0.622 to 0.714 after adding CIRC-PHT. For E3FP and UniMol, the combined embedding also exceeds the standalone four-direction CIRC-PHT value of 0.706.

Relation to the controlled chirality experiments. The reflected-copy retrieval protocol and the long-range experiment of Section 5.2 isolate the proper-rotation versus reflection symmetry by construction. 3DCS complements them: it evaluates the frozen descriptor on an external stereoisomer-separation benchmark and tests whether it provides useful information when combined with existing molecular representations.