

Too Many or Too Few? Sampling Bounds for Topological Descriptors

Brittany Terese Fasy* Maksym Makarchuk* Samuel Micka†

David L. Millman‡

November 18, 2025

1 Abstract

Topological descriptors, such as the Euler characteristic function and the persistence diagram, have grown increasingly popular for representing complex data. Recent work showed that a carefully chosen set of these descriptors encodes all of the geometric and topological information about a shape in $\mathbb{R}^d$. In practice, ϵ -nets are often used to find samples in one of two extremes. On one hand, making strong geometric assumptions about the shape allows us to choose ϵ small enough (corresponding to a high-enough density sample) in order to guarantee a faithful representation, resulting in over-sampling. On the other hand, if we choose a larger ϵ in order to allow faster computations, this leads to an incomplete description of the shape and a discretized transform that lacks theoretical guarantees. In this work, we investigate how many directions are really needed to represent geometric simplicial complexes, exploring both synthetic and real-world datasets. We provide constructive proofs that help establish size bounds and an experimental investigation giving insights into the consequences of over- and under-sampling.

2 Introduction

Topological data analysis (TDA) provides tools for extracting complex features from complex data. Representing shapes using sets of topological descriptors, such as the persistence diagram and the Euler characteristic function, has drawn substantial interest in recent years; see [37, 13, 21, 3, 18, 12, 6, 30]. More generally, several projects have found success in using topological descriptors as

*Montana State University, Bozeman, MT, USA

†Western Colorado University, Gunnison, CO, USA

‡Blocky, Bozeman, MT, USA

features for machine learning algorithms and statistical frameworks; see [12, 23, 31, 26, 25, 10, 34].

Much of the success of applying techniques from TDA comes from the properties of the transformations that take data and transform it into a set of topological descriptors. One fundamental question is *do these topological transforms help accurately capture the data?* [30] argue that, because TDA techniques are theoretically grounded, there is a strong relationship between the representation of a shape and the shape itself.

In this paper, we study the topological descriptors known as the persistent homology transform (PHT) and the Euler characteristic transform (ECT), and finite representations of them. The PHT (or ECT) is a function that maps a simplicial complex in $\mathbb{R}^{d+1}$ to the family of persistence diagrams (Euler characteristic functions, respectively) parameterized by the set of directions $\mathbb{S}^d$. [37] showed that this map is injective for simplicial complexes in $\mathbb{R}^2$ and $\mathbb{R}^3$.

We turn to finite representations or discretizations of the PHT by parameterizing the family of descriptors by $\Delta \subseteq \mathbb{S}^d$ instead of $\mathbb{S}^d$ itself. Thus, another natural question is: *how small can Δ be?* That is, how many directional persistence diagrams or Euler characteristic functions are necessary to fully represent a shape? [19] established an upper bound on the size of the set, which is exponential in the dimension of the simplicial complex, with a slight improvement for graphs embedded in $\mathbb{R}^d$. In Section 4, we investigate lower bounds for the size of the descriptor set of several different transforms.

Recent results show that there are finite, faithful representations of topological transforms when assumptions are made on the curvature of the underlying shape [13], when using the verbose descriptors of simplicial complexes in $\mathbb{R}^d$ [3, 4, 19], or when using grayscale images and cubical complexes [6]. However, the resulting sets are often too large to be used in practice, as they are exponential in dimension (unless additional assumptions are made). In Section 5, we investigate how large these sets become in practice on synthetic and real-world data sets.

Alternatively, one could use a coarse discretization of the sphere to create a finite representation. For example, [23] found success on shape recognition tasks by selecting 16 directions from $\mathbb{S}^1$. For each direction s , they computed one persistence diagram of a height or lower-star filtration in direction s . Unfortunately, as we explore in Section 5, the simple discretization schemes would not necessarily capture all of the geometry of the simplicial complex. In Section 6, we provide bounds on the error of the discretization in terms of the Hausdorff distance.

3 Background and Preliminary Results

We explore topological descriptors that are created from lower-star filtrations of (finite) simplicial complexes. In this section, we give a brief background. We defer to [28, Ch. 1] for a rigorous introduction to homology and to [15, Ch. VII] for an introduction to filtrations and the computation of persistent homology.

The Data: Embedded Simplicial Complexes Throughout this paper, we assume $d \in \mathbb{N} := \{0, 1, \dots\}$ and we consider geometric simplicial complexes in $\mathbb{R}^{d+1}$ (i.e., simplicial complexes embedded in $\mathbb{R}^{d+1}$). Letting K be a simplicial complex, we use K_k to denote the set of k -simplices in K . We call elements of K_0 vertices and K_1 edges. We assume that the simplicial complexes are not geometrically degenerate. Because we refer to this assumption often, we state it formally:

Assumption 1 (General Position). *Let $V \subset \mathbb{R}^{d+1}$ be a finite set. We say that V is in general position if no three points are collinear and no two points share a coordinate value. We say a geometric simplicial complex is in general position if its vertices are in general position.*

The Descriptors: Invariants of Filtrations Simplicial complexes themselves are difficult to compare [33]. Instead, researchers turn to using descriptors of the data for summarization and comparison. We focus on descriptors that summarize changes in homology of a filtered topological space. A *filtration* of a geometric simplicial complex K in $\mathbb{R}^{d+1}$ is a sequence of subcomplexes starting at the empty set and ending with K ; in particular, given a direction $s \in \mathbb{S}^d$ and threshold $t \in [0, \infty]$, the lower-level set of K with respect to s at threshold t is:

$$\text{LLS}_{K,s}(t) := \{\sigma \in K \mid \max_{v \in \sigma} s \cdot v \leq t\} \quad (1)$$

We often refer to the threshold t as the *height* parameter. Noticing that, for all $t \leq t'$, the lower-level sets nest $\text{LLS}_{K,s}(t) \subseteq \text{LLS}_{K,s}(t')$, the *lower-star filtration* of K with respect to direction s is the following parameterized family of simplicial complexes:

$$\text{Filt}_K(s) := \{\text{LLS}_{K,s}(t)\}_{t \in \mathbb{R}}. \quad (2)$$

We say that two directional filtrations, $\text{Filt}_K(s)$ and $\text{Filt}_{K'}(s')$, are equivalent if $K = K'$ and s and s' see the vertices in the same order (but perhaps at different heights); we denote this as $\text{Filt}_K(s) \cong \text{Filt}_{K'}(s')$.

For a lower-star filtration, we observe that topological changes only happen at the heights of vertices (i.e., when $t = s \cdot v$ for some $v \in K_0$). However, while every topological change happens at the height of a vertex, not every vertex witnesses a topological change. For example, consider the simplicial complex K shown in Figure 1, and let $s = e^{i\pi/2}$ (the y -direction). The vertices are seen in the following order: v_1 , then v_2 , then v_3 . For ϵ small enough, $\text{LLS}_{K,s}(s \cdot v_2 + \epsilon) = \{v_1, v_2, [v_1, v_2]\}$ deformation retracts onto $\text{LLS}_{K,s}(s \cdot v_2 - \epsilon) = \{v_1\}$, which means that the map $\text{LLS}_{K,s}(s \cdot v_2 - \epsilon) \hookrightarrow \text{LLS}_{K,s}(s \cdot v_2 + \epsilon)$ induces an isomorphism on homology (i.e., there is no topological change at the height of v_2).

We define a **topological descriptor** of a filtration to be any topological invariant of the filtration. Examples of topological descriptor types include persistence diagrams [16, Ch. 7], persistence landscapes [9], Betti functions [38], and Euler characteristic functions [36, 39]. The *persistence diagram* is a graded multiset of points in the extended plane, where each point is graded, or labeled, with a dimension. The points correspond to homological features (colloquially




Figure 1: Example of a simplicial complex with three vertices, three edges, and one triangle. In the y -direction, the corresponding filtration sees three distinct simplicial complexes: $\{v_1\} \subset \{v_1, v_2, [v_1, v_2]\} \subset \{v_1, v_2, v_3, [v_1, v_2], [v_1, v_3], [v_2, v_3], [v_1, v_2, v_3]\}$. However, the lower-star filtration in the y -direction only sees one topological change, when a new connected component is introduced at the height of v_1 .

described as connected components, tunnels, and voids in low dimensions) that are present for an interval of heights. The endpoints of that interval are the coordinates of the point in the persistence diagram. The persistence diagram is a complete, discrete invariant of a filtration, up to homology. In particular, from the persistence diagram alone, at each height t , we know the homology of $\text{LLS}_{K,s}(\mathbf{t})$. Naturally, functional summaries of persistence diagrams are also topological descriptors, and provide access to the tools of functional data analysis [12, 5]. For example, the persistence landscape (defined in [9]) is a function $L: \mathbb{R} \times \mathbb{Z} \rightarrow \mathbb{R}$ and is in one-to-one correspondence with persistence diagrams, but has the advantage that means are well-defined, and, in fact, can be taken pointwise (however, those means might not map back to a persistence diagram). For each $i \in \mathbb{N}$, the *Betti function* $\beta_i: \mathbb{R} \rightarrow \mathbb{R}$ assigns, for each $t \in \mathbb{R}$, the rank of the i -th homology group, $H_i(\text{LLS}_{K,s}(\mathbf{t}))$. The *Euler characteristic function* $\chi: \mathbb{R} \rightarrow \mathbb{R}$ maps each height to the Euler characteristic of the corresponding lower-level set: $\chi(t) = \sum_{i=0} (-1)^i \beta_i(t)$. The Euler characteristic function and a variant of it called the smooth Euler characteristic function are used by diverse groups of researchers today; see [12, 7].

Topological descriptors are invariants of the *filtration*, not of the *underlying shape* itself. For example, chirality is not captured by a single descriptor. However, given a descriptor in all directions, the infinite set of descriptors (labeled with their directions) contains more information. A **directional topological transform** of type $\mathcal{D}$ (simply, a $\mathcal{D}$ -transform in this paper) of a simplicial complex K is the following family of topological descriptors:

$$\text{TT}^{\mathcal{D}}(K) := \{(\mathbf{s}, \mathcal{D}_K(\mathbf{s})) \mid \mathbf{s} \in \mathbb{S}^d\}, \quad (3)$$

where $\mathcal{D}_K(s)$ is the topological descriptor of K of type $\mathcal{D}$ in direction s .

For many common topological descriptor types, the directional topological transform is a complete invariant of the shape. We call this set of descriptor types the **fundamental descriptor types**; in particular, we note that the persistence diagram and the Euler characteristic function are fundamental descriptor types.

While directional topological transforms use an infinite set of descriptors, a finite set of well-chosen directions often suffices to represent the transform.

Definition 2 (Discrete Topological Transform). *Let $\mathcal{D}$ be a topological descriptor type and let $\Delta \subset \mathbb{S}^d$ be a discrete set. We call Δ a set of directions. Then, we define $\text{DTT}_{\Delta}^{\mathcal{D}}$ as the function from the space of geometric simplicial complexes in $\mathbb{R}^{d+1}$ to a space of topological descriptors of type $\mathcal{D}$ parameterized by Δ :*

$$\text{DTT}_{\Delta}^{\mathcal{D}}(K) := \{(\mathbf{s}, \mathcal{D}_K(\mathbf{s})) \mid \mathbf{s} \in \Delta\}. \quad (4)$$

Equivalently, $\text{DTT}_{\Delta}^{\mathcal{D}}(K)$ is a set of direction / topological descriptor pairs. We call $\text{DTT}_{\Delta}^{\mathcal{D}}(K)$ faithful if it contains enough information to recover K .

In order to quantify how much of a simplicial complex is represented by a set of topological descriptors, we define *observable* and *θ -observable*, generalizing [13, Defs. 7.2 and 7.3].¹

Definition 3 (Observable). *Let K be a geometric simplicial complex in $\mathbb{R}^{d+1}$. Let $v \in K_0$ and $s \in \mathbb{S}^d$. Then, we say that v is observable from s if there is a topological change at v in direction s .*

Definition 4 (θ -Observable). *Let K be a geometric simplicial complex in $\mathbb{R}^{d+1}$ and $v \in K_0$. We say that v is θ -observable if there exists a direction $s_0 \in \mathbb{S}^d$ such that v is observable from s_0 and for all s such that $\cos^{-1}(s_0 \cdot s) \leq \theta$ (i.e., the angle between s_0 and s is at most θ).*

For a vertex $v \in K_0$, the largest value of θ such that v is θ -observable relates the “flatness” of the simplicial complex at v . To represent a simplicial complex with a set of topological descriptors, a necessary condition is that each vertex is observed. Hence, we define:

Definition 5 (Observing Region). *Let K be a geometric simplicial complex in $\mathbb{R}^{d+1}$ and $v \in K_0$. The observing region of v is:*

$$\text{obs}_K(v) := \{s \in \mathbb{S}^d \mid v \text{ is observable from } s\}. \quad (5)$$

The *size* of the observing region is its Lebesgue measure [27]. Furthermore, the observing regions are related to a stratification of $\mathbb{S}^d$.

Definition 6 (Coarse Stratification). *Let K be a geometric simplicial complex in $\mathbb{R}^{d+1}$. The coarse stratification of $\mathbb{S}^d$ is the partition induced by the equivalence relation $\simeq$, where, for all $s_1, s_2 \in \mathbb{S}^d$ in the same strata,*

$$s_1 \simeq s_2 \iff \text{order}_{K_0}(\mathbf{s}_1) = \text{order}_{K_0}(\mathbf{s}_2), \quad (6)$$

where $\text{order}_K(\mathbf{s})$ denotes the partial order of vertices in K_0 with respect to direction s .

¹We generalized Euler observable from [13] to simply observable by any topological descriptor. Letting $\delta = \sqrt{2 - 2 \cos \frac{\theta}{2}}$, what we define as θ -observable implies ‘at least’ δ -observable, as defined in [13].

By this definition, knowing the vertex set is sufficient for knowing the coarse stratification. Furthermore, we note that this definition uses only a partial order for an arbitrary direction (as opposed to a total order), as there are directions where two or more vertices are at the same height. The cells whose partial order is a total order are the d -dimensional cells.

Corollary 7 (Relation Between K and θ). *Let K be a geometric simplicial complex in $\mathbb{R}^{d+1}$. The diameter of the smallest d -dimensional stratum in the coarse stratification of K is equal to the smallest angle between vectors with endpoints in the vertex set. Furthermore, if $\theta := \min\{\angle u, v, w \mid uvw \in K_0\}$, then all vertices in K_0 are θ -observable.*

In [13, Theorem 7.14], a representative persistence diagram (or Euler characteristic function) from each d -cell of the coarse stratification, along with the coarse stratification itself, are used as a finite representation of the PHT or ECT. In fact, this sufficiency result holds for all fundamental descriptor types with the same proof as in [13, Theorem 7.14], which allows us to use stratifications in order to compare sets of simplicial complexes:

Lemma 8 (Sufficient Faithful Set). *Let K be a geometric simplicial complex in $\mathbb{R}^{d+1}$, $\mathcal{D}$ a fundamental descriptor type, and $P \subseteq \mathbb{S}^d$ be a set that hits all the highest-dimensional strata of the coarse stratification. Then, $\text{DTT}_{\Delta}^{\mathcal{D}}(K)$ along with the coarse stratification is faithful.*

Leveraging this lemma, we know that in order to distinguish different shapes, it would suffice to select a descriptor in each strata.

Theorem 9 (Sufficient Comparison Set). *Let $\mathcal{D}$ be a fundamental descriptor type, with $\delta: \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}$ a distance metric.² Let $\mathcal{K}$ be a set of geometric simplicial complexes in $\mathbb{R}^{d+1}$. Let $P \subset \mathbb{S}^d$ be a finite set that hits every d -cell of the coarse stratifications for each $K \in \mathcal{K}$. Then, the following function is a distance metric:*

$$d_{\mathcal{K}}: \mathcal{K} \times \mathcal{K} \rightarrow \mathbb{R} \tag{7}$$

$$(K, K') \mapsto \sum_{s \in P} \delta(\mathcal{D}_K(s), \mathcal{D}_{K'}(s)). \tag{8}$$

Proof. Non-negativity, symmetry, the triangle inequality, and the fact that, for all $K \in \mathcal{K}$, $d_{\mathcal{K}}(K, K) = 0$ follow directly from δ being a distance metric.

To show the final property (positivity), let $K \neq K' \in \mathcal{K}$ and suppose, by contradiction, that $d_{\mathcal{K}}(K, K') = 0$. Because P hits all of the d -cells of the coarse stratification for K , by Lemma 8, there must exist some direction $s \in P$ such that $\text{Filt}_{\mathcal{K}}(s) \not\cong \text{Filt}_{\mathcal{K}'}(s)$. Because δ is a metric, we then know that $\delta(\mathcal{D}_K(s), \mathcal{D}_{K'}(s)) > 0$. Thus, by the definition of $d_{\mathcal{K}}$, we conclude $d_{\mathcal{K}}(K, K') > 0$. $\square$

²Here, we slightly abuse notation and let $\mathcal{D}$ denote both the descriptor type and the set of all descriptors of that type.

One way to apply Theorem 9 is to define a set of directions for a fixed database of shapes. Let $\mathcal{K} = \{K_1, K_2, \dots, K_n\}$ be a set of geometric simplicial complexes. For $K_i \in \mathcal{K}$, let ϵ_i be the size of the minimal strata in the coarse stratification of K_i . Let $\epsilon = \min_i \epsilon_i$ and let Δ be ϵ -net over $\mathbb{S}^d$. First, observe that for each $K_i \in \mathcal{K}$ and $\mathcal{D}$ a fundamental descriptor type, the transform $\text{DTT}_\Delta^{\mathcal{D}}(K_i)$ produces a set of descriptors that uniquely represents the shape. Second, observe that Δ satisfies the assumptions on the direction set of Theorem 9, allowing us to compare our complexes. Thus, understanding the minimal stratum size is particularly relevant.

4 Lower Bounds on Representation

In this section, we provide lower bounds on the number of topological descriptors necessary to guarantee that the observing region of each vertex is hit by at least one descriptor. This lower bound grows at least linearly with the size of the vertex set, as the observing region for a degree-two vertex grows arbitrarily small as it approaches the affine space spanned by its two adjacent vertices.

Lemma 10 (Observing Regions Witness Local Max and Min). *Let K be a geometric simplicial complex in $\mathbb{R}^{d+1}$. Let $v \in K_0$ be a degree-two vertex with adjacent vertices u and w . Define the following two sets:*

$$\begin{aligned} P_1 &:= \{s \in \mathbb{S}^d \mid s \cdot u < s \cdot v \text{ and } s \cdot w < s \cdot v\}, \\ P_2 &:= \{s \in \mathbb{S}^d \mid s \cdot u > s \cdot v \text{ and } s \cdot w > s \cdot v\}. \end{aligned}$$

Then, $\text{obs}_K(v) = P_1 \cup P_2$.

Proof. Let $s \in \text{obs}_K(v)$. Because v is degree two, when looking at v in direction s , we encounter one of three cases: (1) v is a local maximum in direction s . This happens if and only if $s \in P_1$; (2) v is seen as a local minimum in direction s . This happens if and only if $s \in P_2$; (3) one edge goes into v from below and the other edge starts at v and goes up with respect to direction s . In this case, v is not observable from s . Hence, we conclude $\text{obs}_K(v) = P_1 \cup P_2$. $\square$

Consider the simplicial complex in Figure 2 comprised of five triangles, along with their edges and vertices. This complex has the additional feature that no two edges are parallel. As a result, we imagine moving the vertices v_2, v_5, v_8, v_{11} , and v_{14} so close to the complementary edge of their respective triangles that the three edges defining the triangle are nearly parallel. We generalize this idea as follows:

Theorem 11 (Linear Size Descriptor Set). *There exist simplicial complexes that require at least a linear number of topological descriptors (with respect to the number of vertices) for unique representation.*

Proof. We give an example of such a simplicial complex. Let $n \geq 3$ be a multiple of three. Let $\mathcal{I} = \{2, 5, 8, \dots, 3n - 1\}$. For each $i \in \mathcal{I}$, define the following two




Figure 2: A simplicial complex with 15 vertices. For any fundamental descriptor type, at least five topological descriptors are needed to uniquely represent this simplicial complex; in general, there exist configurations of n simplices that require $\Omega(n_0)$ topological descriptors for a faithful representation.

vertices:

$$v_{i-1} := (i-1, i-1) \text{ and } v_{i+1} := (i+1, i+1 + \frac{1}{i}). \quad (9)$$

Let $\delta = \frac{1}{3n}$, and let ℓ_i be the line that goes through v_{i-1} and v_{i+1} . Let $i, j \in \mathcal{I}$ such that $j \neq i$. Then, by construction, the difference of slope between ℓ_i and ℓ_j is at least δ . By an application of the intermediate value theorem, we choose an $\epsilon > 0$ so that the point $v_i = (i, i + \frac{1}{2i} + \epsilon)$ is chosen such that the slopes of lines through v_{i-1} and v_i and through v_i and v_{i+1} differ by at most $\frac{\delta}{2}$. Then, we define a simplicial complex as follows:

$$\begin{aligned} K_0 &:= \{[v_i]\}_{i \in \{1, 2, \dots, 3n\}} \\ K_1 &:= \{[v_{i-1}, v_i], [v_i, v_{i+1}], [v_{i-1}, v_{i+1}]\}_{i \in \mathcal{I}} \\ K_2 &:= \{[v_{i-1}, v_i, v_{i+1}]\}_{i \in \mathcal{I}} \end{aligned}$$

No higher dimensional simplices are in K .

In order to represent K , each vertex needs to be observed. In particular, for each $i \in \mathcal{I}$, we must use a topological descriptor from a direction sampled from $\text{obs}_K(v_i)$. Let $i, j \in \mathcal{I}$. Then, by choice of ϵ , $\text{obs}_{v_i}(K) \cap \text{obs}_{v_j}(K) = \emptyset$. In other words, each such v_i has a disjoint observing region, so we need at least $\Omega(n)$ topological descriptors to represent K . $\square$

We note here that, if the i -th vertex is not observed, then K cannot be distinguished from the simplicial complex that removes the i -th vertex and its cofaces (i.e., replaces the i -th triangle with a line segment).

5 Experimental Investigation

In this section, we perform a series of experiments. First, we use an ϵ -net for sampling to highlight the issue of oversampling in geometry-based approaches. Then, we examine another commonly used technique in the field: uniformly sampling a fixed number of directions. Our goal is to analyze how many of the

observing regions are missed as the number of vertices increases. At the end of this section, we demonstrate how the theoretical lower bounds on representation work in practice. All code is available in a public git repo.³

5.1 Data

In our experiments, three planar data sets: random point clouds, the MPEG7 dataset [32], and a subset of the EMNIST dataset (which is an extension of the MNIST dataset) [11]. For the MPEG7 and MNIST datasets, we describe a simple but reasonable set of preprocessing steps for a TDA pipeline based on contours, similar to that found in [1, 2, 20].

RANDPTS The point cloud dataset, named RANDPTS, consists of varying point set sizes in which the points are drawn i.i.d. from the uniform distribution over the $[0, 10]^2$ box. In particular, for each k in 3, 5, 10, 20, $\dots$, 100, we generate 100 point clouds of size k .

MPEG7 and EMNIST The MPEG7 dataset contains 70 classes of binary images of shapes, such as animals. The EMNIST by class dataset contains 62 classes of grayscale images of handwritten characters; we use the first 100 images in each class.

For these images, we extract the boundary from the object in the image such that the vertices conform to Assumption 1 and the boundary is a simple polygon. We simplify the boundaries at two levels using tools from OpenCV [24] and NetworkX [22], which results in a total of four sets of graphs named MPEG7._{.001}, EMNIST._{.001}, MPEG7._{.005}, and EMNIST._{.005}, where the subscript corresponds to the simplification parameter. Specifically, first, we convert to a binary image based on a threshold. In MPEG7, images are black and white, so we use a threshold of zero. In EMNIST, we compute a global threshold and apply it to all images. To compute the global threshold, we find the optimal threshold of the first image in each class using Otsu’s algorithm [29], then we set the global threshold as the average of these thresholds. The global threshold was 102.951 with a standard deviation of 7.631. Second, for each binary image, we compute one contour curve at two simplification levels as follows. We first apply [35] to extract a set of contours and keep the contour with the longest arc length. Afterwards, we simplify each contour using Douglas–Peucker [14]. Given a parameter $\epsilon > 0$, Douglas–Peucker iteratively simplifies a curve by replacing a subcurve with a line segment, provided that the change in Hausdorff distance is smaller than ϵ . For a contour with arc length s , we produce the two simplifications by setting $\epsilon \in .001s, .005s$. Third, because the contours have vertices that lie on a grid, and Assumption 1 states that x - and y -coordinates are distinct, we perturb each coordinate by a value chosen uniformly from $[-.01, .01]$. Fourth, we clean the data by dropping graphs that contain three collinear vertices or are not a simplicial complex. For example, when simplifying contours in the second

³Code available at <https://github.com/compTAG/topo-descriptor-experiments>

step, a vertex removal may cause the contour to self-overlap. More sophisticated simplification algorithms, such as [17], avoid overlaps. We also removed erroneous data, such as one MPEG7 image with inverted colors. After generating and cleaning the data, we convert each contour to a graph and build our output sets. After the preprocessing and pruning, MPEG7_{.001} has 1387 graphs, EMNIST_{.001} has 5563 graphs, MPEG7_{.005} has 1364 graphs, and EMNIST_{.005} has 5751 graphs.

5.2 Too Many: Geometric Based Discretization

By Theorem 9, one way to build a faithful descriptor set is to find a finite sample P such that every stratum in the coarse stratification is “hit” by a point in P . If we know the size of the smallest stratum in a coarse stratification, then an alternate approach is to use an ϵ -net.

In this experiment, we investigate the relationship between the size of the smallest one-stratum (measured in radians) and the number of vertices n_0 as it increases across the datasets RANDPTS, MPEG7_{.001}, and EMNIST_{.001}. We show that creating such an ϵ -net results in too many descriptors to be analyzed. We depict the results as a scatter plot on a log-log scale with best-fit lines in Figure 3. Specifically, we examine how the size of the smallest one-stratum ($\log(m)$) changes with the logarithm of the number of vertices ($\log(n_0)$).




Figure 3: Log-log plot of the smallest stratum size versus the number of vertices for datasets RANDPTS, EMNIST_{.001}, and MPEG7_{.001}.

Across all three datasets, the smallest stratum size decreases proportionally as the number of vertices increases. Even for graphs with thousands of vertices, the minimum stratum size drops below 10^{-5} , making uniform discretization impractical due to the excessive number of descriptors it would generate. The best-fit regression lines confirm a strong negative correlation between the smallest one-stratum size and the number of vertices:

- RANDPTS, $\log(m) = 3.17515 - 4.28122 \log(n_0)$.
- EMNIST_{.001}, $\log(m) = -1.5887 - 3.5964 \log(n_0)$.




Figure 4: Log-log plot of the smallest stratum size versus the number of vertices for dataset EMNIST_{.005}.

- MPEG7_{.001}, $\log(m) = -0.89784 - 3.55102 \log(n_0)$.

In all three experiments, as the number of vertices increases, we see smaller strata appearing with more drastic changes for smaller values of n_0 . The observation produces a particularly pessimistic view for trying to discretize a transformation using the minimum observing region of the dataset. Indeed, we already see that for relatively small graphs with 1000's of vertices, the minimum stratum size drops below 10^{-5} , which would produce descriptor sets containing hundreds of thousands of descriptors, far too large to be of any use in practice.

Additional Considerations Because the contour approximations remove vertices, it is natural to ask whether a coarser approximation yields different dependencies between m and n_0 . To test this, we consider the dataset EMNIST_{.005}. The resulting model still had a negative coefficient for $\log(n_0)$, and the results are depicted in Figure 4. We note that, similar to the main experiments of this section, a different coarsening does not substantially decrease the size of a descriptor set created by uniformly discretizing based on the minimal observing region.

5.3 Too Few: Constant Size Discretization

We demonstrated that sampling directions from every cell of the ϵ -net yields a computationally intensive volume of diagrams. Such an approach is impractical for analyzing real-world data. Instead of sampling based on the geometry of the data, it is common to use a fixed-size set of directions. But how does it affect the results of the analysis?

To address this question, we conducted an experiment to examine how much information about a simplicial complex can be captured using a fixed-size direction set. Given a simplicial complex K , its coarse stratification S_K , and a

discretization $\Delta = \{s_1, s_2, \dots, s_n\} \subset \mathbb{S}^1$, a stratum $k \in S_K$ is considered *hit* by the discretization if $|\{s_i \in \Delta : s_i \in k\}| > 0$.

In this experiment, we fix a discretization of $\mathbb{S}^1$ and study the proportion of strata from coarse stratifications that are hit as we vary n_0 . Specifically, we investigate how the number of captured strata from the coarse stratification changes as the number of vertices within the simplicial complex varies.

For our experiments, we choose $2^{14} = 16384$ directions for the discretization, which is much larger than is common in practice. For example, [37] uses a discretization of 64 directions, and [12] use 72 directions. We picked such a large set of directions to study the best-case scenario for a constant discretization, as in practice, any reasonable discretization would be coarser and miss a higher proportion of strata.

It is worth noting that for a graph with n_0 vertices, the number of strata in the coarse stratification is $N = 2^{\binom{n_0}{2}}$. To ensure that there was ever a chance of hitting every stratum, we only include graphs from our datasets whose coarse stratification contains fewer than 5000 strata. The resulting sizes of the datasets were RANDPTS containing 900 graphs, EMNIST_{.001} with 5543 graphs, and MPEG7_{.001} reduced to 521 samples. By utilizing scatterplots, we visualize the relationship between the number of hits and the number of vertices in the simplicial complex.




Figure 5: Plot of the ratio of hit stratum over the total number of strata versus the number of vertices for RANDPTS, EMNIST_{.001}, and MPEG7_{.001}.

We select 16384 uniformly distributed directions Δ from $\mathbb{S}^1$. For each graph, we compute the coarse stratification and compute the proportion of strata hit by Δ . We observe that in all cases, even for small graphs with 60 or more vertices, we miss more than 10% of the strata, implying that if we want to use a small set of descriptors, we must be willing to miss some strata.

In summary, across all datasets (RANDPTS, EMNIST_{.001}, and MPEG7_{.001}), using uniform sampling to hit each stratum becomes less effective as the number of vertices increases, with varying degrees of correlation strength.

- RANDPTS: $\log(m) = 0.956464 - 1.347025 \log(n_0) - 0.314594 \log^2(n_0)$

- EMNIST_{.001}: $\log(m) = 1.017 - 0.006265 \log(n_0)$
- MPEG7_{.001}: $\log(m) = 1.0141388 - 0.0027408 \log(n_0)$

5.4 Small Graph Experiment




Figure 6: Bozeman, MT OSMnx Street Network


Using too many directions ensures accuracy, but makes computation impractical due to the large number of descriptors. On the other hand, selecting directions without considering the ϵ -net for simplicial complexes risks missing important topological features. This raises a key question: how can we effectively select directions for discretizing topological descriptors?

Our goal in this experiment is to demonstrate that, in practice, only a small number of directions are needed to distinguish real-world data. We consider plane graphs as our dataset. A *plane graph* is a graph embedded in the plane (i.e., a graph G with a bijective continuous map $G \rightarrow \mathbb{R}^2$).

Figure 6 depicts the street network of Bozeman, MT, obtained from Python’s OSMnx package [8]. We extract graphs from this network using bounding boxes of dimensions 60×60 , 80×80 , and 30×30 to obtain subgraphs with four, five, and six vertices, respectively. For each sampled graph, we generate all possible plane subgraphs. We restrict our study to graphs with at most six vertices because, as the number of vertices increases, enumerating all possible plane subgraphs becomes computationally intense.

At the start, we consider a subgraph of the Bozeman road network and generate all possible plane subgraphs from it. We create a coarse stratification

associated with the main subgraph. From this stratification, we randomly select a stratum and then a random direction within that stratum, denoted d_0 . Using d_0 , we compute the directional persistence diagram for every plane subgraph. We compare these diagrams to determine whether the original Bozeman subgraph can be uniquely identified from its plane subgraphs. If the identification is not unique—that is, if multiple plane subgraphs yield identical persistence diagrams—we select another random stratum (different from previously chosen ones) and a random direction within it. We then use persistence diagrams from both directions to attempt unique identification. This process continues by sampling additional random directions from different strata until the original subgraph can be uniquely identified from its set of plane subgraphs.



(a) Subgraphs on 4 Vertices (b) Subgraphs on 5 Vertices (c) Subgraphs on 6 Vertices

Figure 7: For each subgraph, we randomly sample from a discretization of $\mathbb{S}^1$ and add equally spaced directions until no other plane graphs exist with the same set of persistence diagrams.

As depicted in Figure 7, majority of the subgraphs sampled required only three directions. Generally, we find that between two and four directions are sufficient; however, we observed a small number of road networks that require a larger number of directions. These networks are depicted in Figure 8. As Assumption 1 suggests, all the networks that require more directions have (near) collinearities, depicted as white vertices in the figure.




Figure 8: Graphs with 4, 5, and 6 nodes that require more than four directions.

The final step involves plotting a histogram depicting the number of subgraphs versus the number of directions required. The standard distribution

of this histogram serves as the average number of directions from the coarse stratification needed to uphold the faithfulness of the descriptor set.

6 Implications of Too Few

Our experiments show that using *too many* directions can cause oversampling, while using *too few* directions often fails to capture the full complexity of data. Balancing the two, researchers often err on the side of too few, for what is the worth of a descriptor set if there are too many descriptors to realistically compute? Thus, we address one final question: *how far apart can two simplicial complexes be if their topological descriptors are the same on a fixed set of directions?*

Theorem 12 (Lost Vertex). *Let K be a simplicial complex such that the following conditions are met:*

- (a) *The vertices are in general position (Assumption 1),*
- (b) *There exist distinct vertices u, v, w such that $\|u - v\| = \|w - v\|$,*
- (c) *The open ball $\mathbb{B}(v, \|w - v\|) \cap K = \{v, (u, v), (v, w)\}$.*

Let $\Delta \subset \mathbb{S}^d \setminus \text{obs}_K(v)$. Let $\mathcal{K}_\Delta$ denote the set of geometric simplicial complexes in $\mathbb{R}^d$ such that, for all $s \in \Delta$, the corresponding filtrations are equivalent: $\text{Filt}_K(\mathbf{s}) \cong \text{Filt}_{\tilde{K}}(\mathbf{s})$. Let $\theta = \pi - \angle uvw$. Then, there exists a simplicial complex $\tilde{K} \in \mathcal{K}_\Delta$ such that:

$$d_H(K, \tilde{K}) = \|v - w\| \cos \frac{\theta}{2} > 0. \quad (10)$$

Proof. Let $\tilde{K} = \{(u, w)\} \cup K \setminus \{v, (u, v), (v, w)\}$. For any $s \notin \text{obs}_K(v)$, we know that we always see the three named vertices in one of two orders: u before v before w , or w before v before u . This property, and by the definition of $\tilde{K}$, we know that $\text{Filt}_K(\mathbf{s}) \cong \text{Filt}_{\tilde{K}}(\mathbf{s})$ for all $s \in \Delta$. Thus, $\tilde{K} \in \mathcal{K}_\Delta$.

Let $\pi(v)$ denote the orthogonal projection of v onto the edge (u, w) . By the definition of K and by Condition (c), we know that the Hausdorff distance between K and $\tilde{K}$ is as follows:

$$d_H(K, \tilde{K}) = d_H(v, \tilde{K}) = d_H(v, (u, w)). \quad (11)$$

In other words, the Hausdorff distance between K and $\tilde{K}$ is the Euclidean distance between v and $\pi(v)$.

By trigonometric properties of isosceles triangles, we know that $\|v - \pi(v)\| = \|v - w\| \cos \frac{\theta}{2}$.

Thus, we conclude:

$$d_H(K, \tilde{K}) = \|v - w\| \cos \frac{\theta}{2}. \quad (12)$$

Finally, because $\theta = \pi - \angle uvw$ and by Condition (a), we know that $\theta \in (0, \pi)$. Furthermore, because $v \neq w$ by Condition (b), we obtain $\|v - w\| \cos \frac{\theta}{2} > 0$. $\square$

7 Discussion

The experiments in this paper are all in $\mathbb{R}^2$. They could, of course, be extended to $\mathbb{R}^3$ or even $\mathbb{R}^d$, as all of the theory and definitions allow for arbitrary dimensions. However, in this paper, we sought to find simple, yet common, data examples that demonstrate how bad the behavior of under- or over-sampling directions could be. Extending these experiments to $\mathbb{R}^3$ or higher would add computational complexity, as data structures defined over $\mathbb{S}^1$ are relatively simple, but over $\mathbb{S}^2$ require more nuances. The observed behavior would get worse in higher dimensions (e.g. Figure 8c).

While, in theory, faithful representations of shape need many directions, in practice, smaller sets of directions seem to suffice. Of course, loss of information is expected in data summaries, and the impact of that loss depends on the particular application and the machine learning or statistical tools used. For example, in cluster analysis, classification in (small) descriptor space is much easier than classification in (full) shape space, and approximate clustering is often acceptable. We hope that this paper provides some insight into how much information about data is lost with the smaller sets used, so that, when used in applications, scientists better understand the limits of these descriptors.

Acknowledgments The authors thank Luke Padula and James Wilson for their assistance with implementing prototypes of the experiments and preprocessing data. We also thank Lorin Crawford for conversations around the use of topological transforms in practice.

This work was partially supported by the National Science Foundation grants DMS 1854336 and CCF 2046730.

References

- [1] Xiang Bai, Wenyu Liu, and Tu Zhuowen. Integrating contour and skeleton for shape classification. In *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*, pages 360–367, 2009.
- [2] Serge Belongie, Jitendra Malik, and Jan Puzicha. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(4):509–522, 2002.
- [3] 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. In *Canadian Conference on Computational Geometry*, U. Manitoba, Winnipeg, Canada, 2018.
- [4] Robin Lynne Belton, Brittany Terese Fasy, Rostik Mertz, Samuel Micka, David L. Millman, Daniel Salinas, Anna Schenfisch, Jordan Schupbach, and Lucia Williams. Reconstructing embedded graphs from persistence diagrams. 90:101658, 2020.

- [5] Eric Berry, Yen-Chi Chen, Jessi Cisewski-Kehe, and Brittany Terese Fasy. Functional summaries of persistence diagrams. *Journal of Applied and Computational Topology*, 4:211–262, 2020.
- [6] Leo M. Betthauser. *Topological Reconstruction of Grayscale Images*. PhD thesis, University of Florida, 2018.
- [7] Omer Bobrowski and Shmuel Weinberger. On the vanishing of homology in random čech complexes. *Random Structures & Algorithms*, 51(1):14–51, 2017.
- [8] Geoff Boeing. Modeling and analyzing urban networks and amenities with osmnx. *Geographical Analysis*, 2025. Published online ahead of print.
- [9] Peter Bubenik. Statistical topological data analysis using persistence landscapes. *The Journal of Machine Learning Research*, 16(1):77–102, 2015.
- [10] Mathieu Carrière, Frédéric Chazal, Yuichi Ike, Théo Lacombe, Martin Royer, and Yuhei Umeda. A general neural network architecture for persistence diagrams and graph classification. *arXiv preprint arXiv:1904.09378*, 2019.
- [11] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and André van Schaik. Emnist: an extension of MNIST to handwritten letters. In *International Joint Conference on Neural Networks (IJCNN)*, pages 2921–2926. IEEE, 2017.
- [12] 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.
- [13] 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(32):1006–1043, 2022.
- [14] David H. Douglas and Thomas K. Peucker. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *Cartographica: the International Journal for Geographic Information and Geovisualization*, 10(2):112–122, 1973.
- [15] Herbert Edelsbrunner and John Harer. *Computational Topology: An Introduction*. American Mathematical Society, 2010.
- [16] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 454–463. IEEE, 2000.

- [17] Regina Estkowski and Joseph S. B. Mitchell. Simplifying a polygonal subdivision while keeping it simple. In *Proceedings of the Seventeenth Annual Symposium on Computational Geometry*, pages 40–49. Association for Computing Machinery, 2001.
- [18] Brittany Terese Fasy, Samuel Micka, David L. Millman, Anna Schenfisch, and Lucia Williams. Challenges in reconstructing shapes from Euler characteristic curves. In *Fall Workshop on Computational Geometry*, 2018. Also available at arXiv:1811.11337.
- [19] Brittany Terese Fasy, Samuel Micka, David L. Millman, Anna Schenfisch, and Lucia Williams. A faithful discretization of verbose directional transform, 2019. Preprint available at arXiv:1912.12759.
- [20] Pedro F. Felzenszwalb and Joshua D. Schwartz. Hierarchical matching of deformable shapes. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007.
- [21] Robert Ghrist, Rachel Levanger, and Huy Mai. Persistent homology and Euler integral transforms. *Journal of Applied and Computational Topology*, 2:55–60, 2018.
- [22] Aric Hagberg, Pieter Swart, and Daniel S. Chult. Exploring network structure, dynamics, and function using NetworkX. Technical report, Los Alamos National Lab. (LANL), 2008.
- [23] Christoph D Hofer, Roland Kwitt, and Marc Niethammer. Learning representations of persistence barcodes. *Journal of Machine Learning Research*, 20(126):1–45, 2019.
- [24] Itseez. Open source computer vision library. <https://github.com/itseez/opencv>, 2015.
- [25] Aditi S. Krishnapriyan, Maciej Haranczyk, and Dmitriy Morozov. Topological descriptors help predict guest adsorption in nanoporous materials. *The Journal of Physical Chemistry C*, 124(17):9360–9368, 2020.
- [26] Peter Lawson, Jordan Schupbach, Brittany Terese Fasy, and John W Shepard. Persistent homology for the automatic classification of prostate cancer aggressiveness in histopathology images. In *Medical Imaging 2019: Digital Pathology*, volume 10956, page 109560G. International Society for Optics and Photonics, 2019.
- [27] Henri Lebesgue. *Intégrale, Longueur, Aire*. PhD thesis, Université de Paris, 1902.
- [28] James R. Munkres. *Elements of Algebraic Topology*. CRC Press, [1930] 2018.

- [29] Nobuyuki Otsu. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1):62–66, 1979.
- [30] Steve Oudot and Elchanan Solomon. Inverse problems in topological persistence. In *Topological Data Analysis: The Abel Symposium 2018*, pages 405–433. Springer, 2020.
- [31] Talha Qaiser, Yee-Wah Tsang, Daiki Taniyama, Naoya Sakamoto, Kazuaki Nakane, David Epstein, and Nasir Rajpoot. Fast and accurate tumor segmentation of histology images using persistent homology and deep convolutional features. *Medical Image Analysis*, 55:1–14, 2019.
- [32] Richard Ralph. Mpeg-7 core experiment ce-shape-1 test set, accessed 2019.
- [33] Matthew Ricci, Junkyung Kim, and Thomas Serre. Same-different problems strain convolutional neural networks. *Annual Meeting of the Cognitive Science Society*, 2018. arXiv preprint arXiv:1802.03390.
- [34] Soheil Rostami, Walid Saad, and Choong Seon Hong. Deep learning with persistent homology for orbital angular momentum (oam) decoding. *IEEE Communications Letters*, 2019.
- [35] Satoshi Suzuki and Keiichi Abe. Topological structural analysis of digitized binary images by border following. *Computer Vision, Graphics, and Image Processing*, 30:32–46, 1985.
- [36] Sergio Torres. Topological analysis of COBE-DMR cosmic microwave background maps. *Astrophysical Journal, Part 2-Letters*, 423(1):L9–L12, 1994.
- [37] Katharine Turner, Sayan Mukherjee, and Doug M. Boyer. Persistent homology transform for modeling shapes and surfaces. *Information and Inference: A Journal of the IMA*, 3(4):310–344, 2014.
- [38] Rien Van De Weygaert, Gert Vegter, Herbert Edelsbrunner, Bernard J. T. Jones, Pratyush Pranav, Changbom Park, Wojciech A. Hellwing, Bob Eldering, Nico Kruithof, E. G. Patrick Bos, Johan Hidding, Job Feldbrugge, Eline ten Have, Matti van Engelen, Manuel Caroli, and Monique Teillaud. Alpha, betti and the megaparsec universe: on the topology of the cosmic web. *Transactions on Computational Science XIV: Special Issue on Voronoi Diagrams and Delaunay Triangulation*, pages 60–101, 2011.
- [39] Michael S. Vogeley, Changbom Park, Margaret J. Geller, John P. Huchra, and J. Richard Gott III. Topological analysis of the CfA redshift survey. *The Astrophysical Journal, Part 1*, 420(2):525–544, 1994.