

Topological Data Analysis with ϵ -net Induced Lazy Witness Complex

Naheed Anjum Arafat¹, Debabrota Basu², Stéphane Bressan¹

¹ School of Computing, National University of Singapore, Singapore

² Data Science and AI Division, Chalmers University of Technology, Göteborg, Sweden

Abstract. Topological data analysis computes and analyses topological features of the point clouds by constructing and studying a simplicial representation of the underlying topological structure. The enthusiasm that followed the initial successes of topological data analysis was curbed by the computational cost of constructing such simplicial representations. The lazy witness complex is a computationally feasible approximation of the underlying topological structure of a point cloud. It is built in reference to a subset of points, called landmarks, rather than considering all the points as in the Čech and Vietoris-Rips complexes. The choice and the number of landmarks dictate the effectiveness and efficiency of the approximation.

We adopt the notion of ϵ -cover to define ϵ -net. We prove that ϵ -net, as a choice of landmarks, is an ϵ -approximate representation of the point cloud and the induced lazy witness complex is a 3-approximation of the induced Vietoris-Rips complex. Furthermore, we propose three algorithms to construct ϵ -net landmarks. We establish the relationship of these algorithms with the existing landmark selection algorithms. We empirically validate our theoretical claims. We empirically and comparatively evaluate the effectiveness, efficiency, and stability of the proposed algorithms on synthetic and real datasets.

1 Introduction

Topological data analysis computes and analyses topological features of generally high-dimensional and possibly noisy data sets. Topological data analysis is applied to eclectic domains namely shape analysis [6], images [2, 22], sensor network analysis [8], social network analysis [3, 25, 27, 24], computational neuroscience [21], and protein structure study [30, 19].

The enthusiasm that followed the initial successes of topological data analysis was curbed by the computational challenges posed by the construction of an exact simplicial representation, the Čech complex, of the point cloud. A simplicial representation facilitates computation of basic topological objects such as simplicial complexes, filtrations, and persistent homologies. Thus, researchers devised approximations of the Čech complex as well as its best possible approximation the Vietoris-Rips complex [7, 26, 10]. One of such computationally feasible approximate simplicial representation is the *lazy witness complex* [7]. The lazy witness complex is built in reference to a subset of vertices, called *landmarks*. The choice and the number of landmarks dictate the effectiveness and efficiency of the approximation.




Fig. 1: Components of topological data analysis.

We adopt the notion of ϵ -cover [18] from analysis to define and present the notions of ϵ -sample, ϵ -sparsity, and ϵ -net (Section 4) to capture bounds on the loss of topological features induced by the choice of landmarks. We prove that an ϵ -net is an ϵ -approximate representation of the point cloud with respect to the Hausdorff distance. We prove that the lazy witness complex induced by an ϵ -net, as a choice of landmarks, is a 3-approximation of the induced Vietoris-Rips complex. ϵ -net allows us to provide such approximation guarantees for lazy witness complex (Section 4.2) which was absent in the literature. Furthermore, we propose three algorithms to construct ϵ -net as landmarks for point clouds (Section 5). We establish their relationship with the existing landmark selection algorithms, namely *random* and *maxmin* [7]. We empirically and comparatively show that the size of the ϵ -net landmarks constructed by the algorithms varies inversely with ϵ and agrees with the known bound on the size of ϵ -net [20].

We empirically and comparatively validate our claim on the topological approximation quality of the lazy witness complex induced by the ϵ -net landmarks (Section 6). Furthermore, we empirically and comparatively validate the effectiveness, efficiency and stability of the proposed algorithm on representative synthetic point clouds as well as a real dataset. Experiments confirm our claims by showing equivalent effectiveness of the algorithms constructing ϵ -net landmarks with the existing maxmin algorithm. We also show the ϵ -net landmarks to be more stable than those selected by the algorithms maxmin and random as ϵ -net incurs narrower confidence band in the persistent landscape topological descriptor. We conclude (Section 7) with the theoretical and experimental pieces of evidence that validate the ϵ -nets constructed as a stable and effective way to construct landmarks and to induce lazy witness complexes.

2 Topological Data Analysis

Topological data analysis is the study of computational models for efficient and effective computation of topological features, such as persistent homology classes, from different datasets, and representation of the topological features using different topological descriptors, such as persistence barcodes, for further analysis and application [12, 23]. In this section, we represent the computational blocks of topological data analysis in Figure 1 and further describe each of the blocks.

Topological data analysis computes the topological features, such as persistent homology classes, by computing the topological objects called *simplicial complex* for a given dataset. A **simplicial complex** is constructed using simplices. Formally, a k -*simplex* is the convex-hull of $(k + 1)$ data points. For instance, A 0-simplex $[v_0]$ is a single point, a 1-simplex $[v_0v_1]$ an edge, and a

2-simplex $[v_0v_1v_2]$ a filled triangle. A **k -homology class** is an equivalent class of such k -simplicial complexes that cannot be reduced to a lower dimensional simplicial complex [12]. In order to compute the k -homology classes, a practitioner does not have direct access to the underlying space of the point cloud and it is combinatorially hard to compute the exact simplicial representation of Čech complex [31]. Thus, different approximations of the exact simplicial representation are proposed: *Vietoris-Rips* complex [17] and *lazy witness* complex [7].

The **Vietoris-Rips complex** $R_\alpha(D)$, for a given dataset D and real number $\alpha > 0$, is an abstract simplicial complex representation consisting of such k -simplices, where any two points v_i, v_j in any of these k -simplices are at distance at most α . Vietoris-Rips complex is the best possible ($\sqrt{2}$ -)approximation of the Čech complex, computable with present computational resources, and is extensively used in topological data analysis literature [23]. Thus, we use the Vietoris-Rips complex as the baseline representation in this paper. In the worst case, the number of simplices in the Vietoris-Rips complex grows exponentially with the number of data points [23, 31]. Lazy witness complex [7] approximates the Vietoris-Rips complex by constructing the simplicial complexes over a subset of data points L , referred to as the landmarks. Formally, given a positive integer ν and real number $\alpha > 0$, the **lazy witness complex** $LW_\alpha(D, L, \nu)$ of a dataset D is a simplicial complex over a landmark set L where for any two points v_i, v_j of a k -simplex $[v_0v_1 \cdots v_k]$, there is a point w whose $(d^\nu(w) + \alpha)$ -neighborhood³ contains v_i, v_j . In the worst case, the size of the lazy witness complexes grows exponentially with the number of landmarks. Less number of landmarks facilitates computational acceleration while produces a bad approximation of Vietoris-Rips with loss of topological features. Thus, the trade-off between the approximation of topological features and available computational resources dictates the choice of landmarks. We provide a quantification on such loss of topological features (Section 4.2) which was absent in the literature.

As the value of filtration parameter α increases, new simplices arrive and the topological features, i.e. the homology classes, start to appear. Some of the homology classes merge with the existing classes in a subsequent simplicial complex, and some of them persist indefinitely [12]. In order to capture the evolution of topological structure with scale, topological data analysis techniques construct a sequence of simplicial complex representations, called a *filtration* [12], for an increasing sequence of α 's. In a given filtration, the persistence interval of a homology class is denoted by $[\alpha_b, \alpha_d)$, where α_b and α_d are the filtration values of its appearance and merging respectively. The persistence interval of an indefinitely persisting homology class is denoted as $[\alpha_b, \infty)$.

Topological descriptors, such as barcodes [6], persistence diagram [12], and persistence landscapes [1], represent persistence intervals as stacked lines, points and functions respectively in order to draw qualitative and quantitative inference about the topological features. Distance measures between persistent diagrams such as the bottleneck and Wasserstein distances [12] are often used to draw

³ $d^\nu(w)$ is the distance from point $w \in D$ to its ν -th nearest point in L .

quantitative inference. The bottleneck distance between two diagrams is the smallest distance d for which there is a perfect matching between the points of the two diagrams such that any pair of matched points are at distance at most d [12]. The Wasserstein distance between two diagrams is the cost of the optimal matching between points of the two diagrams [12].

3 Related Work

Applications of TDA. TDA is applied in different domains mostly on relatively small datasets and up to dimension 2 due to computational intractability of the popular Čech and Vietoris-Rips complexes. [6] computed homology classes at dimension 0 for their proposed tangential filtration of point clouds of handwritten digits for image classification (dataset size ~ 69 – 294). [28] used the persistence pairs at dimension 0 for segmenting mesh on benchmark mesh segmentation datasets (size ~ 50000). Researchers applying TDA to network analysis focus on characterising networks using features computed from persistence homology classes. [3] and [25] computed persistence homology at dimension 0,1 and 2 of the clique filtration to study weighted collaboration networks (size ~ 36000) and weighted networks from different domains (size ~ 54000) respectively. In biological networks, [11] clustered gene co-expression networks (size ~ 400) based on distances between Vietoris-Rips persistence diagram computed on each network. In molecular biology, persistent homology reveals different conformations of proteins [30, 19] based on the strength of the bonds of the molecules.

Topological Approximation. Computational infeasibility of the Čech complex and Vietoris-Rips complex motivates the development of approximate simplicial representations such as the lazy witness complexes, sparse-Rips complex [26] and graph induced complex (GIC) [10]. Sparse-Rips complex [26] perturbs the distance metric in such a way that when the regions covered by a point can be covered by its neighbouring points, that point can be deleted without changing the topology. Given a graph constructed on a point cloud as input, the graph induced complex is a simplicial complex built on a subset of vertices, where the vertices of a k -simplex are the nearest neighbours of a clique-vertices of a k -clique [10]. Due to their computational benefits, lazy witness complex and graph induced complexes have found applications in studying natural image statistics [7] and image classification [9].

Applications of ϵ -net. The concept of ϵ -net is a standard concept in analysis and topology [18] originating from the idea of (δ, ϵ) -limits formulated by Cauchy. Nets have been used in nearest-neighbor search [20]. [15] proposed the Net-tree data structure to represent ϵ -nets at all scales of ϵ . Net-tree is used to construct approximate well-separated pair decompositions [15] and approximate geometric spanners [15]. The simplicial complexes in the graph induced complex are nets. Sparse-Rips filtration constructs a net-tree on the point-cloud to decide which neighbouring points to delete. [14] used ϵ -net for manifold reconstruction.

4 ϵ -net

As we discussed in Section 2, topological data analysis of a dataset begins with the computation of simplicial complex representations. Though Vietoris-Rips is

the best possible approximation of the Čech complex, it incurs an exponential computational cost with respect to the size of the point cloud. Thus, lazy witness complex is often used as a practical solution for scalable topological data analysis. Computation of lazy witness complex is dependent on the selection of landmarks. Selection of landmarks dictates the trade-off between *effectiveness*, i.e. the quality of approximation of topological features, and *efficiency*, i.e. the computational cost of computing the lazy witness complex.

ϵ -cover is a construction used in topology to compute the inherent properties of a given space [18]. In this paper, we import the concept of ϵ -cover to define ϵ -net of a point cloud. We use the ϵ -net of the point cloud as the landmarks for constructing lazy witness complex. We show that ϵ -net, as a choice of landmarks, has guarantees such as being an ϵ -approximate representation of the point cloud, its induced lazy witness complex being a 3-approximation of its induced Vietoris-Rips complex, and also bounding the number of landmarks for a given ϵ . These guarantees are absent for the other existing landmark selection algorithms (Section 5) such as random and maxmin algorithms.

4.1 ϵ -net of a Point Cloud

ϵ -cover is a set of subsets of a point cloud in an Euclidean space such that these subsets together cover the point cloud, but none of the subsets has a diameter more than ϵ .

Definition 1 (ϵ -cover [18]). An ϵ -cover of a point cloud P is the set of P_i 's such that $P_i \subseteq P$, $P = \cup_i P_i$, and diameter^A of P_i is at most $\epsilon \geq 0$ for all i .

When the sets in the 2ϵ -cover of P are Euclidean balls of radius ϵ , the set of centres of the balls is termed as an ϵ -sample of set P .

Definition 2 (ϵ -sample [14]). A set $L \subseteq P$ is an ϵ -sample of P if the collection $\{B_\epsilon(x) : x \in L\}$ of ϵ -balls of radius ϵ covers P , i.e. $P = \cup_{x \in L} B_\epsilon(x)$.

According to the definition of ϵ -sample, P is an ϵ -sample of itself for $\epsilon = 0$. For decreasing further computational expense, it is desirable to have an ϵ -sample is sparse that means it contains as less number of points as possible. An ϵ -sparse subset of P is a subset where any two points are at least ϵ apart from each other.

Definition 3 (ϵ -sparse). A set $L \subset P$ is ϵ -sparse if for all $x, y \in L$, $d(x, y) > \epsilon$.

An ϵ -net of set P is an ϵ -sparse subset of P which is also an ϵ -sample of P .

Definition 4 (ϵ -net [18]). Let (P, d) be a metric space and $\epsilon \geq 0$. A subset $L \subset P$ is called an ϵ -net of P if L is ϵ -sparse and an ϵ -sample of P .

4.2 Properties of ϵ -nets

ϵ -net of a point cloud comes with approximation guarantees irrespective of its algorithmic construction. An ϵ -net of a point cloud of diameter Δ in Euclidean space \mathbb{R}^D is an ϵ -approximation of the point cloud in Hausdorff distance. The lazy

^A The diameter $diam(P_i)$ of a set $P_i \subseteq P$ is defined as the largest distance $d(x, y)$ between any two points in $x, y \in P_i$.

witness complex induced by an ϵ -net is a 3-approximation of the Vietoris-Rips complex on that ϵ -net. Furthermore, the size of an ϵ -net is at most $(\frac{\Delta}{\epsilon})^{\theta(D)}$ [20]. Here, we establish the first two approximation guarantees of ϵ -net theoretically.

Point-cloud Approximation Guarantee of an ϵ -net. We use Lemma 1 to prove that the ϵ -net of a point cloud P is an ϵ -approximate representation of that point cloud in Hausdorff distance.

Lemma 1. *Let L be an ϵ -net of point cloud P . For any point $p \in P$, there exists a point $q \in L$ such that the distance $d(p, q) \leq \epsilon$*

Theorem 1. *The Hausdorff distance between the point cloud P and its ϵ -net $L \subseteq P$ is at most ϵ .*

Proof. For any $l \in L$, there exists a vertex $p \in P$ such that $d(l, p) \leq \epsilon$, by definition of $B_\epsilon(l)$. Hence, $\min_{l \in L} d(l, p) \leq \epsilon$, and thus, $\max_{p \in P} \min_{l \in L} d(l, p) \leq \epsilon$. For any $p \in P$, there exists a landmark $l \in L$ such that $d(l, p) \leq \epsilon$, by Lemma 1. Thus, $\max_{l \in L} \min_{p \in P} d(l, p) \leq \epsilon$. Hence the Hausdorff distance $d_H(P, L)$ between P and L , defined as the maximum of $\max_{l \in L} \min_{p \in P} d(l, p)$ and $\max_{p \in P} \min_{l \in L} d(l, p)$ is bounded by ϵ .

Topological Approximation Guarantee of an ϵ -net Induced Lazy Witness Complex. In addition to an ϵ -net being an ϵ -approximation of the point-cloud, we prove that the lazy witness complex induced by the ϵ -net landmarks is a good approximation (Theorem 2) to the Vietoris-Rips complex on the landmarks. This approximation ratio is independent of the algorithm constructing the ϵ -net. As a step towards Theorem 2, we state Lemma 2 that follows from the definition of the lazy witness complex and ϵ -sample. Lemma 2 establishes the relation between 1-nearest neighbour of points in an ϵ -net.

Lemma 2. *If L is an ϵ -net landmark of point cloud P , then the distance $d(p, p')$ from any point $p \in P$ to its 1-nearest neighbour $p' \in P$ is at most ϵ .*

Theorem 2 shows that the lazy witness complex induced by the landmarks in an ϵ -net is a 3-approximation of the Vietoris-Rips complex on those landmarks above the value 2ϵ of filtration parameter.

Theorem 2. *If L is an ϵ -net of the point cloud P for $\epsilon \in \mathbb{R}^+$, $LW_\alpha(P, L, \nu = 1)$ is the lazy witness complex of L at filtration value α , and $R_\alpha(L)$ is the Vietoris-Rips complex of L at filtration α , then $R_{\alpha/3}(L) \subseteq LW_\alpha(P, L, 1) \subseteq R_{3\alpha}(L)$ for $\alpha \geq 2\epsilon$.*

Proof. In order to prove the first inclusion, consider a k -simplex $\sigma_k = [x_0 x_1 \cdots x_k] \in R_{\alpha/3}(L)$. For any edge $[x_i x_j] \in \sigma_k$, let w_t be the point in P that is nearest to the vertices of $[x_i x_j]$ and wlog, let that vertex be x_j . Since w_t is the nearest neighbour of x_j , by Lemma 2, $d(w_t, x_j) \leq \epsilon \leq \frac{\alpha}{2}$. Since $[x_i x_j] \in R_{\alpha/3}$, $d(x_i, x_j) \leq \frac{\alpha}{3} < \frac{\alpha}{2}$. By triangle inequality, $d(w_t, x_i) \leq \frac{\alpha}{2} + \frac{\alpha}{2} \leq \alpha$. Hence, x_i is within distance α from w_t . The α -neighbourhood of point w_t contains both x_i and x_j . Since $d^1(w_t) > 0$, the $(d^1(w_t) + \alpha)$ -neighbourhood of w_t also contains x_i, x_j . Therefore, $[x_i x_j]$ is an edge in $LW_\alpha(P, L, 1)$. Since the argument is true for any $x_i, x_j \in \sigma_k$, the k -simplex $\sigma_k \in LW_\alpha(P, L, 1)$.

In order to prove the second inclusion, consider a k -simplex $\sigma_k = [x_0 x_1 \cdots x_k] \in LW_\alpha(P, L, 1)$. Therefore, by definition of lazy witness complex, for any edge

$[x_i x_j]$ of σ_k there is a witness $w \in P$ such that, the $(d^1(w) + \alpha)$ -neighbourhood of w contains both x_i and x_j . Hence, $d(w, x_i) \leq d^1(w) + \alpha \leq \epsilon + \alpha$ (by Lemma 2) $\leq 3\alpha/2$. Similarly, $d(w, x_j) \leq 3\alpha/2$. By triangle inequality, $d(x_i, x_j) \leq 3\alpha$. Therefore, $[x_i x_j]$ is an edge in $R_{3\alpha}(L)$. Since the argument is true for any $x_i, x_j \in \sigma_k$, the k -simplex $\sigma_k \in R_{3\alpha}(L)$.

Discussion. Theorem 2 implies that the interleaving of lazy witness filtration $LW = LW_\alpha(L)$ and the Vietoris-Rips filtration $R = R_\alpha(L)$ occurs when $\alpha > 2\epsilon$. As a consequence, the weak-stability theorem [4] implies that the bottleneck distance between the partial persistence diagrams $Dgm_{>2\epsilon}(LW)$ and $Dgm_{>2\epsilon}(R)$ are 3-approximations of each other in log-scale, and their bottleneck distance is upper bounded by $3 \log 3$. In Section 6, we empirically validate this bound for the lazy witness complex induced by the ϵ -net landmarks.

Size of an ϵ -net. The size of an ϵ -net depends on ϵ , the diameter of the point-cloud and the dimension of the underlying space [20, 15]. If a point cloud $P \subset \mathbb{R}^D$ has diameter Δ , the size of an ϵ -net of P is $(\frac{\Delta}{\epsilon})^{\theta(D)}$ [20]. The size of an ϵ -net does not depend on the size of the point cloud. In Section 6, we empirically validate this bound for the ϵ -net landmarks generated by the proposed algorithms. The framework of ϵ -net along with its approximation guarantees lead to the question of its algorithmic construction as landmarks.

5 Construction of an ϵ -net

The naïve algorithm [16] to construct an ϵ -net selects the first point l_1 uniformly at random. In i -th iteration, it marks the points at a distance less than ϵ from the previously selected landmark l_{i-1} as covered, and selects the new point l_i from the unmarked points arbitrarily until all points are marked [15]. The fundamental principle is to choose, at each iteration, a new landmark from the set of yet-to-cover points such that it retains the ϵ -net property. We propose three algorithms where this choice determines the algorithm.

5.1 Three Algorithms: ϵ -net-rand, ϵ -net-maxmin, and $(\epsilon, 2\epsilon)$ -net

The algorithm **ϵ -net-rand**, at each iteration, marks the points at a distance less than ϵ from the previously chosen landmark as covered and chooses a new landmark uniformly at random from the unmarked points. The algorithm **ϵ -net-maxmin**, at each iteration, marks the points at a distance less than ϵ from the previously chosen landmark as covered and chooses the farthest unmarked point from the already chosen landmarks. It terminates when the distance to the farthest unmarked point is no more than ϵ . The algorithm **$(\epsilon, 2\epsilon)$ -net**, at each iteration, marks the points at a distance less than ϵ from the previously chosen landmark as covered, and chooses a landmark uniformly at random from those unmarked points whose distance to the previously chosen landmark is at most 2ϵ . If there are no unmarked points at a distance in-between ϵ and 2ϵ from the previous landmark, it searches for unmarked points at a distance between 2ϵ and 4ϵ , 4ϵ and 8ϵ , and so on, until it either finds one to continue as before or all points are marked. The pseudo-code for $(\epsilon, 2\epsilon)$ -net is in Algorithm 1.

$(\epsilon, 2\epsilon)$ -net attempts to cover the point-cloud with intersecting balls of radius ϵ , whereas ϵ -net-maxmin attempts to cover the point-cloud with non-intersecting balls of radius ϵ . ϵ -net-rand does not maintain any invariant.

Algorithm 1 Algorithm $(\epsilon, 2\epsilon)$ -net**Input:** Point cloud $P = \{p_1, p_2, \dots, p_n\}$, $n \times n$ Distance matrix D , parameter ϵ .**Output:** Set of Landmarks L .

- 1: Select the initial landmark l_1 uniformly at random from P .
- 2: Initialize $L = \{l_1\}$.
- 3: Let $N_{(\epsilon, 2\epsilon)}^1$ be the set of points at a distance between ϵ and 2ϵ from l_1 .
- 4: Initialize candidate landmarks $C_1 = N_{(\epsilon, 2\epsilon)}^1$.
- 5: $i = 1$.
- 6: **repeat**
- 7: Let $N_{\leq \epsilon}^i$ be the set of points at a distance less than ϵ from l_i .
- 8: Mark all the points in $N_{\leq \epsilon}^i$ as covered.
- 9: Let C_i^u be the set of unmarked points in C_i .
- 10: **if** C_i^u is empty **then**
- 11: Find the first $\delta = [1, 2, \dots, \log(\lceil \Delta/2\epsilon \rceil)]$ for which $N_{(2^\delta \epsilon, 2^{\delta+1} \epsilon)}^i$ contains any unmarked point.
- 12: Set $C_i = C_i \cup N_{(2^\delta \epsilon, 2^{\delta+1} \epsilon)}^i$.
- 13: **end if**
- 14: Select l_{i+1} uniformly at random from C_i^u .
- 15: Insert l_{i+1} to L .
- 16: $C_{i+1} = C_i \cup N_{(\epsilon, 2\epsilon)}^{i+1}$.
- 17: $i = i + 1$.
- 18: **until** all points are marked

In terms of time-complexity, ϵ -net-rand and $(\epsilon, 2\epsilon)$ -net have the complexity of $O(\frac{1}{\epsilon^D})$ and $O(\frac{1}{\epsilon^D} \log(\frac{1}{\epsilon}))$ respectively. Their run-time does not depend on the size of the input point cloud. On the other hand, the run-time of ϵ -net-maxmin depends on the size of the point-cloud as it has to search for the farthest point from the landmarks at each iteration. If the size of the point cloud is n , ϵ -net-maxmin has $O(\frac{n}{\epsilon^D})$ time-complexity.

5.2 Connecting ϵ -net to Random and Maxmin algorithms

De Silva et al.[7] proposed random and maxmin algorithms for point clouds.

Random. Given a point cloud P , the algorithm *random* selects $|L|$ vertices uniformly at random from the set of points P . This algorithm is closely related to ϵ -nets. Given the number of landmarks $K > 1$, the set of landmarks selected by random is δ -sparse where δ is the minimum of the pairwise distances among the landmarks. However, the same choice of K may not necessarily make the landmarks a δ -sample of the point cloud.

The ϵ -net-rand algorithm is a modification of random that takes ϵ as a parameter instead of K and use ϵ to put a constraint on the domain of random choices. It continues to select landmarks until all points are marked to ensure the ϵ -sample property. The proof sketch of the fact that the constructed landmarks are ϵ -sparse and ϵ -sample is as follows:

Proof. The ϵ -net-rand algorithm does not terminate until all points are marked as covered. Hence the set of landmarks selected by ϵ -net-rand is an ϵ -sample, since otherwise, there would have been unmarked points. The pairwise distance between any two landmarks cannot be less than ϵ ; otherwise, one of them would

have been marked by the other and the marked point would not be a landmark. Hence the set landmarks selected by ϵ -net-rand is ϵ -sparse.

Maxmin. The *maxmin* algorithm selects the first landmark l_1 uniformly at random from the set of points, P . Following that; it selects the point which is furthest to the present set of landmarks at each step till a given number of landmarks, say $|L|$, are chosen. If $L_{i-1} = \{l_1, l_2, \dots, l_{i-1}\}$ is the set of already chosen landmarks, it selects such a point $u \in P \setminus L_{i-1}$ as the i^{th} landmark that maximises the minimum distance from the present set of landmarks L_{i-1} . Mathematically, $l_i \triangleq \arg \max_{u \in P \setminus L_{i-1}} \min_{v \in L_{i-1}} d(u, v)$. The maxmin algorithm selects landmarks such that the point cloud is covered as vastly as possible.

The maxmin algorithm is closely related to ϵ -net. Given the number of landmarks $K > 1$, the set of landmarks selected by maxmin is δ -sparse where δ is the minimum of the pairwise distances among the landmarks chosen. However, that choice of K may not necessarily make the landmarks a δ -sample of the point cloud. The ϵ -net-maxmin algorithm is a modification of maxmin that takes ϵ as a parameter instead of K and uses ϵ to control sparsity among the landmarks. It terminates when the minimum of the pairwise distances among the landmarks drops below ϵ to ensure the ϵ -sample property of the landmarks chosen. The proof sketch of the resulting landmarks being ϵ -sparse and ϵ -sample is as follows:

Proof. The ϵ -net-maxmin algorithm, at each iteration, selects only that point as a landmark whose minimum distance to the other landmark points is the largest among all unmarked points. If such a point's minimum distance to the other landmark points is no more than ϵ , the algorithm terminates. Hence the set of landmarks selected by ϵ -net-maxmin must be ϵ -sparse. A point that is not a landmark must be covered by some landmark point already. Otherwise, its minimum distance to the landmark set would have been at least ϵ , and hence would have been the only unmarked point available to be selected as a new landmark by ϵ -net-maxmin. Therefore the set of landmarks selected by ϵ -net-maxmin is also ϵ -sample of the point cloud.

6 Empirical Performance Evaluation

We implement the pipeline illustrated in Figure 1 to empirically validate our theoretical claims and also the effectiveness, efficiency, and stability of the algorithms that construct ϵ -net landmarks compared to that of the random and maxmin algorithms. We test and evaluate these algorithms on two synthetic point cloud datasets, namely Torus and Tangled-torus, and a real-world point cloud dataset, namely lgrm. On each input point cloud, we compute the lazy witness filtration and Vietoris-Rips filtration induced by the landmarks, as well as the Vietoris-Rips filtration induced by the point cloud.

On each dataset, as we vary parameter ϵ of the algorithms constructing ϵ -nets, we study the relationship between ϵ to the number of landmarks, the quality of the topological features approximated by the lazy witness filtration induced by those landmarks, as well as the stability of those approximated features. As the algorithms maxmin and random require the number of landmarks a priori,




Fig. 2: (left) Torus, (middle) Tangled-torus, and (right) 1grm Dataset for the sake of comparison, we use the same number of landmarks as that of the corresponding ϵ -net algorithm for a given ϵ .

We compute the quality of the features approximated by an algorithm in terms of the 1-Wasserstein distance between the lazy witness filtration induced by the landmarks selected by that algorithm to those of a Vietoris-Rips filtration on the same dataset. As there are elements of randomness in the algorithms, we run each experiment 10 times and compute distances averaged over the runs.

We compute the stability of the features approximated by the algorithms in terms of the 95% confidence band corresponding the rank 1 persistence landscape using bootstrap [5]. We use persistence landscape to validate the stability of the filtrations because unlike persistence diagrams and barcodes, two sets of persistence landscapes always have unique mean and by strong law of large numbers the empirical mean landscapes of sufficiently large collection converge to its expected landscapes [1].

6.1 Datasets and Experimental Setup

Datasets. We use the datasets illustrated in Figure 2 for experimentation. The dataset **Torus** is a point cloud of size 500 sampled uniformly at random from the surface of a torus in \mathbb{R}^3 . The torus has a major radius of 2.5 and minor radius of 0.5. The dataset **Tangled-torus** is a point cloud of size 1000 sampled uniformly at random from two tori tangled with each other in \mathbf{R}^3 . Both tori has a major radius of 2.5 and minor radius of 0.5. The dataset **1grm** is the conformation of the gramicidin-A protein. It has a helical shape. Gramicidin-A has two disconnected chains of monomers consisting of 272 atoms.

Experimental Setup. We implement the experimental workflow in Matlab 2018a (with 80GB memory limit). All experiments are run on a machine with an Intel(R) Xeon(R)@2.20GHz CPU and 196 GB memory. We use the Javaplex library [29] to construct lazy witness filtrations and to compute their persistence intervals. We use the Ripser library to construct the Vietoris-Rips filtrations and to compute their persistence intervals. We use R-TDA package [13] to compute bottleneck and Wasserstein distances, and 95% confidence band for the landscapes. We set the lazy witness parameter $\nu = 1$ in all computations. We set the maximum value of the filtration parameter to the diameter of the dataset in question.

6.2 Validation of Theoretical Claims

Number of Landmarks Generated by the ϵ -net Algorithms. In Figure 3, we illustrate the relation between the number of landmarks generated by the




Fig. 3: Number of landmarks generated by the ϵ -net algorithms vs. ϵ on Torus dataset.

ϵ -net algorithms and ϵ on Torus dataset. Each algorithm is run 10 times for each ϵ , and the mean and standard deviation are plotted. We observe that the number of landmarks decreases as ϵ increases. We also observe that, for a fixed ϵ , the average number of landmarks selected by the ϵ -net algorithms is more or less stable across different algorithms. We use the number of landmarks of an ϵ -net-maxmin to fit a curve with values $\Delta = 5.9$ (the diameter of Torus) and coefficient $\theta(D) = 1.73$ (found from fitting with 95% confidence). This observation supports the theoretical bound of $(\frac{\Delta}{\epsilon})^{\theta(D)}$.






Fig. 4: Topological approximation guarantee of ϵ -net constructed by the algorithms on Torus dataset.

Topological Approximation Guarantee. In order to validate Theorem 2 on dataset Torus, we compute the bottleneck distance between the persistence diagram of the lazy witness filtration and that of the Vietoris-Rips filtration induced by the ϵ -net landmarks for different values of ϵ . For each ϵ and algorithm, we generate 10 sets of ϵ -net landmarks, compute their corresponding persistence diagrams and plot the mean and standard deviation of the bottleneck distances in Figure 4. Since the theorem is valid for $\alpha \geq 2\epsilon$, we exclude the homology classes born below 2ϵ before the distance computation. The algorithms satisfy the bound as the distances are always less than the theoretical bound of $3 \log 3$. Since the plots on the other datasets support these claims, for the sake of brevity, we omit them.

6.3 Effectiveness and Efficiency of Algorithms Constructing ϵ -nets


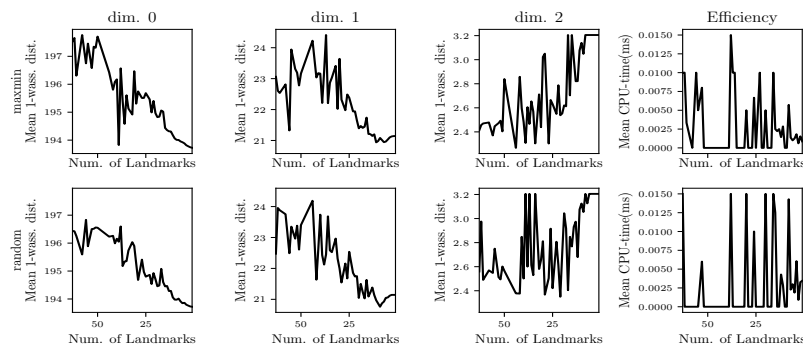
For each ϵ , we compute the 1-Wasserstein distance between the persistent diagrams of the lazy witness filtration induced by each ϵ -net landmarks and that of the Vietoris-Rips filtration induced by the whole point cloud. We compute the mean distance and mean CPU-times across 10 runs. Unlike ϵ -net algorithms, the existing landmark selection algorithms take the number of landmarks as input. Since the average number of landmarks selected by the ϵ -net algorithms does not vary much across different algorithms (Figure 3), for each ϵ , we take the same number of ϵ -net-maxmin landmarks as parameters to select the random and maxmin landmarks. Figure 5 illustrates result on Torus dataset.

We observe that maxmin performs well in dimensions 0 and 2 whereas $(\epsilon, 2\epsilon)$ -net has competitive effectiveness. In dimension 1, we observe that ϵ -net-maxmin achieves the lowest minimum, whereas random achieves the highest minimum. All the ϵ -net algorithms has two local minima, the first of which at around $\epsilon = 0.5$

(a) Effectiveness and Efficiency of the ϵ -net algorithms on Torus dataset.

(b) Effectiveness and Efficiency of the existing algorithms on Torus dataset.

Fig. 5: Torus dataset.

(a) Effectiveness and Efficiency of the ϵ -net algorithms on 1grm dataset.

(b) Effectiveness and Efficiency of the existing algorithms on 1grm dataset.

Fig. 6: 1grm dataset.

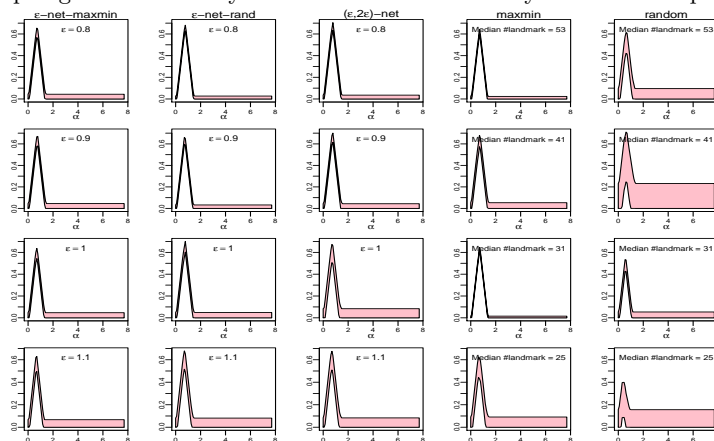


Fig. 7: 95% confidence band of the rank one persistence landscape at dimension 1 of the lazy witness filtration induced by the landmark selection algorithms on Tangled-torus dataset.

and the second in between $\epsilon = 2$ to $\epsilon = 4$. The first local minimum is due to the minor radius. As for the explanation of the second local minimum, it is sufficient to either cover the inner diameter of 5 or the outer diameter of 6 to capture the cycle. A 2.5- to 3-sparse sample suffices to do so. The performance of the maxmin and random landmarks is not as explainable as the ϵ -net landmarks. In terms of efficiency, we observe that that $(\epsilon, 2\epsilon)$ -net algorithm has the lowest run-time among all the ϵ -net algorithms. The $(\epsilon, 2\epsilon)$ -net algorithm has competitive effectiveness and better efficiency among the proposed algorithms. Figure 6 illustrates the results for 1grm dataset. We observe that $(\epsilon, 2\epsilon)$ -net achieves the smallest loss in dimensions 0 and 1. In dimension 2, maxmin achieves the smallest loss. ϵ -net-rand takes the smallest CPU-time among all the ϵ -net algorithms. We observe that the effectiveness of $(\epsilon, 2\epsilon)$ -net and efficiency of ϵ -net-rand in the results on Tangled-torus dataset. We omit the plots due to space limitation.

Despite providing better efficiency and equivalent effectiveness on the datasets under study, the performance of the maxmin algorithm is less predictable and less explainable than the ϵ -net algorithms. Among the ϵ -net algorithms, $(\epsilon, 2\epsilon)$ -net has better effectiveness at the expense of little loss in efficiency, whereas ϵ -net-rand has better efficiency than the others with effectiveness comparable to ϵ -net-maxmin.

6.4 Stability of the ϵ -net Landmarks

In Figure 7, we vary ϵ and plot the rank 1 persistence landscape at dimension 1 and its 95% confidence band corresponding to the lazy witness filtration induced by the different landmark selection algorithms. For maxmin and random, we take the same number of landmarks as that in the corresponding ϵ -net-maxmin landmarks. The rank 1 persistent landscape is a functional representation of the most persistent homology class, which we observe from Figure 7 in the form of a peak for all the algorithms. The x -axis represents the value of filtration parameter and y -axis represents function values. We observe that the ϵ -net-maxmin has similar confidence bands as maxmin, whereas the confidence bands

of ϵ -net-rand and $(\epsilon, 2\epsilon)$ -net are often narrower than both maxmin and random. Random has the widest confidence band among all. The confidence bands of maxmin are in-between these two extremes. This observation implies that the ϵ -net algorithms are more stable than the existing algorithms. We observe similar stability results on other datasets that we omit due to space limitation.

7 Conclusion

We use the notion of ϵ -net to capture bounds on the loss of the topological features of the induced lazy witness complex. We prove that ϵ -net is an ϵ -approximation to the original point cloud and the lazy witness complex induced by ϵ -net is a 3-approximation to the Vietoris-Rips complex on the landmarks for values of filtration parameter beyond 2ϵ . Such quantification of approximation for lazy witness complex was absent in literature and is not derivable for algorithms limiting the number of landmarks.

We propose three algorithms to construct ϵ -net landmarks. We show that the proposed ϵ -net-rand and ϵ -net-maxmin algorithms are variants of the algorithm random and maxmin respectively, which ensures the constructed landmarks to be an ϵ -sample of the point cloud. We empirically and comparatively show that the sizes of the landmarks that our algorithms construct agree with the bound on the size of ϵ -net. We empirically validate our claim on the topological approximation guarantee by showing that beyond 2ϵ filtration value, the bottleneck distances are bounded by $3 \log 3$. Furthermore, we empirically and comparatively validate the effectiveness, efficiency and stability of the proposed algorithms on representative synthetic point clouds as well as a real dataset. Experiments confirm our claims by showing equivalent effectiveness of the algorithms constructing ϵ -net landmarks at the cost of a little decrease in efficiency but offering better stability.

References

1. Bubenik, P.: Statistical topological data analysis using persistence landscapes. *The Journal of Machine Learning Research* 16(1), 77–102 (2015)
2. Carlsson, G., Ishkhanov, T., de Silva, V., Zomorodian, A.: On the local behavior of spaces of natural images. *International Journal of Computer Vision* 76(1), 1–12 (2008)
3. Carstens, C.J., Horadam, K.J.: Persistent Homology of Collaboration Networks. *Mathematical Problems in Engineering* 2013(6), 1–7 (jun 2013)
4. Chazal, F., Cohen-Steiner, D., Glisse, M., Guibas, L.J., Oudot, S.Y.: Proximity of persistence modules and their diagrams. In: *Proceedings of the twenty-fifth annual symposium on Computational geometry*. pp. 237–246. ACM (2009)
5. Chazal, F., Fasy, B., Lecci, F., Michel, B., Rinaldo, A., Wasserman, L.: Subsampling methods for persistent homology. In: *International Conference on Machine Learning*. pp. 2143–2151 (2015)
6. Collins, A., Zomorodian, A., Carlsson, G., Guibas, L.J.: A barcode shape descriptor for curve point cloud data. *Computers & Graphics* 28(6), 881–894 (2004)
7. De Silva, V., Carlsson, G.: Topological estimation using witness complexes. In: *Proceedings of the First Eurographics conference on Point-Based Graphics*. pp. 157–166. Eurographics Association (2004)

8. De Silva, V., Ghrist, R., et al.: Coverage in sensor networks via persistent homology. *Algebraic & Geometric Topology* 7(1), 339–358 (2007)
9. Dey, T., Mandal, S., Varcho, W.: Improved image classification using topological persistence. In: *Proceedings of the conference on Vision, Modeling and Visualization*. pp. 161–168. Eurographics Association (2017)
10. Dey, T.K., Fan, F., Wang, Y.: Graph induced complex on point data. *Computational Geometry* 48(8), 575–588 (2015)
11. Duman, A.N., Pirim, H.: Gene coexpression network comparison via persistent homology. *International journal of genomics* 2018 (2018)
12. Edelsbrunner, H., Harer, J.: *Computational topology: an introduction*. American Mathematical Soc. (2010)
13. Fasy, B.T., Kim, J., Lecci, F., Maria, C.: Introduction to the r package tda. arXiv preprint arXiv:1411.1830 (2014)
14. Guibas, L.J., Oudot, S.Y.: Reconstruction using witness complexes. *Discrete & computational geometry* 40(3), 325–356 (2008)
15. Har-Peled, S., Mendel, M.: Fast construction of nets in low-dimensional metrics and their applications. *SIAM Journal on Computing* 35(5), 1148–1184 (2006)
16. Har-Peled, S., Raichel, B.: Net and prune: A linear time algorithm for euclidean distance problems. *Journal of the ACM (JACM)* 62(6), 44 (2015)
17. Hausmann, J.C., et al.: On the vietoris-rips complexes and a cohomology theory for metric spaces. *Ann. Math. Studies* 138, 175–188 (1995)
18. Heinonen, J.: *Lectures on analysis on metric spaces*. Springer Science & Business Media (2012)
19. Kovacev-Nikolic, V., Bubenik, P., Nikolić, D., Heo, G.: Using persistent homology and dynamical distances to analyze protein binding. *Statistical applications in genetics and molecular biology* 15(1), 19–38 (2016)
20. Krauthgamer, R., Lee, J.R.: Navigating nets: simple algorithms for proximity search. In: *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*. pp. 798–807 (2004)
21. Lee, H., Chung, M.K., Kang, H., Kim, B.N., Lee, D.S.: Discriminative persistent homology of brain networks. In: *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. pp. 841–844. IEEE (2011)
22. Letscher, D., Fritts, J.: Image segmentation using topological persistence. In: *International Conference on Computer Analysis of Images and Patterns*. pp. 587–595. Springer (2007)
23. Otter, N., Porter, M.A., Tillmann, U., Grindrod, P., Harrington, H.A.: A roadmap for the computation of persistent homology. *EPJ Data Science* 6(1), 17 (2017)
24. Patania, A., Petri, G., Vaccarino, F.: The shape of collaborations. *EPJ Data Science* 6(1), 18 (2017)
25. Petri, G., Scolamiero, M., Donato, I., Vaccarino, F.: Topological strata of weighted complex networks. *PloS one* 8(6), e66506 (2013)
26. Sheehy, D.R.: Linear-size approximations to the vietoris-rips filtration. *Discrete & Computational Geometry* 49(4), 778–796 (2013)
27. Sizemore, A., Giusti, C., Bassett, D.S.: Classification of weighted networks through mesoscale homological features. *Journal of Complex Networks* 5(2), 245–273 (2017)
28. Skraba, P., Ovsjanikov, M., Chazal, F., Guibas, L.: Persistence-based segmentation of deformable shapes. In: *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition-Workshops*. pp. 45–52. IEEE (2010)
29. Tausz, A., Vejdemo-Johansson, M., Adams, H.: JavaPlex: A research software package for persistent (co)homology. In: Hong, H., Yap, C. (eds.) *Proceedings of ICMS 2014*. pp. 129–136. *Lecture Notes in Computer Science* 8592 (2014)

30. Xia, K., Wei, G.W.: Persistent homology analysis of protein structure, flexibility, and folding. *International Journal for Numerical Methods in Biomedical Engineering* 30(8), 814–844 (8 2014)
31. Zomorodian, A.: Fast construction of the vietoris-rips complex. *Computers & Graphics* 34(3), 263–271 (2010)