Diss-I-ECT: Dissecting Graph Data with Local Euler Characteristic Transforms

Julius von Rohrscheidt¹² Bastian Rieck¹²³

Abstract

The Euler Characteristic Transform (ECT) is an efficiently-computable geometrical-topological invariant that characterizes the *global* shape of data. In this paper, we introduce the Local *Euler Characteristic Transform* (ℓ -ECT), a novel extension of the ECT particularly designed to enhance expressivity and interpretability in graph representation learning. Unlike traditional graph neural networks (GNNs), which may lose critical local details through aggregation, the ℓ -ECT provides a lossless representation of local neighborhoods. This approach addresses key limitations in GNNs by preserving local structures while maintaining global interpretability. Moreover, we construct a rotation-invariant metric based on *l*-ECTs for spatial alignment of data spaces. Our method exhibits superior performance compared to standard GNNs on a variety of benchmark nodeclassification tasks, while also offering theoretical guarantees that demonstrate its effectiveness.

1. Introduction

Traditionally, graph neural network (GNNs) rely on message-passing schemes to aggregate node features. While effective for many tasks, this approach often leads to the loss of critical *local* information, as the aggregation process can diffuse and obscure the original node vector representations (Topping et al., 2022). This limitation makes it challenging to preserve local characteristics that may be essential for some applications. To address this limitation, we harness the *Euler Characteristic Transform* (Turner et al., 2014, ECT), an expressive geometrical-topological invariant that can be computed efficiently. Relying only on (weighted) sums, the ECT can be computed efficiently, making it a powerful tool for representation learning (Röell & Rieck, 2024). Moreover, the ECT is known to be *invertible* for data in \mathbb{R}^n , ensuring that the original data can always be reconstructed (Curry et al., 2022; Ghrist et al., 2018). In this paper, we extend the ECT to local neighborhoods, presenting the Local Euler Characteristic *Transform* (ℓ -ECT), a novel method designed to preserve *local* structure while retaining *global* interpretability. The *l*-ECT captures both topological (i.e., structural) and geometrical (i.e., spatial) information around each data point, making it particularly advantageous for graph-based or higher-order data.¹ The ℓ -ECT thus becomes an expressive *fingerprint* of local neighborhoods, specifically addressing the challenge of neighborhood aggregation in featured graphs while ensuring the lossless representation of local node neighborhoods. We theoretically investigate how ℓ -ECTs maintain critical local details, and therefore provide a nuanced representation that can be used for downstream graph-learning tasks such as node classification. Our method is highly effective, particularly for tasks where node feature aggregation may obscure essential differences, such as in graphs with high heterophily. Additionally, the *l*-ECT framework's natural vector representation makes it compatible with a wide range of machine-learning models, facilitating both performance and interpretability.

As our main contributions, we (i) construct ℓ -ECTs in the context of embedded simplicial complexes (and graphs) and theoretically investigate their expressivity in the special case of featured graphs, (ii) empirically show that this expressivity positions ℓ -ECTs as a powerful general tool for interpretable node classification, often superior to standard GNNs, and (iii) introduce an efficiently computable rotationinvariant metric based on ℓ -ECTs that facilitates the spatial alignment of geometric graphs.

2. Background

We define our method in the most general setting, i.e., that of a simplicial complex, while also providing a brief introduction to graph neural networks.

¹Institute of AI for Health, Helmholtz Munich, Germany ²Technical University of Munich, Germany ³University of Fribourg, Switzerland. Correspondence to: Julius von Rohrscheidt <julius.rohrscheidt@helmholtz-munich.de>, Bastian Rieck <bastian.grossenbacher@unifr.ch>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

¹Our experiments deliberately focus on graphs, but we note that the method can be extended to novel higher-order datasets based on *simplicial complexes*, for instance (Ballester et al., 2025).

Simplicial Complexes A simplicial complex K is a mathematical structure that generalizes graphs to model higher-order (non-dyadic) relationships and interactions. While graphs model pairwise (dyadic) connections between entities using nodes and edges, simplicial complexes extend this representation to higher dimensions by including triangles (2-simplices), tetrahedra (3-simplices), and their higher-dimensional analogues. Let $v_0, \ldots, v_k \in \mathbb{R}^n$ be affinely independent points. The (geometric) k-simplex determined by these vertices is the convex hull

$$\sigma = [v_0 v_1 \cdots v_k] := \Big\{ \sum_{i=0}^k \lambda_i v_i \ \Big| \ \lambda_i \ge 0, \ \sum_{i=0}^k \lambda_i = 1 \Big\}.$$

The points v_0, \ldots, v_k are called the *vertices* of σ . Simplices determined by a subset of v_0, \ldots, v_k are called *faces* of σ . Formally, a simplicial complex is a finite collection of simplices such that every face of a simplex in the collection is also in the collection, and the intersection of any two simplices is either empty or a common face.

Euler Characteristic The *Euler characteristic* χ is a topological invariant that provides a summary of the "shape" or structure of a topological space, such as a simplicial complex. It is defined as the alternating sum of the number of simplices in each dimension, i.e.,

$$\chi(K) = \sum_{k=0}^{d} (-1)^k \sigma_k(K),$$
(1)

where $\sigma_k(K)$ is the number of k-dimensional simplices in the simplicial complex K, and d is the dimension of K. As a topological invariant, the Euler characteristic remains unchanged under transformations like homeomorphisms, making it—despite its simplicity—a fundamental tool for distinguishing topological spaces.

Graph Neural Networks and Message Passing *Graph neural networks* (GNNs) are a class of neural network models designed to operate on graph-structured data. They extend neural networks by incorporating the relational structure inherent to graphs, enabling the learning of tasks such as *node classification*. The core mechanism of many GNNs is *message passing*, an iterative procedure that propagates information through the graph to update node representations based on their local neighborhood. Given a graph G = (V, E), where V is the set of nodes and E is the set of edges, each node $v \in V$ is associated with a *feature vector* \mathbf{x}_v . At each layer t, i.e., each message-passing step, a new node embedding $\mathbf{h}_v^{(t+1)}$ is calculated via

UPDATE
$$\left(\mathbf{h}_{v}^{(t)}, \operatorname{AGG}\left(\left\{\mathbf{h}_{u}^{(t)} \mid u \in \mathcal{N}(v)\right\}\right)\right)$$
, (2)

where $\mathcal{N}(v)$ denotes the set of neighbors of node v, and AGG and UPDATE are learnable functions parameterized

by the model. The AGG function combines information from neighboring nodes, while the UPDATE function refines node embeddings. Popular choices for these functions include mean, sum, and attention mechanisms. Through multiple layers of message passing, GNNs aggregate information from larger neighborhoods, capturing both local and global graph structure (Veličković, 2023).

3. Related Work

GNNs have revolutionized the field of graph representation learning by enabling end-to-end learning of node/graph embeddings through message passing (Kipf & Welling, 2017). However, traditional GNNs face theoretical limitations that pose fundamental obstructions to learning expressive and general representations of graph data (Xu et al., 2019). Related to the latter phenomenon, GNNs are known to suffer from issues like *oversmoothing* (Rusch et al., 2023; Zhang et al., 2024) and oversquashing (Di Giovanni et al., 2023). Hamilton et al. (2017) and Veličković et al. (2018) have addressed these issues by incorporating sampling and attention mechanisms into the messagepassing paradigm. However, even these advancements often show limited performance, particularly in graphs with high heterophily, and there is no "general" GNN capable of handling both heterophilous and homophilous graphs. Recent work in graph machine learning thus started incorporating additional inductive biases into architectures, such as geometric information (Joshi et al., 2023; Pei et al., 2020; Southern et al., 2023) or topological information (Horn et al., 2022; Verma et al., 2024), with the ultimate goal of improving the expressivity of a model, i.e., its capability to distinguish between non-isomorphic families of graphs (Morris et al., 2023).

Many such endeavors arise from the field of *topological* deep learning (Papamarkou et al., 2024), which aims to develop models that are "aware" of the underlying topology of a space, and thus also capable of handling data with higher-order relations. Other constructions include special architectures for heterophilous tasks that are not based on message-passing (Lim et al., 2021), or *modifications* of the graph itself to improve predictive performance. Suresh et al. (2021), for instance, use edge rewiring to raise graph assortativity and thus gain accuracy under low homophily, whereas Luan et al. (2022) mix feature channels during aggregation to obtain state-ofthe-art results on heterophilous benchmarks. Finally, as Rampášek et al. (2022) show, a hybrid model, combining message-passing (local information) with attention (global information) via structural encodings, may exhibit high expressivity and high scalability. Subsequently, Müller et al. (2024) extended these results by providing a taxonomy of elements related to the "design space" of graph transformers. As a geometrical-topological invariant, the ECT is poised to contribute to more expressive architectures. Contributing Being already a popular tool in topological data analysis (Ghrist et al., 2018; Turner et al., 2014), recent extensions started tackling the integration into deep-learning architectures (Röell & Rieck, 2024) or the incorporation of additional invariance properties (Curry et al., 2022; Marsh et al., 2024). Despite advantageous performance in shapeclassification tasks, however, all existing contributions solely focus on global ECTs and do not discuss any local aspects, which are crucial for our approach. In addition to the ECT, some works also use other topology-based tools in graph learning, primarily persistent homology (PH), an expressive but computationally expensive geometricaltopological invariant. Examples of this approach include Rieck et al. (2019) and Hofer et al. (2020), who use PH for graph classification, or Zhao & Wang (2019), who learn a weighted kernel on topological descriptors arising during PH computations. Closest to our approach in spirit is Zhao et al. (2020), who include topological features of graph neighborhoods into a GNN, again leveraging PH. However, to the best of our knowledge, ours is the first work to develop local variants of the ECT for graph-learning tasks, analyze the theoretical properties of such local variants, and finally show their empirical utility for node classification.

4. Methods

Euler Characteristic Transform (ECT) The *Euler Characteristic Transform* (ECT) of a simplicial complex $X \subset \mathbb{R}^n$ is a function ECT(X): $S^{n-1} \times \mathbb{R} \to \mathbb{Z}$, given by

$$ECT(X)(v,t) := \chi(\{x \in X \mid x \cdot v \le t\}), \qquad (3)$$

where χ denotes the Euler characteristic and $x \cdot v$ denotes the Euclidean dot product. The interpretation of ECT(X)is that it scans the ambient space of X in every direction and records the Euler characteristic of the sublevel sets. The ECT(X) is *invertible*, meaning that X can be recovered from ECT(X), as long as X is a so-called constructible set (Curry et al., 2022; Ghrist et al., 2018). The main focus of this work are compact geometric simplicial complexes (like geometric graphs), which are constructible, and thus the invertibility theorem applies in our setting. Note that in practice, we *approximate* ECT(X)via $\overline{\text{ECT}}(X)_{(m,l)} := \text{ECT}(X)_{|\{v_1,\ldots,v_m\}\times\{t_1,\ldots,t_l\}}$ for uniformly-distributed directions $v_1,\ldots,v_m \in S^{n-1}$ and filtration steps $t_1, \ldots, t_l \in \mathbb{R}$. Since X is compact, t_1, \ldots, t_l can be chosen to lie in a compact interval [a, b]with $t_1 = a$ and $t_l = b$, and so that the sequence $\{t_i\}_i$ forms a uniform partition of [a, b]. We note that this approximation is efficiently computable and has a natural representation as a vector of dimension $m \cdot l$. Regarding the choice of the magnitudes of m, l we notice that the *expected* nearestneighbor distance for uniform samples on S^n scales as

 $\mathcal{O}((\log m/m)^{1/n})$ (Beck, 1987), and that the equidistant partitioning of a compact interval scales as $\mathcal{O}(1/l)$, leading to $\mathcal{O}((\log m/m)^{1/(n-1)}l^{-1})$ for the total approximation error of the domain of ECT(X). Curry et al. (2022) prove that the aforementioned approximation actually determines the *true* value, provided that m, l are sufficiently large. We notice that both translations and scalings of X in the ambient space lead to a reparametrization of ECT(X). Hence, ECT(X) remains essentially unaltered (up to a parameter change) under these two types of transformations.

Local ECT (ℓ -ECT) Given a geometric simplicial complex $X \subset \mathbb{R}^n$ and a vertex $x \in X$, we define the *local ECT* of x with respect to $k \ge 0$ as

$$\ell\text{-ECT}_k(x;X) := \text{ECT}(N_k(x;X)), \tag{4}$$

where $N_k(x; X)$ denotes an appropriate local neighborhood of x in X, whose locality scale is controlled by a parameter k. Usually, $N_k(x; X)$ will be either the full subcomplex of X, which is spanned by the k-hop neighbors of x, or the full subcomplex of X, which is spanned by the knearest vertices of x. The first important special case arises when X is a 0-dimensional simplicial complex, i.e., a point cloud. In this case, the full subcomplex of X, which is spanned by the k-nearest vertices of x, $N_k(x; X)$, is given by the k-nearest neighbors of x. Being based on the Euler Characteristic, the construction of ℓ -ECTs appears to be purely topological at first glance. However, in light of the invertibility theorem, we note that ℓ -ECT(x; X) can be interpreted as a *fingerprint* of a local neighborhood of x in X. The upshot is that this fingerprint can be well approximated in practice, making it possible to obtain local representations of combinatorial data embedded in Euclidean space. Similar to the approximation of the ECT, this approximation works by sampling $v_1, \ldots, v_m \in S^{n-1}$ and $t_1, \ldots, t_l \in \mathbb{R}$, and considering $\overline{\text{ECT}}(N_k(x; X))_{(m,l)}$, instead of ℓ -ECT_k(x; X). The latter quantity is wellcomputable in practice, and the approximation error can be controlled by the sample sizes m and l, as we discussed above. Again, this approximation has a natural representation as a vector of dimension $m \cdot l$, enabling us to encode local structural information of point neighborhoods in an approximate lossless way that can readily be used by machine-learning algorithms for downstream tasks.

4.1. Properties of *l*-ECTs

Our formulation of ℓ -ECTs provides a natural representation of local neighborhoods of geometric simplicial complexes. One important special case is that of *featured graphs*, meaning graphs in which every node admits a feature vector. The latter data structure forms the basis of many modern graph-learning tasks, such as node classification, graph classification, or graph regression.

The predominant class of methods to deal with these graph learning problems are message-passing graph neural networks. We develop an alternative procedure for dealing with featured graph data, built on ℓ -ECTs and we show that ℓ -ECTs provide sufficient information to perform message passing, which we explain in the following.

Definition 1. A featured graph is a (non-directed) graph \mathcal{G} such that every node $v \in \mathcal{G}$ admits a feature vector $x(v) \in \mathbb{R}^n$. We denote the set of nodes of \mathcal{G} by $V(\mathcal{G})$, and the set of edges by $E(\mathcal{G})$.

We notice that a featured graph \mathcal{G} can naturally be interpreted as a graph embedded in \mathbb{R}^n , by representing each node feature vector as a point in \mathbb{R}^n , and by drawing an edge between two embedded points if and only if there is an edge between the underlying nodes in \mathcal{G} . This construction yields a graph isomorphism between \mathcal{G} and the embedded graph if and only if for any pair of nodes $v, w \in \mathcal{G}$ with $v \neq w$ we have $x(v) \neq x(w)$ for their associated feature vectors. In practice, the latter assumption can always be achieved by adding an arbitrarily small portion of Gaussian noise to each feature vector, and we therefore may restrict ourselves to featured graphs that yield an isomorphism on their Euclidean embeddings.² We now show that ℓ -ECTs are in fact expressive graph-learning representations.

Theorem 1. Let \mathcal{G} be a featured graph and let $\{\ell \operatorname{ECT}_1(x; \mathcal{G})\}_x$ be the collection of local ECTs with respect to the 1-hop neighborhoods in \mathcal{G} . Then the collection $\{\ell \operatorname{ECT}_1(x; \mathcal{G})\}_x$ provides the necessary (non-learnable) information for performing a single message-passing step on \mathcal{G} , in the sense that for a given vertex $x \in \mathcal{G}$ one can reconstruct the feature vectors of its 1-hop neighborhood from $\ell \operatorname{-ECT}_1(x; \mathcal{G})$.

Theorem 1 tells us that for a featured graph \mathcal{G} , the collection $\{\ell$ -ECT₁ $(x; \mathcal{G})\}_x$ already contains sufficient information to perform a single step of message passing. The advantage of using ℓ -ECTs instead of message passing to represent featured graph data lies in the possibility to *additionally* use $\{\ell$ -ECT_k $(x; \mathcal{G})\}_x$ for $k \ge 2$, which contain both structural and feature vector information of larger neighborhoods of nodes in the graph. This type of information is typically *not* explicitly available through message passing since passing messages to non-direct neighbors depends on prior message passing steps, which solely produce an aggregation of neighboring feature vectors.

In addition to essentially subsuming the information from one message-passing step, we can also show that the ℓ -ECT is "aware" of local structures like subgraphs. As shown by Chen et al. (2020), message-passing graph neural networks *cannot* perform counting of induced subgraphs for *any* connected substructure consisting of 3 or more nodes. By contrast, we will now show that ECTs for featured graphs and their local variants can indeed be used to perform subgraph counting. We start with the definitions of the necessary concepts.

Definition 2. Two featured graphs \mathcal{G}_1 and \mathcal{G}_2 are isomorphic if there is a bijection $\pi \colon V(\mathcal{G}_1) \to V(\mathcal{G}_2)$, such that $(v, w) \in E(\mathcal{G}_1)$ if and only if $(\pi(v), \pi(w)) \in E(\mathcal{G}_2)$ and so that for all $v \in \mathcal{G}_1$ one has $x(v) = x(\pi(v))$ for the respective feature vectors.

A featured graph \mathcal{G}_S is called a *subgraph* of \mathcal{G} if $V(\mathcal{G}_S) \subset V(\mathcal{G})$ and $E(\mathcal{G}_S) \subset E(\mathcal{G})$, such that the respective node features remain unaltered under the induced embedding. A featured graph \mathcal{G}_S is called an *induced subgraph* of \mathcal{G} , if \mathcal{G}_S is a subgraph of \mathcal{G} , and if $E(\mathcal{G}_S) = E(\mathcal{G}) \cap \mathcal{G}_S$. For two featured graphs \mathcal{G} and \mathcal{G}_S , we define $C_{\text{Sub}}(\mathcal{G};\mathcal{G}_S)$ to be the number of subgraphs in \mathcal{G} that are isomorphic to \mathcal{G}_S . Similarly, we define $C_{\text{Ind}}(\mathcal{G};\mathcal{G}_S)$ to be the number of induced subgraphs in \mathcal{G} which are isomorphic to \mathcal{G}_S .

Theorem 2. Two featured graphs \mathcal{G}_1 and \mathcal{G}_2 are isomorphic if and only if $ECT(\mathcal{G}_1) = ECT(\mathcal{G}_2)$.

An immediate consequence of the previous Theorem is: **Corollary 1.** ECT*s can perform subgraph counting.*

We therefore conclude that ECT-based methods for graphrepresentation learning can be more powerful than messagepassing-based approaches, suggesting the development of *hybrid* architectures, making use of both message passing *and* ECT variants.

4.2. Rotation-Invariant Metric based on Local ECTs

The aforementioned invariance properties of ECTs with respect to translations and scalings naturally raise the question if ℓ -ECTs may be used to compare the local neighborhoods of two distinct points/vertices. However, the ECT is sensitive to rotations since rotating the underlying simplicial complex leads to a misalignment of the respective directions in S^{n-1} . Because a local comparison should *not* depend on the choice of a coordinate system, this property is a fundamental obstruction of using ℓ -ECT as a local similarity measure. We therefore construct a novel *rotationinvariant metric* as follows. Let $X, Y \subset \mathbb{R}^n$ be two finite geometric simplicial complexes. Since X, Y are finite, ECT(X) and ECT(Y) only take finitely many values, and we may define a similarity measure $d_{ECT}(X, Y)$ as

$$d_{\mathrm{ECT}}(X,Y) := \inf_{\rho \in \mathrm{SO}(n)} \|(\mathrm{ECT}(X) - \mathrm{ECT}(\rho Y))\|_{\infty}.$$
 (5)

²Alternatively, we can drop the requirement of an embedding by noting that a featured graph can be considered as an abstract simplicial complex \mathcal{G} with an *arbitrary* function $f: \mathcal{G} \to \mathbb{R}^n$ defined on its vertices and edges. In this case, we may define the ECT as $ECT(X)(v, t) := \chi(\{f^{-1}\{x \in \mathbb{R}^n \mid x \cdot v \leq t\})$. This formulation, developed by Marsh & Beers (2023), demonstrates that the ECT is *generally* applicable and does not require node features to provide an embedding of a graph.

We first prove that this similarity measure satisfies the definitions of a metric.

Theorem 3. d_{ECT} is a metric on the collection of rotation classes of finite simplicial complexes embedded in \mathbb{R}^n .

Theorem 3 ensures that we may use d_{ECT} as a metric that measures the similarity between embedded simplicial complexes up to rotation. In particular, for a simplicial complex $X \subset \mathbb{R}^n$ and $x, y \in X$, we have a rotation-invariant measure to compare local neighborhoods of x and y by defining $d_{\text{ECT}}^k(x,y;X)$ as $\inf_{\rho \in SO(n)} \|\ell \text{-ECT}_k(x;X) - \ell \text{-ECT}_k(y;\rho X)\|_{\infty}$. In practice, we approximate $d_{\text{ECT}}(X,Y)$ by

$$\inf_{\rho \in \mathrm{SO}(n)} \left\| \overline{\mathrm{ECT}}(X)_{(m,l)} - \overline{\mathrm{ECT}}(\rho Y)_{(m,l)} \right\|_{\infty}$$
(6)

for a choice of samples $v_1, \ldots, v_m \in S^{n-1}$ and $t_1, \ldots, t_l \in \mathbb{R}$; this works analogously for the local version $d_{\text{ECT}}^k(x, y; X)$. As discussed before, the approximations of the ECTs used in Eq. (6) have a natural vector representation, so that the $\|\bullet\|_{\infty}$ in Eq. (6) is in fact the maximum of the entry-wise absolute differences between the two respective representation vectors. Hence, the approximation shown in Eq. (6) is *efficiently computable*. However, our experiments in Section 5 use the Euclidean metric for differentiability reasons.

4.3. Limitations

While ℓ -ECTs present clear advantages in preserving *local* details, there are some trade-offs to consider. In certain cases, message-passing GNNs, which aggregate information across neighbors, may be preferable, in particular for tasks where *global* context is more important than local details (see Coupette et al. (2025) for a recent analysis of graph-learning datasets under different perspectives). Furthermore, while our method is computationally feasible on medium-sized datasets (as demonstrated in our experiments), the complexity of "naïvely" calculating ℓ -ECTs increases for larger k and with the size and density of the graph, suggesting a need for improved methods (see Section A.2 for an extended discussion).

5. Experiments

In this section, we present experiments to empirically evaluate the performance of the ℓ -ECT-based approach in graph representation learning, focusing on nodeclassification tasks. We aim to demonstrate how ℓ -ECT representations can capture structural information more effectively than traditional message-passing mechanisms, especially in scenarios with high heterophily (even though we consider other scenarios as well). Our experiments compare the performance of ℓ -ECT-based models to several standard GNN models, namely graph attention networks (Veličković et al., 2018, GAT), graph convolutional networks (Kipf & Welling, 2017, GCN), graph isomorphism networks (Xu et al., 2019, GIN), as well as a heterophily-specific architecture (Zhu et al., 2020, H2GCN). Furthermore, we showcase how the rotationinvariant metric from Section 4 may be used for spatial alignment of graph data.

5.1. *l*-ECTs in Graph Representation Learning

The link between message passing and ℓ -ECTs (cf. Theorem 1) encourages us to empirically validate the expressivity of ℓ-ECTs for node-classification tasks. Given a featured graph \mathcal{G} and fixed $k \geq 0$, we assign ℓ -ECT_k(x; \mathcal{G}) to every node $v \in \mathcal{G}$. We then use the ℓ -ECT corresponding to a node together with the respective node feature vector as the *input* to classification models. Subsequently, we focus on XGBoost (Chen & Guestrin, 2016), as we found it to outperform more complex models. However, our ℓ -ECT can be used with *any* model. Notice that our experiments are not about claiming state-of-the-art performance but rather about showcasing that an approach based on ℓ -ECT yields results that are on a par with and often superior to more complex graph-learning techniques based on message passing, while at the same time working well in *both* heterophilous and homophilous settings.

Implementation details We assume that we are given a featured graph \mathcal{G} such that there is an assignment $V(\mathcal{G}) \rightarrow \mathcal{Y}$, with $V(\mathcal{G})$ being the node set of \mathcal{G} and \mathcal{Y} being the space of classes w.r.t. the underlying nodeclassification task. For a fixed $k > 0, x \in V(\mathcal{G})$ and $N_k(x;\mathcal{G})$ being the k-hop neighborhood of x in \mathcal{G} , we then approximate ℓ -ECT_k $(x; \mathcal{G})$ via $\overline{\text{ECT}}(N_k(x; \mathcal{G}))_{(m,l)}$ for sampled directions and filtration steps, as explained in Section 4. We use m = l = 64 (but the number of samples may be tuned in practice) and use the the resulting $m \cdot l$ -dimensional vector(s) $\overline{\text{ECT}}(N_k(x;\mathcal{G}))_{(m,l)}$, together with the feature vector of x, as additional inputs for the classifier. The architecture of our baseline models includes a two-layer MLP after every graph-neighborhood aggregation layer, as well as skip connections and layer normalization. We train each model for 1000 epochs and report the test accuracy corresponding to the state of the model that admits the *maximum validation accuracy* during training. This makes the predictive performance of our baseline models directly comparable with Platonov et al. (2023).

WebKB Datasets For all datasets of the WebKB collection (Pei et al., 2020), our ℓ -ECT-based approach outperforms the baseline GNNs by far (cf. Table 1; GraphSAGE results from Xu et al. 2024). While the combination of both ℓ -ECT₁ and ℓ -ECT₂ performs best for

Model	Cornell	Wisconsin	Texas
GCN	45.0 ± 2.2	44.2 ± 2.6	47.3 ± 1.5
GAT	44.7 ± 2.9	48.2 ± 2.0	51.7 ± 3.2
GIN	46.5 ± 3.1	49.7 ± 2.5	54.2 ± 2.9
GraphSAGE	76.0 ± 3.5	72.9 ± 1.9	71.8 ± 2.4
H2ĜCN	66.2 ± 3.5	70.2 ± 2.3	72.3 ± 3.0
ℓ -ECT ₁	66.8 ± 4.2	81.2 ± 2.9	74.6 ± 0.5
ℓ -ECT ₂	67.0 ± 4.9	76.1 ± 2.8	73.8 ± 2.6
$\ell\text{-}\mathrm{ECT}_1+\ell\text{-}\mathrm{ECT}_2$	67.1 ± 4.1	78.5 ± 2.6	74.8 ± 3.1

Table 1: Performance (accuracy, in percent) of graph-learning models on WebKB datasets (5 training runs).

Table 2: Performance (*accuracy*, in percent) of graphlearning models on *heterophilous* datasets (5 training runs).

Model	Amazon Ratings	Roman Empire
GCN	42.3 ± 0.7	73.3 ± 0.8
GAT	44.6 ± 0.9	76.4 ± 1.2
GIN	44.1 ± 0.8	56.8 ± 1.0
GraphSAGE	42.2 ± 0.5	77.4 ± 0.8
H2GCN	40.1 ± 0.7	64.2 ± 0.9
ℓ -ECT ₁	48.4 ± 0.3	80.4 ± 0.4
ℓ -ECT ₂	49.6 ± 0.3	78.0 ± 0.3
$\ell\text{-}\mathrm{ECT}_1+\ell\text{-}\mathrm{ECT}_2$	49.8 ± 0.3	81.1 ± 0.4

"Texas," using only ℓ -ECT₁ leads to best performance for "Wisconsin." However, for the two aforementioned datasets, the combination of ℓ -ECT₁ and ℓ -ECT₂ only slightly improves the performance in comparison to ℓ -ECT₁, suggesting that 1-hop neighbor information is already sufficiently informative here.

Heterophilous Datasets Platonov et al. (2023) introduced several *heterophilous* datasets; we validate our method on "Amazon Ratings" and "Roman Empire," again observing that the combination of ℓ -ECT₁ + ℓ -ECT₂ performs best, substantially outperforming baseline models (cf. Table 2). The results are closely aligned with findings by Platonov et al. (2023), i.e., that specialized architectures like H2GCN often perform less well than "standard" architectures. Moreover, ℓ -ECT₁ outperforms ℓ -ECT₂ on "Roman Empire," while ℓ -ECT₂ outperforms ℓ -ECT₁ on "Amazon Ratings." We interpret this as 1-hop neighborhoods being particularly informative for "Roman Empire," while 2-hop neighborhoods are more informative for "Amazon Ratings."

Amazon dataset The Amazon dataset (Shchur et al., 2018) consists of the two co-purchase graphs "Computers" and "Photo." While GAT outperforms all methods on "Computers," the combination of ℓ -ECT₁ and ℓ -ECT₂ performs best on "Photo" (cf. Table 3). Overall, ℓ -ECT-based methods exhibit competitive performance here, given that they are *not* based on message passing.

Table 3: Performance (*accuracy*, in percent) of graphlearning models on Amazon datasets (5 training runs).

Model	Computers	Photo
GCN	91.6 ± 1.6	93.6 ± 1.7
GAT	92.4 ± 1.3	94.8 ± 1.1
GIN	55.9 ± 1.5	82.2 ± 1.3
GraphSAGE	89.2 ± 0.9	92.5 ± 0.7
H2GCN	84.5 ± 1.4	92.8 ± 1.2
ℓ -ECT ₁	89.6 ± 0.3	94.1 ± 0.3
ℓ -ECT ₂	90.1 ± 0.5	94.4 ± 0.7
$\ell\text{-}\mathrm{ECT}_1 + \ell\text{-}\mathrm{ECT}_2$	92.2 ± 0.6	94.9 ± 0.6

Table 4: Performance (*accuracy*, in percent) of graphlearning models on *heterophilous* datasets (5 training runs).

Model	Actor	Squirrel	Chameleon
GCN	30.7 ± 2.1	28.9 ± 1.4	42.8 ± 1.8
GAT	31.1 ± 1.8	31.8 ± 1.3	47.3 ± 1.3
GIN	26.5 ± 2.0	35.4 ± 1.5	43.1 ± 1.7
GraphSAGE	30.2 ± 1.4	33.3 ± 0.7	45.2 ± 1.3
H2GCN	30.7 ± 1.9	40.8 ± 1.4	62.7 ± 1.6
ℓ -ECT ₁	31.4 ± 1.9	35.6 ± 0.7	43.5 ± 1.7
ℓ -ECT ₂	30.1 ± 1.3	35.6 ± 0.8	40.4 ± 1.5
$\ell\text{-}\mathrm{ECT}_1 + \ell\text{-}\mathrm{ECT}_2$	30.9 ± 0.7	35.3 ± 1.5	43.9 ± 0.7

Actor/Wikipedia Datasets Moving to additional heterophilous datasets with high feature dimensionality, we compare predictive performance on "Actor" (Pei et al., 2020) as well as "Chameleon" and "Squirrel" (Rozemberczki et al., 2021); cf. Table 4. For "Actor", the ℓ -ECT₁ model achieves the highest accuracy, while ℓ -ECT₁ + ℓ -ECT₂ performs slightly worse. ℓ -ECT₂ performs the lowest on this dataset, suggesting that larger neighborhoods are detrimental here. For the other datasets, the heterophily-specific model H2GCN performs best. However, ℓ -ECT₁ and ℓ -ECT₁ and ℓ -ECT₂ exhibit similar (or even better) performance as *all other standard baselines*, showing the advantages of ℓ -ECT methods even in the absence of hyperparameter tuning.

Table 5: Performance (*accuracy*, in percent) of graphlearning models on "Planetoid" datasets (5 training runs).

Model	Cora	CiteSeer	PubMed
GCN	88.1 ± 1.2	74.6 ± 1.5	85.3 ± 4.7
GAT	88.3 ± 1.1	75.3 ± 1.5	85.7 ± 4.2
GIN	85.0 ± 1.5	72.2 ± 1.7	87.0 ± 0.5
GraphSAGE	82.2 ± 1.2	68.1 ± 1.2	84.3 ± 0.7
H2GCN	85.4 ± 1.6	72.4 ± 1.9	86.4 ± 0.5
ℓ -ECT ₁	87.6 ± 0.6	72.1 ± 0.6	90.2 ± 0.5
ℓ -ECT ₂	87.2 ± 0.7	72.3 ± 0.8	90.3 ± 0.5
$\ell\text{-}\mathrm{ECT}_1 + \ell\text{-}\mathrm{ECT}_2$	87.8 ± 0.6	72.5 ± 0.7	90.3 ± 0.5

Planetoid Datasets We also analyze node-classification performance on datasets from the "Planetoid" collection (Yang et al., 2016), comprising "Cora," "CiteSeer," and "PubMed." We trained all models using a random 75-25 split; cf. Table 5. Although GCN and GAT perform slightly better than ℓ -ECT methods for "Cora" and "CiteSeer," the gap is surprisingly small. For "PubMed," the ℓ -ECT-based models even outperform both all comparison partners. These findings suggest that the theoretical expressivity of ℓ -ECTs, which we formally established in Section 4, also has practical implications, providing an alternative way of dealing with graph data that is not restricted by the underlying model architecture and therefore allows for interpretability.

Post-hoc Evaluation To evaluate our methods across datasets, we used critical difference diagrams, enabling us to compare the performance of various models on both homophilic and heterophilic graph datasets (cf. Section A.7). The results highlight the superior performance of *l*-ECT-based approaches over standard baselines and heterophily-specific architectures such as H2GCN. Notably, the ℓ -ECT₁ + ℓ -ECT₂ method achieved the best average rank of 2 and even the least effective *l*-ECT-based model (*l*-ECT₂) outperformed all non*l*-ECT-based methods, including GAT. Further evaluation against heterophily-specific models reported in the literature corroborates these findings. In comparison to state-of-theart methods such as H2GCN, GPR-GNN, and GloGNN, ℓ -ECT₁ + ℓ -ECT₂ achieved a competitive rank, matching GloGNN and surpassing both GAT and GT. Despite being a general-purpose approach not specifically designed for heterophilic graphs, *l*-ECT-based methods demonstrated exceptional adaptability and robustness across diverse graph structures. These results establish ℓ -ECT-based architectures as a versatile and high-performing solution for node classification tasks, suitable for tackling challenges across a wide range of graph data.

Summary of Node-Classification Experiments We find that ℓ -ECTs work particularly well in situations where aggregating neighboring information is *inappropriate*, such as when dealing with graphs that exhibit a high degree of heterophily. In such contexts, our approach may outperform message-passing-based methods. The upshot of our method is that local graph information can be incorporated *without* the architectural necessity to diffuse information along the graph structure, as it is the case for message-passing-based models. While this discrete diffusion process induced by message passing is useful for a plethora of graph-learning tasks, it can also be an obstruction in learning the right representation for tasks where node features of neighbors in the graph should not be aggregated (cf. Coupette et al. 2025 for a recent analysis of graph-learning datasets in the context of message passing). In this sense, ℓ -ECTs naturally overcome a fundamental limitation inherent to message-passing methods. Another advantage of ℓ -ECTs is that they are *agnostic* to the choice of the downstream model. This permits us to use models that are easy to tune, enabling practitioners to make use of their graph data without necessarily having specialized knowledge in GNN training and tuning while at the same time also working well in the small-sample regime. Moreover, it permits using models that are interpretable, making our method wellsuited for domains where regulatory demands often ask for levels of interpretability that cannot readily be achieved by (graph) neural networks. In fact, by using feature importance values (which are directly available for treebased algorithms like XGBoost) and since the entries of the *l*-ECT vectors that are used as the input for the model can be linked to the directions in the calculation of the ℓ -ECTs, one may obtain a deeper understanding of how the model arrives at predictions (see Section A.3 for a more in-depth discussion and an ablation on the number of directions used to calculate ℓ -ECTs).

5.2. Learning Spatial Alignment of Geometric Graphs

In the following, we use the approach described in Section 4 in order to learn the spatial alignment of two data spaces by re-rotating one into the other. We start by showing that synthetic data, which only differs up to a rotation, can be re-aligned using ℓ -ECTs. Moreover, we show that this alignment is stable with respect to noise, making it a robust measure for the comparison of local neighborhoods in data. In comparison to other spatial alignment methods like the iterative closest point algorithm, ours does not necessitate the computation of all pairwise distances between points in the respective spaces. The latter is often a computational bottleneck, especially for large datasets, thus positioning our method for spatial alignment as a computationally more efficient method in practice. While alignment methods like Procrustes alignment are restricted to pointcloud data, we observe that our approach is also capable of aligning embedded graph data (or, more generally, simplicial complexes). This makes it particularly useful for dealing with geometric graphs, constituting a highlyefficient alternative to more involved machine-learning models like geometric GNNs (Joshi et al., 2023).

For the subsequent learning problem (see Section A.5 for an example of how to align point-cloud data), we assume that we are given two embedded simplicial complexes $X, Y \subset \mathbb{R}^n$. In light of Section 4, the metric properties of d_{ECT} ensure that $d_{\text{ECT}}(X, Y) = 0$ if X and Y only differ up to a rotation. We therefore approximate $d_{\text{ECT}}(X, Y)$ via

$$\min_{\rho \in \mathrm{SO}(n)} \left\| \overline{\mathrm{ECT}}(X)_{(m,l)} - \overline{\mathrm{ECT}}(\rho Y)_{(m,l)} \right\|_2^2 \tag{7}$$



(b) Alignment results

Figure 1: A comparison of the Hausdorff distances of aligned graphs. The black dots represents the Hausdorff distance between the original graph and a randomly-rotated version of itself. Our ℓ -ECT-based alignment always results in substantially lower distances, with a median distance close to zero.

for a choice of directions $v_1, \ldots, v_m \in S^{n-1}$ and filtration steps $t_1, \ldots, t_l \in \mathbb{R}$. As explained in Section 4, the ECT approximations are given by vectors, making it feasible to approach the above learning problem by any gradient-based learning algorithm. The advantage of this formulation is that it yields both the rotation-invariant loss and the rotation that leads to this minimum loss.

Geometric graphs provide a compelling example of the utility of our method, particularly in addressing the challenging problem of graph re-alignment. For this analysis, we focus on a specific type of geometric graph known as the k-star graph. A k-star graph is defined as a tree with one internal node and k leaves, i.e., a simple structure with relevant geometric properties (cf. Figure 1a). To embed such graphs into a 2D space, we assign a unique 2D vector to each node, ensuring the assigned vectors are equidistant to maintain structural symmetry. Furthermore, to introduce variability and assess robustness, we subject each embedded graph to a random 2D rotation, simulating realistic perturbations encountered in practical settings. The central goal is to recover the original graph's orientation by learning the rotation matrix using the metric defined Section 4. We evaluate the performance of our approach by measuring the similarity between the original graph and its re-rotated version using the Hausdorff distance, a metric that quantifies the maximal deviation between two sets of points. To ensure significance, we repeated the learning procedure

200 times, maintaining consistent initializations for both the graph and the rotation matrix. Figure 1 shows the results; we observe that our realignment procedure *consistently* achieves small Hausdorff distances with medians near zero, indicating the successful recovery of the original graph's orientation. By contrast, the Hausdorff distances between the original graph and its perturbed version are significantly larger. While equivariant GNNs have been shown to struggle with distinguishing the orientation of rotationally-symmetric structures (Joshi et al., 2023), our method generates graph representations that are initially sensitive to rotations but can be made rotation-invariant through alignment of the underlying ECTs.

These findings highlight the potential of ECT-based metrics for *robust* geometric graph alignment, paving the way for broader applications in domains requiring precise graphbased comparisons such as the analysis of geometric graphs with constrained parameter budgets (Maggs et al., 2024). This result is particularly notable since the ℓ -ECT easily outperforms more involved architectures, pointing towards its overall utility as an alternative to message-passing graph neural networks. At the same time, we believe that the ℓ -ECT could also help in aligning higher-order data like geometric simplicial complexes.

6. Discussion

We introduced the Local Euler Characteristic Transform (l-ECT), providing a novel approach to graph-representation learning that preserves local structural information without relying on aggregation. Our method addresses fundamental limitations in message-passing neural networks, particularly in tasks where aggregating neighboring information is suboptimal, such as in graphs with heterophily. By retaining critical local details, ℓ-ECTs enable more nuanced and expressive representations, offering significant advantages in node classification tasks. One key strength of our approach is its model-agnostic nature, allowing it to be paired with interpretable machine learning models (in our experiments, an XGBoost model was used to provide feature importance values, for instance). This is particularly useful in domains such as healthcare, finance, and legal applications, where regulatory frameworks demand high levels of transparency and interpretability that are often difficult to achieve with blackbox neural networks. By leveraging ℓ -ECTs, we can satisfy these requirements while maintaining the high expressivity and high predictive performance required for graph-learning tasks. In this way, ℓ -ECT-based representations offer a novel pathway toward interpretable machine learning on graph data: they yield topologically-grounded, vectorized encodings of local structure that not only retain predictive power but also support downstream diagnostics.

Future Work: Higher-Order Domains Being situated at the intersection of geometry and topology, our ℓ -ECT method is part of the nascent field of *topological deep learning* (Papamarkou et al., 2024), which aims to develop novel inductive biases that are capable of leveraging additional structural information from data, both in the context of graphs and in the context of higher-order domains like *simplicial complexes*. It is in this context where we believe that future work could be beneficial, in particular since recent research (Ballester et al., 2025) showed that tasks on such domains are highly challenging for existing GNNs. Given the advantageous scalability properties of the ECT (Röell & Rieck, 2024; Turner et al., 2014), we believe that this constitutes a useful avenue for future research.

Future Work: Comparing Representations Containing both geometrical and topological components, we also believe ECT-based methods to be suitable in navigating different representations. Since the ECT can be considered a compression technique with controllable fidelity properties (Röell & Rieck, 2025), it could be useful in condensing latent spaces, thus permitting simple and efficient comparisons of models as hyperparameters are being varied. Such *multiverse analyses* are vital for ensuring reproducibility (Bell et al., 2022; Wayland et al., 2024).

Future Work: Hybrid Models Beyond representation learning on graphs, our *l*-ECT framework also opens up new applications in domains where local structure is critical, such point-cloud analysis (including sensor data or computer-graphics data), 3D shape analysis, or data from the life sciences (like molecular data or biological networks). Future work could thus explore more efficient algorithms for computing ECTs and ℓ -ECTs at scale, as well as hybrid approaches that balance local and global information more effectively. A highly-relevant direction would be the integration of ℓ -ECTs into existing message-passing neural networks, similar to recent work that leverages persistent homology (Verma et al., 2024). Moreover, heterophilyspecific mechanisms such as a separation of neighborhood aggregation (as used in specialized GNN architectures) may be incorporated into our ℓ -ECT-based framework to further strengthen its expressivity in the presence of highheterophily graphs.

Software and Data

Our code is available under https://github.com/ aidos-lab/Diss-l-ECT. We make use of standard benchmarking datasets, loaded and processed via the PyTorch Geometric library (Fey & Lenssen, 2019).

Acknowledgments

The first author acknowledges the use of ChatGPT for grammar suggestions. Both authors are very grateful for the discussions with the anonymous reviewers, in particular reviewer nYTM, and the area chair, who also believed in the merits of this work. This work has received funding from the Swiss State Secretariat for Education, Research, and Innovation (SERI).

Impact Statement

This paper presents work whose goal is to advance the field of graph machine learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- Ballester, R., Röell, E., Schmid, D. B., Alain, M., Escalera, S., Casacuberta, C., and Rieck, B. MANTRA: The Manifold Triangulations Assemblage. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum? id=X6y5CC44HM.
- Beck, J. Irregularities of distribution. I. *Acta Mathematica*, 159:1–49, 1987.

- Bell, S. J., Kampman, O., Dodge, J., and Lawrence, N. Modeling the machine learning multiverse. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 18416– 18429. Curran Associates, Inc., 2022.
- Bo, D., Wang, X., Shi, C., and Shen, H. Beyond lowfrequency information in graph convolutional networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 3950–3957, 2021.
- Borgwardt, K., Ghisu, E., Llinares-López, F., O'Bray, L., and Rieck, B. Graph kernels: State-of-the-art and future challenges. *Foundations and Trends*® in Machine *Learning*, 13(5–6):531–712, 2020.
- Chen, T. and Guestrin, C. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 785–794, 2016.
- Chen, Z., Chen, L., Villar, S., and Bruna, J. Can graph neural networks count substructures? In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 10383–10395. Curran Associates, Inc., 2020.
- Chien, E., Peng, J., Li, P., and Milenkovic, O. Adaptive universal generalized PageRank graph neural network. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum? id=n6j17fLxrP.
- Coupette, C., Wayland, J., Simons, E., and Rieck, B. No metric to rule them all: Toward principled evaluations of graph-learning datasets. *arXiv:2502.02379*, 2025.
- Curry, J., Mukherjee, S., and Turner, K. How many directions determine a shape and other sufficiency results for two topological transforms. *Transactions of the American Mathematical Society, Series B*, 9(32):1006–1043, 2022.
- Di Giovanni, F., Giusti, L., Barbero, F., Luise, G., Lio, P., and Bronstein, M. M. On over-squashing in message passing neural networks: The impact of width, depth, and topology. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 7865–7885. PMLR, 2023.
- Du, L., Shi, X., Fu, Q., Ma, X., Liu, H., Han, S., and Zhang, D. GBK-GNN: Gated bi-kernel graph neural networks for modeling both homophily and heterophily.

In Proceedings of the ACM Web Conference 2022, pp. 1550–1558, 2022.

- Fey, M. and Lenssen, J. E. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop* on *Representation Learning on Graphs and Manifolds*, 2019.
- Ghrist, R., Levanger, R., and Mai, H. Persistent homology and Euler integral transforms. *Journal of Applied and Computational Topology*, 2:55–60, 2018.
- Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of* the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778, 2016.
- Hofer, C. D., Graf, F., Rieck, B., Niethammer, M., and Kwitt, R. Graph filtration learning. In Daumé III, H. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 4314–4323. PMLR, 2020.
- Horn, M., De Brouwer, E., Moor, M., Moreau, Y., Rieck, B., and Borgwardt, K. Topological graph neural networks. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum? id=oxxUMeFwEHd.
- Joshi, C. K., Bodnar, C., Mathis, S. V., Cohen, T., and Liò, P. On the expressive power of geometric graph neural networks. In *International Conference on Learning Representations*, 2023. URL https://openreview. net/forum?id=Rkxj1GXn9_.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum? id=SJU4ayYgl.
- Li, X., Zhu, R., Cheng, Y., Shan, C., Luo, S., Li, D., and Qian, W. Finding global homophily in graph neural networks when meeting heterophily. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S. (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 13242– 13256. PMLR, 17–23 Jul 2022.

- Lim, D., Hohne, F., Li, X., Huang, S. L., Gupta, V., Bhalerao, O., and Lim, S. N. Large scale learning on nonhomophilous graphs: New benchmarks and strong simple methods. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), Advances in Neural Information Processing Systems, volume 34, pp. 20887–20902. Curran Associates, Inc., 2021.
- Luan, S., Hua, C., Lu, Q., Zhu, J., Zhao, M., Zhang, S., Chang, X.-W., and Precup, D. Revisiting heterophily for graph neural networks. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 1362–1375. Curran Associates, Inc., 2022.
- Maggs, K., Hacker, C., and Rieck, B. Simplicial representation learning with neural *k*-forms. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum? id=Djw0XhjHZb.
- Marsh, L. and Beers, D. Stability and inference of the Euler characteristic transform. *arXiv:2303.13200*, 2023.
- Marsh, L., Zhou, F. Y., Qin, X., Lu, X., Byrne, H. M., and Harrington, H. A. Detecting temporal shape changes with the Euler characteristic transform. *Transactions of Mathematics and Its Applications*, 8(2), 2024.
- Maurya, S. K., Liu, X., and Murata, T. Improving graph neural networks with simple architecture design. *arXiv:2105.07634*, 2021.
- Mernyei, P. and Cangea, C. Wiki-cs: A wikipedia-based benchmark for graph neural networks. *arXiv preprint arXiv:2007.02901*, 2020.
- Morris, C., Lipman, Y., Maron, H., Rieck, B., Kriege, N. M., Grohe, M., Fey, M., and Borgwardt, K. Weisfeiler and Leman go machine learning: The story so far. *Journal of Machine Learning Research*, 24(333):1–59, 2023.
- Müller, L., Galkin, M., Morris, C., and Rampášek, L. Attending to graph transformers. *Transactions on Machine Learning Research*, 2024. URL https://openreview.net/forum?id=HhbqHBBrfZ.
- Papamarkou, T., Birdal, T., Bronstein, M., Carlsson, G., Curry, J., Gao, Y., Hajij, M., Kwitt, R., Liò, P., Lorenzo, P. D., Maroulas, V., Miolane, N., Nasrin, F., Ramamurthy, K. N., Rieck, B., Scardapane, S., Schaub, M. T., Veličković, P., Wang, B., Wang, Y., Wei, G.-W., and Zamzmi, G. Position: Topological deep learning is the new frontier for relational learning. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume

235 of *Proceedings of Machine Learning Research*, pp. 39529–39555. PMLR, 2024.

- Pei, H., Wei, B., Chang, K. C.-C., Lei, Y., and Yang, B. Geom-GCN: Geometric graph convolutional networks. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum? id=S1e2agrFvS.
- Platonov, O., Kuznedelev, D., Diskin, M., Babenko, A., and Prokhorenkova, L. A critical look at the evaluation of GNNs under heterophily: Are we really making progress?, 2023. URL https://openreview.net/ forum?id=tJbbQfw-5wv.
- Rampášek, L., Galkin, M., Dwivedi, V. P., Luu, A. T., Wolf, G., and Beaini, D. Recipe for a general, powerful, scalable graph transformer. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), Advances in Neural Information Processing Systems, volume 35, pp. 14501–14515. Curran Associates, Inc., 2022.
- Rieck, B., Bock, C., and Borgwardt, K. A persistent Weisfeiler–Lehman procedure for graph classification. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings* of the 36th International Conference on Machine Learning, number 97 in Proceedings of Machine Learning Research, pp. 5448–5458. PMLR, 2019.
- Röell, E. and Rieck, B. Differentiable Euler characteristic transforms for shape classification. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum? id=M0632iPq3I.
- Röell, E. and Rieck, B. Point cloud synthesis using inner product transforms. arXiv:2410.18987, 2025.
- Rozemberczki, B., Allen, C., and Sarkar, R. Multiscale attributed node embedding. *Journal of Complex Networks*, 9(2):cnab014, 2021.
- Rusch, T. K., Bronstein, M. M., and Mishra, S. A survey on oversmoothing in graph neural networks. *arXiv:2303.10993*, 2023.
- Shchur, O., Mumme, M., Bojchevski, A., and Günnemann, S. Pitfalls of graph neural network evaluation. In *Relational Representation Learning Workshop (R2L) at NeurIPS*, 2018.
- Shi, Y., Huang, Z., Feng, S., Zhong, H., Wang, W., and Sun, Y. Masked label prediction: Unified message passing model for semi-supervised classification. In Zhou, Z.-H. (ed.), *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, pp. 1548–1554. International Joint Conferences on Artificial Intelligence Organization, 2021.

- Southern, J., Wayland, J., Bronstein, M., and Rieck, B. Curvature filtrations for graph generative model evaluation. In Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 63036–63061. Curran Associates, Inc., 2023.
- Suresh, S., Budde, V., Neville, J., Li, P., and Ma, J. Breaking the limit of graph neural networks by improving the assortativity of graphs with local mixing patterns. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 1541–1551, 2021.
- Topping, J., Giovanni, F. D., Chamberlain, B. P., Dong, X., and Bronstein, M. M. Understanding oversquashing and bottlenecks on graphs via curvature. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum? id=7UmjRGzp-A.
- Turner, K., Mukherjee, S., and Boyer, D. M. Persistent homology transform for modeling shapes and surfaces. *Information and Inference: A Journal of the IMA*, 3(4): 310–344, 12 2014.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. Graph attention networks. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum? id=rJXMpikCZ.
- Veličković, P. Everything is connected: Graph neural networks. *Current Opinion in Structural Biology*, 79: 102538, 2023.
- Verma, Y., Souza, A. H., and Garg, V. Topological neural networks go persistent, equivariant, and continuous. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 49388–49407. PMLR, 2024.
- von Rohrscheidt, J. and Rieck, B. Topological singularity detection at multiple scales. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *Proceedings of the 40th International Conference* on Machine Learning, volume 202 of Proceedings of Machine Learning Research, pp. 35175–35197. PMLR, 2023.
- Wang, X. and Zhang, M. How powerful are spectral graph neural networks. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S. (eds.), *Proceedings of the 39th International Conference*

on Machine Learning, volume 162 of Proceedings of Machine Learning Research, pp. 23341–23362. PMLR, 2022.

- Wayland, J., Coupette, C., and Rieck, B. Mapping the multiverse of latent representations. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, number 235 in Proceedings of Machine Learning Research, pp. 52372–52402, 2024.
- Wu, F., Souza, A., Zhang, T., Fifty, C., Yu, T., and Weinberger, K. Simplifying graph convolutional networks. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 6861–6871. PMLR, 2019.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL https:// openreview.net/forum?id=ryGs6iA5Km.
- Xu, Y., Huang, H., and State, R. CTQW-GraphSAGE: Trainabel continuous-time quantum walk on graph. In *Proceedings of the 33rd International Conference on Artificial Neural Networks*, pp. 79–92. Springer, 2024.
- Yang, Z., Cohen, W., and Salakhudinov, R. Revisiting semisupervised learning with graph embeddings. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The* 33rd International Conference on Machine Learning, volume 48 of Proceedings of Machine Learning Research, pp. 40–48. PMLR, 2016.
- Zhang, X., Xu, Y., He, W., Guo, W., and Cui, L. A comprehensive review of the oversmoothing in graph neural networks. In Sun, Y., Lu, T., Wang, T., Fan, H., Liu, D., and Du, B. (eds.), *Computer Supported Cooperative Work and Social Computing*, pp. 451–465. Springer, 2024.
- Zhao, Q. and Wang, Y. Learning metrics for persistencebased summaries and applications for graph classification. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019.
- Zhao, Q., Ye, Z., Chen, C., and Wang, Y. Persistence enhanced graph neural network. In *International Conference on Artificial Intelligence and Statistics*, pp. 2896–2906. PMLR, 2020.
- Zheleva, E. and Getoor, L. To join or not to join: the illusion of privacy in social networks with mixed public

and private user profiles. In *Proceedings of the 18th International Conference on World Wide Web*, pp. 531– 540, 2009.

- Zhu, J., Yan, Y., Zhao, L., Heimann, M., Akoglu, L., and Koutra, D. Beyond homophily in graph neural networks: Current limitations and effective designs. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 7793–7804. Curran Associates, Inc., 2020.
- Zhu, J., Rossi, R. A., Rao, A., Mai, T., Lipka, N., Ahmed, N. K., and Koutra, D. Graph neural networks with heterophily. In *Proceedings of the AAAI Conference* on Artificial Intelligence, volume 35, pp. 11168–11176, 2021.

A. Appendix

A.1. Proofs

We briefly restate all theorem from the main text for the reader's convenience before providing proofs.

Theorem 1. Let \mathcal{G} be a featured graph and let $\{\ell \text{-}\mathrm{ECT}_1(x;\mathcal{G})\}_x$ be the collection of local ECTs with respect to the 1-hop neighborhoods in \mathcal{G} . Then the collection $\{\ell \text{-}\mathrm{ECT}_1(x;\mathcal{G})\}_x$ provides the necessary (non-learnable) information for performing a single message-passing step on \mathcal{G} , in the sense that for a given vertex $x \in \mathcal{G}$ one can reconstruct the feature vectors of its 1-hop neighborhood from $\ell \text{-}\mathrm{ECT}_1(x;\mathcal{G})$.

Proof. By the remark in Section 4 (in the paragraph right above the original statement of this Theorem), we may assume that the natural embedding of \mathcal{G} into \mathbb{R}^n is a graph isomorphism. Then, making use of the invertibility theorem, the 1-hop neighborhood of a point x in the embedding of \mathcal{G} can be reconstructed from ℓ -ECT₁(x; \mathcal{G}). Therefore, the feature vectors of x and its 1-hop neighbors can be deduced from ℓ -ECT₁(x; \mathcal{G}), which is the only non-learnable information one needs to perform a message-passing step.

Theorem 2. Two featured graphs \mathcal{G}_1 and \mathcal{G}_2 are isomorphic if and only if $ECT(\mathcal{G}_1) = ECT(\mathcal{G}_2)$.

Proof. When two featured graphs are isomorphic in the sense of Definition 2, their respective Euclidean embeddings produce equal ECTs by construction because the node feature vectors of two corresponding points under the isomorphism are equal. By contrast, let us assume that $ECT(\mathcal{G}_1) = ECT(\mathcal{G}_2)$. Then by the invertibility theorem, the Euclidean embeddings of \mathcal{G}_1 and \mathcal{G}_2 are equal. Therefore, the only information that may tell apart the two graphs are their node labels, but this means that \mathcal{G}_1 and \mathcal{G}_2 are isomorphic.

Theorem 3. d_{ECT} is a metric on the collection of rotation classes of finite simplicial complexes embedded in \mathbb{R}^n .

Proof. $d_{\text{ECT}}(X, X) = 0$ holds for ρ being the identity. Now assume that $d_{\text{ECT}}(X, Y) = 0$. Then there exists $\rho \in \text{SO}(n)$ with $\|(\text{ECT}(X) - \text{ECT}(\rho Y))\|_{\infty} = 0$. As $\|\bullet\|_{\infty}$ is a norm, it follows that $\text{ECT}(X) = \text{ECT}(\rho Y)$, and by the invertibility theorem we obtain $X = \rho Y$. This shows the first property of a metric (note that positivity follows from $\|\bullet\|_{\infty}$). For symmetry, note that $\|(\text{ECT}(X) - \text{ECT}(\rho Y))\|_{\infty} = \|(\text{ECT}(\rho^{-1}X) - \text{ECT}(Y))\|_{\infty}$ since rotations are invertible. For the triangle inequality, let Z be another finite simplicial complex. $d_{\text{ECT}}(X, Z)$ then reads $\inf_{\rho \in \text{SO}(n)} \|(\text{ECT}(X) - \text{ECT}(\rho Z))\|_{\infty}$, which is less than or equal to $\inf_{\rho, \rho' \in \text{SO}(n)} (\|(\text{ECT}(X) - \text{ECT}(\rho'Y))\|_{\infty} + \|(\text{ECT}(Y) - \text{ECT}((\rho')^{-1}\rho Z))\|_{\infty})$. This term, however, is equal to $\inf_{\rho, \rho' \in \text{SO}(n)} (\|(\text{ECT}(X) - \text{ECT}(\rho'Y))\|_{\infty} + \|(\text{ECT}(Y) - \text{ECT}((\rho')^{-1}\rho Z))\|_{\infty})$, which is equal to $\inf_{\rho \in \text{SO}(n)} \|(\text{ECT}(X) - \text{ECT}(\rho'Y))\|_{\infty} + \inf_{\rho \in \text{SO}(n)} \|(\text{ECT}(X) - \text{ECT}(\rho Z))\|_{\infty}$. But this final term is precisely the definition of $d_{\text{ECT}}(X, Y) + d_{\text{ECT}}(Y, Z)$. □

A.2. Computational Complexity

For a fixed node x, the computational complexity of ℓ -ECT_k(x) is $O(m \cdot l \cdot |N_k(x)|)$, where: (i) m is the number of sampled directions, (ii) l is the number of filtration steps, and (iii) $|N_k(x)|$ is the number of vertices (or simplices) in the k-hop neighborhood of x.



Figure 2: Feature importance scores of an XGBoost model for the "Coauthor Physics" dataset (using ℓ -ECT₁). Only a small number of features admit high importance scores.

A.3. Ablation on Directions and Interpretability

Coming back to our approximation of ECT(X) via $\overline{\text{ECT}}(X)_{(m,l)} := \text{ECT}(X)_{\{v_1,...,v_m\}\times\{t_1,...,t_l\}}$ for uniformlydistributed directions $v_1, \ldots, v_m \in S^{n-1}$ and filtration steps $t_1, \ldots, t_l \in \mathbb{R}$, we notice that the $(l \cdot (j-1) + 1)$ -th till $(l \cdot j)$ -th entries of $\overline{\text{ECT}}(X)_{(m,l)}$ correspond to the direction v_j . The latter gives us the opportunity to get a deeper understanding of how the model predicts its outcome, by analyzing its feature importance values (which are available for tree-based algorithms like XGBoost). Therefore, our approach enables us to analyze which features, i.e., directions, of the underlying ECT vector are most important. In practice, we often observe that a small number of features admits high feature importance with respect to the corresponding model (cf. Figure 2). This raises the question if we may use a *smaller* random collection of features and still obtain reasonably useful results. We therefore ran experiments for a collection of datasets for a varying number of randomly-sampled entries of the ℓ -ECT₁ vector; cf. Table 6. Here, 4096 corresponds to the whole vector. We observe that for certain datasets, such as "Coauthor CS," "Coauthor Physics," and "Amazon Ratings," the performance of the model only slightly changes when using a reduced version of the ℓ -ECT₁ vector. In light of the results by Curry et al. (2022), this observation is not entirely surprising—one main claim therein is that the ECT can be determined using a small number of directions.

Table 6: Mean accuracy (in percent, 5 runs each) for different node-classification tasks, and varying numbers of randomlysampled entries of the corresponding ℓ -ECT₁ vectors.

Dataset	0	50	100	500	1000	4096
WikiCS	67.8	69.2	70.5	71.3	72.7	74.6
Coauthor CS	92.1	92.3	92.4	92.5	92.6	92.6
Coauthor Physics	95.2	95.6	95.6	95.8	95.9	96.1
Roman Empire	64.7	73.7	75.8	78.3	79.7	80.4
Amazon Ratings	47.9	47.9	48.2	48.4	48.2	48.4



Table 7: Performance (accuracy, in percent) of graph-learning models on WikiCS dataset (5 training splits).

Figure 3: A comparison of the squared L^2 distances of ℓ -ECTs of aligned and non-aligned MNIST digits of "1," respectively.

A.4. Additional Node Classification Experiment

WikiCS Dataset To further validate the effectiveness of our approach, we consider the WikiCS dataset (Mernyei & Cangea, 2020), a medium-sized co-occurrence graph derived from Wikipedia articles on computer science topics. Nodes represent articles, and edges reflect mutual links between them. Each node is equipped with a 300-dimensional embedding, and the task is to classify articles into one of several predefined categories. As shown in Table 7, our ℓ -ECT₁-based method achieves competitive performance compared to message-passing baselines. While GAT obtains the best accuracy overall, ℓ -ECT₁ performs on par with H2GCN and GCN, despite not relying on neighborhood aggregation. This supports the idea that ℓ -ECT-based representations can serve as effective input features in classification settings, even for graphs with moderately homophilous structures. The small standard deviation further illustrates the stability of our method across splits.

A.5. Spatial Alignment of High-Dimensional Data

Following our previous observations that d_{ECT} enables us to align two spaces, we now use it to investigate its effect on high-dimensional data. We start this discussion with the well-known MNIST benchmark dataset, following an analysis of local geometrical-topological structures that we performed previously (von Rohrscheidt & Rieck, 2023). We thus first represent each (gray-scale) image in the dataset as a 784-dimensional vector, by flattening the image. In this way, we obtain a high-dimensional point cloud corresponding to the dataset. Subsequently, we sample 300 points of digits of "1" and calculate the pairwise distances of their respective ℓ -ECT (with respect to the whole point cloud), for k = 10. Finally, we calculate the pairwise distances of the respective aligned ℓ -ECTs (by using the approach of Eq. (7) with k = 10). Figure 3 shows the results; we observe that the aligned ℓ -ECTs have a significantly lower squared L^2 distance (with a median of ≈ 112) than the non-aligned ones (with a median of ≈ 224), showcasing that rotations cause dissimilarity between small neighborhoods of points, in many cases.

A.6. Homophily Scores of Node-Classification Experiments

Table 8 reveals that our benchmark suite spans the entire range from extreme homophily to extreme heterophily. At the homophilic end lie the Amazon co-purchase graphs "Computers" and "Photo," together with the "Planetoid" citation graphs "Cora," "CiteSeer," and "PubMed." In each of these networks, at least *seven of every ten* edges connect nodes with identical class labels, replicating the conditions under which early message-passing GNNs achieved their seminal successes. Near the middle of the spectrum, "Cornell" and the "Amazon Ratings" datasets exhibit mixed behavior, with roughly one third of their edges being heterophilous. Thus, neighborhood aggregation still conveys useful class-specific information, but the signal is noticeably diluted. The lower end is populated by the remaining datasets, where fewer than *one edge in three* is homophilic, and by "Roman Empire," the most extreme case in our experimental suite, where only about *one edge in twenty* links same-label endpoints.

Table 8: Edge-homophily ratios for every dataset used in our node-classification experiments, sorted in ascending order. As the table shows, our experiments comprise a wide variety of datasets.

Dataset	$H_{\rm edge}$
Roman Empire	0.047
Texas	0.110
Wisconsin	0.210
Actor	0.217
Squirrel	0.220
Chameleon	0.230
Cornell	0.300
Amazon Ratings	0.380
CiteSeer	0.736
Amazon Computers	0.777
PubMed	0.802
Cora	0.810
Amazon Photo	0.827

Because six datasets are heterophilic, five are strongly homophilic, and two occupy the transition zone, we believe our experimental suite to be effectively balanced. A model must therefore operate reliably across sharply different structural regimes to achieve consistently high average rank. Classical message-passing architectures depend on homophily and tend to deteriorate as the ratio falls, whereas our empirical results from Section 5 demonstrate that our proposed ℓ -ECT representations retain competitive-and often even superioraccuracy regardless of homophily level. The most conspicuous gains appear precisely on the graphs where neighbor aggregation is least informative, namely "Roman Empire," "Texas," and "Wisconsin," confirming that ℓ -ECT features capture structural cues that message passing alone fails to exploit. Hence, the numerical landscape mapped out in Table 8 substantiates the claim that our experimental design both stresses the limits of common GNNs while at the same time showcasing the robustness of ℓ -ECT-based approaches in settings where label agreement along edges is sparse.



Figure 4: *Critical difference diagram* showing the ranks of different models across all node-classification tasks from Section 5. Even the worst-performing ℓ -ECT-based approach (ℓ -ECT₂) exhibits superior performance to *all other methods*, when averaged across all tasks.

A.7. Post-hoc Evaluation of Node-Classification Experiments

A critical difference diagram arranges the average ranks of multiple models across a set of datasets in order to facilitate overall performance comparisons between the model performances. Such diagrams are commonly used when comparing a suite of models on different datasets (cf. Borgwardt et al. (2020) for similar plots in the context of *graph kernels*).

Figure 4 shows the results for all node-classification results from Section 4, including both homophilic and heterophilic graph datasets. We observe that the ℓ -ECT-based approaches outperform standard methods and the heterophily-specific architecture H2GCN by far, when averaged over all datasets.³ The best-performing method ℓ -ECT₁ + ℓ -ECT₂ exhibits an average rank of 2, while the worst performing method is GIN with an average rank of 5.7. Even the worst-performing ℓ -ECT-based method (ℓ -ECT₂) performs better than the best non- ℓ -ECT-based method, i.e., GAT. However, the most interesting fact that can be gleaned from the diagram involves the statistical significance of the results. Methods connected over the same bar are not performing statistically significantly differently. This seemingly *negative* result has a positive implication: Despite being orders of magnitude more complex, even specialized graph neural networks do not perform statistically significantly better than ℓ -ECT-based methods. Given that our results are based on a standard XGBoost model without *any* task-specific hyperparameter tuning, we believe that this demonstrates the potential and practical utility of our proposed methods.

Table 9: Ranks (lower is better) of models from Platonov et al. (2023) across the *heterophilic* datasets therein, in comparison to our methods. Notice that our method is a *general-purpose* method for node classification and neither geared towards heterophily nor homophily.

Model	Rank
H2GCN (Zhu et al., 2020)	18.3
CPGNN (Zhu et al., 2021)	16.8
GPR-GNN (Chien et al., 2021)	15.3
ResNet (He et al., 2016)	13.8
1-ECT1	12.4
1-ECT2	12.4
GAT (Veličković et al., 2018)	12.3
GT (Shi et al., 2021)	11.0
l-ECT1 + l-ECT (ours)	11.0
GloGNN (Li et al., 2022)	11.0
ResNet+SGC (Wu et al., 2019)	10.8
FAGCN (Bo et al., 2021)	10.0
JacobiConv (Wang & Zhang, 2022)	9.8
GCN (Kipf & Welling, 2017)	9.6
GBK-GNN (Du et al., 2022)	9.0
ResNet+adj (Zheleva & Getoor, 2009)	7.3
SAGE (Hamilton et al., 2017)	5.9
GAT-sep (Veličković et al., 2018)	5.5
GT-sep (Shi et al., 2021)	5.3
FSGNN (Maurya et al., 2021)	3.0

To further evaluate the performance of our methods in comparison to those reported in the literature, we also included a comparison with the results presented by Platonov et al. (2023), using the ranks of the respective models as the basis for evaluation; cf. Table 9. Among the listed methods, several, such as H2GCN, CPGNN, and GPR-GNN, are explicitly designed for heterophilic graph settings, leveraging specialized architectures to handle the challenges posed by such data. In contrast, our ℓ -ECT₁ + ℓ -ECT₂ method, despite being a general-purpose approach not tailored specifically for heterophilic settings, achieves a competitive rank of 11. This performance is on a par with other top-performing heterophily-specific models, such as GloGNN, and outperforms well-established architectures like GT and GAT by a significant margin. Overall, these results highlight the robustness and adaptability of our method, demonstrating its ability to handle diverse graph structures effectively without requiring customization for heterophilic scenarios. In consideration of the results given in Figure 4, this makes *l*-ECT-based approaches a versatile general-purpose solution for node -classification tasks.

³We used https://github.com/hfawaz/cd-diagram for the creation of the critical difference diagram.

Diss-I-ECT: Dissecting Graph Data with Local Euler Characteristic Transforms



Figure 5: A comparison of the squared L^2 distances of the ECTs of aligned and non-aligned wedged spheres, respectively. We see that alignment results in a median loss of zero, thus effectively showing that the two spaces are the same.



Figure 6: A comparison of two wedged spheres with one being rotated around the wedge point and the points being perturbed by Gaussian noise (left) and the learned re-rotated sphere that is aligned with the original data (right).

A.8. Spatial Alignment of Wedged Spheres

We approximate the optimization problem from Eq. (7) to show that we can learn a spatial alignment of two data spaces, while the distance between ECTs of non-aligned spaces that only differ up to a rotation will generally be high. We start with a so-called *wedged sphere*, meaning two 2-dimensional spheres which are concatenated at a gluing point (cf. Figure 6 and von Rohrscheidt & Rieck 2023). We use 2000 uniformly-sampled points from such a wedged sphere, and compare the squared L^2 loss between the ECTs of this sample and a rotation of the same data space. We repeat this procedure 500 times, where at each step both the sample of the wedged sphere and the rotation matrix which yields the rotated version of the same space are sampled randomly. We notice that the L^2 losses between the non-aligned spaces are high (with a median of around 19), whereas the L^2 losses of the non-aligned spaces are significantly lower, with a median loss close to zero (cf. Figure 5). Moreover, we observe that the ECT of the same space significantly changes when the coordinate system is transformed, which corroborates the necessity of a rotation-invariant metric for the comparison of ECTs. We conclude that an alignment of the ECTs of the two underlying data spaces in fact leads to an alignment of the data spaces itself, as promised by the theoretical results in Section 4.

Robustness Figure 7 and Figure 8 show that the spatial alignment of wedged spheres still works satisfactorily, even in the presence of outliers and noise. This property is an important feature when dealing with real-world data, which is often noisy, and enables us to align spaces that only *approximately* differ up to a rotation. By contrast, the Hausdorff distance, i.e., a widely-used metric between point clouds is (by definition) *highly sensitive* to outliers. We therefore conclude that the proposed metric based on ECTs is a robust metric to compare point clouds of potentially different cardinalities.



Figure 7: A comparison of two wedged spheres, with one being rotated around the wedge point and added 200 outliers (left) and the learned re-rotated sphere that is aligned with the original data (right).



Figure 8: A comparison of two wedged spheres, with one being rotated around the wedge point and the points being perturbed by *Gaussian noise* (left) and the learned re-rotated sphere that is aligned with the original data (right).