

ϵ -net Induced Lazy Witness Complexes on Graphs

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, Sweden

Abstract. Computation of persistent homology of simplicial representations such as the Rips and the Čech complexes do not efficiently scale to large point clouds. It is, therefore, meaningful to devise approximate representations and evaluate the trade-off between their efficiency and effectiveness. The lazy witness complex economically defines such a representation using only a few selected points, called landmarks.

Topological data analysis traditionally considers a point cloud in a Euclidean space. In many situations, however, data is available in the form of a weighted graph. A graph along with the geodesic distance defines a metric space. This metric space of a graph is amenable to topological data analysis.

We discuss the computation of persistent homologies on a weighted graph. We present a lazy witness complex approach leveraging the notion of ϵ -net that we adapt to weighted graphs and their geodesic distance to select landmarks. We show that the value of the ϵ parameter of the ϵ -net provides control on the trade-off between choice and number of landmarks and the quality of the approximate simplicial representation. We present three algorithms for constructing an ϵ -net of a graph. We comparatively and empirically evaluate the efficiency and effectiveness of the choice of landmarks that they induce for the topological data analysis of different real-world graphs.

1 Introduction

Topological data analysis (TDA) [5, 24] involves computation of topological features of datasets, such as persistent homology classes, and the representation of these topological features using such topological descriptors as persistence barcodes [13]. In this section, we elaborate the computational blocks of topological data analysis as shown in Figure 1.

Simplicial Complex. Topological data analysis computes the topological features of a dataset, such as persistent homology classes, by computing the topological objects called *simplicial complex*. 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 $[v_0v_1]$ is an edge, and a 2-simplex $[v_0v_1v_2]$ is 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 [13].




Fig. 1: Components of topological data analysis.

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 [29] and *lazy witness* complex [10].

Approximate Simplicial Representations. 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 u, v 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 [24]. 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 [31].

Lazy witness complex [10] 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 a 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 $[v_0 v_1 \cdots v_k]$, there is a point w whose $(d^\nu(w) + \alpha)$ -neighbourhood contains v_i, v_j . $d^\nu(w)$ is the geodesic distance from point $w \in L$ to its ν -th nearest point in the landmark set L . 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 that was absent in the literature.*

Filtration and Representation of Topological Features. 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 [13]. 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* [13], 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 ap-

pearance and merging respectively. The persistence interval of an indefinitely persisting homology class is denoted as $[\alpha_b, \infty)$. Topological descriptors, such as barcodes [9], persistence diagram [13], and persistence landscapes [4], represent persistence intervals in order to draw qualitative and quantitative inference about the topological features. Distance measures between persistent diagrams such as the q-Wasserstein and Bottleneck distance [13] are often used to draw quantitative inference.

Graph Topological Data Analysis. Topological data analysis (TDA) [5] traditionally considers a point cloud in a Euclidean space. In many situations, however, data is available in the form of a weighted graph. Conveniently, the vertices of the graph with the geodesic distance define a metric space. This metric space is amenable to topological data analysis. This fact does not depend on whether the graph is embeddable in a Euclidean space or not. Though the metric space induced by a graph provides a different structure than point clouds to investigate, this metric space bears the similar issues of scalable construction of representations as well as similar approximate representations, filtrations, and topological descriptors. This motivated us to exploit the generalisability of topological data analysis and to extend the ϵ -net induced lazy witness complexes [1] to graphs.

Our Contributions. We investigate the computation of persistent homologies on a weighted graph. In Section 3, we present a lazy witness complex approach leveraging the notion of ϵ -net that we adapt to weighted graphs and their geodesic distance to select landmarks. We show that the ϵ parameter of the ϵ -net gives a control on the trade-off between choice and number of landmarks and the quality of the approximate simplicial representation.

In Section 3.2, 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.

In Section 4, we present three algorithms, namely Greedy- ϵ -net, Iterative- ϵ -net, and SPTpruning- ϵ -net, for constructing an ϵ -net of a graph. In Section 5, we comparatively and empirically evaluate the efficiency and effectiveness of the choice of landmarks that they induce for the topological data analysis of several real-world graphs.

In Section 6, we summarise the findings and the future directions of research that ϵ -net opens up for graph topological data analysis.

2 Related Works

Graph TDA. Existing applications of TDA to graphs focus on characterizing networks using features computed from persistence homology classes. [7] 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 biology domain, [12] clustered gene co-expression networks (size ~ 400) based on distances be-

tween Vietoris-Rips persistence diagram computed on each network. [21] studied Vietoris-Rips filtration of the functional brain networks computed on ~ 100 region of interests (points) in human brains with different clinical disorders. They focus on homology classes at dimension 0 and 1. A related line of work regarding the topology on graphs involves graphs derived as a representation of point-cloud data (e.g. the neighbourhood graph) and their usage in data clustering [8] and inference of global topology from local information [28].

Approximate Simplicial Complexes. Computational infeasibility of constructing 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) [11].

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. ϵ -net are sets in a metric space that covers the whole space and are well-separated. Nets have been used in nearest-neighbour search [20]. [15] used ϵ -net for manifold reconstruction. Graph induced complex [11] uses the cliques in the neighbourhood graph to construct simplicial complex over an ϵ -net.

[17] proposed net-tree data structure to represent ϵ -nets at all scales of ϵ . Net-tree is used to construct approximate well-separated pair decompositions [17] and approximate geometric spanners [17]. Sparse-Rips filtration [26] constructs a net-tree on the point-cloud to decide which neighbouring points to delete. Contrary to Sheehy [26], we use ϵ -net to select a fixed subset of points, called landmarks, and compute persistent homology using them.

3 ϵ -net of Graphs: Definition and Analysis

[1] proposes the ϵ -net induced lazy witness complex for a point cloud embedded in a Euclidean space for efficient computation of topological data analysis. In practice, the datasets may not be represented as a point cloud in a Euclidean space. The data may have different representations and non-Euclidean geometry. For instance, the dataset with contextual and relational structure is often represented using graphs. In a graphical representation of data, the vertices represent the data objects, edges represent relations among the data objects, and weights on the edges quantify the amplitude of the relation with respect to others.

In this paper, we study both weighted and unweighted simple graphs. Since unweighted graphs are a special case of weighted graph, we construct the definitions for weighted simple graphs. A weighted simple graph [3] $G(V, E, W)$ is a graph with a vertex set V , an edge set E , a weight function $W : V \times V \rightarrow \mathbf{R}^+$, and does not contain any self-edge or multiple edge. The geodesic distance $d_G(u, v)$ between a pair u, v of vertices in a graph is defined as the length of the shortest path between u and v , where the path length is defined as the sum of weights of the edges connecting the vertices u and v [23]. In this paper we treat a graph $G = (V, E, W)$ as a set V endowed with the canonical metric $d_G : V \times V \rightarrow \mathbf{R}^+$. Thus, a weighted simple graph G transforms into a metric space (V, d_G) .

As we substitute the points in the point cloud with the vertices of the graph and the Euclidean distance between points with the geodesic distance between the vertices, we adapt the components of topological data analysis for simple weighted graphs.

ϵ -cover is a construction used in topology to compute inherent properties of a given space [18]. In this paper, we import the concept of ϵ -cover to define ϵ -net of a graph. We use the ϵ -net of a graph as landmarks for constructing lazy witness complex on that graph.

We show that ϵ -net, as a choice of landmarks, has guarantees such as being an ϵ -approximate representation of the graph, its induced lazy witness complex being a 3-approximation of the corresponding 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 such as random and maxmin algorithms [27].

3.1 Defining ϵ -net of a Graph

In this paper, we consider a graph $G = (V, E, W)$ as a finite metric space (V, d_G) . Since there are many graphs that are neither Euclidean nor they have Euclidean embedding, we extend the definitions of neighbourhood and cover from point-set topology as follows before defining ϵ -net.

Definition 1 (ϵ -neighbourhood). *The ϵ -neighbourhood $\mathcal{N}_\epsilon(u)$ of a vertex $u \in V$ is a subset of the vertex set V such that for any vertex $v \in \mathcal{N}_\epsilon(u)$ the distance $d_G(u, v) \leq \epsilon$.*

The notion of ϵ -cover for graph generalises the geometric notion of cover using the set-theoretic notion of ϵ -neighbourhood in Definition 1.

Definition 2 (ϵ -cover). *An ϵ -cover of G is the finite collection $\{\mathcal{N}_{\epsilon/2}(u_i)\}$ of $\epsilon/2$ -neighbourhoods of vertices in G such that $\cup_i \mathcal{N}_{\frac{\epsilon}{2}}(u_i) = V$.*

By triangle inequality, any set in the ϵ -cover has the property that two vertices u, v in the cover have geodesic distance $d_G(u, v) \leq \epsilon$. We differentiate an ϵ -cover from the set of vertices whose $\frac{\epsilon}{2}$ -neighbourhood determines that cover by defining ϵ -sample.

Definition 3 (ϵ -sample). *A set $L = \{u_1, u_2, \dots, u_{|L|}\} \subseteq V$ is an ϵ -sample of graph G if the collection $\{\mathcal{N}_\epsilon(u_i) : u_i \in L\}$ of ϵ -neighbourhoods covers G i.e. $\cup_i \mathcal{N}_\epsilon(u_i) = V$.*

ϵ -neighbourhoods may intersect, which is not an intended property if we want to decrease the size of the ϵ -sample. We combine the notion of ϵ -cover with the notion of ϵ -sparsity to define ϵ -net.

Definition 4 (ϵ -sparse). *A set $L = \{u_1, u_2, \dots, u_{|L|}\} \subset V$ is ϵ -sparse if for any distinct $u_i, u_j \in L$, $d_G(u_i, u_j) > \epsilon$ in graph G .*

An ϵ -net of graph G is such a subset of V which is both an ϵ -sample of G and ϵ -sparse. The ϵ -net while considered as the landmark set L induces a metric subspace (L, d_L) of the metric space of the graph of (V, d_G) , where d_L is the metric induced on set L by the geodesic metric d_G .

Definition 5 (ϵ -net). *A subset $L \subset V$ is an ϵ -net of G if L is ϵ -sparse and an ϵ -sample of G .*

Relating ϵ -net and Other Graph Theoretic Concepts. The definition of ϵ -net on graphs generalises the notion of independent set and dominating set for undirected graphs [3]. Any 1-net of an undirected graph $G = (V, E)$ is an independent set of G and vice versa. Any minimal cardinality 1-net of G is a dominating set of G , and vice versa.

3.2 Analysing Properties of ϵ -net of a Graph

ϵ -net of a simple, weighted, connected graph comes with approximation guarantees irrespective of its algorithmic construction. In this section, we provide following three analysis of ϵ -net:

1. An ϵ -net of a connected graph is an ϵ -approximation of its set of vertices in Hausdorff distance.
2. The lazy witness complex induced by an ϵ -net is a 3-approximation of the Vietoris-Rips complex induced by the same set.
3. There exists an ϵ -net of a graph of diamtere Δ of size at most $(\frac{\Delta}{\epsilon})^{O(\log \frac{|V|}{\epsilon})}$.

Graph Approximation Guarantee of an ϵ -net. We use Lemma 1 to prove that the ϵ -net of a graph $G = (V, E, W)$ is an ϵ -approximation of V in Hausdorff metric (Theorem 1). Lemma 1 follows from the ϵ -sample property of an ϵ -net.

Lemma 1. *Let L be an ϵ -net of graph G . For any vertex $v \in V$, there exists a point $u \in L \subseteq V$ such that the geodesic distance $d_G(u, v) \leq \epsilon$*

Proof. Since $V = \cup_{u \in L} \mathcal{N}_\epsilon(u)$, for any vertex $v \in V$ there exists an $u \in L$ such that $v \in \mathcal{N}_\epsilon(u)$. As $v \in \mathcal{N}_\epsilon(u)$, by definition of ϵ -neighbourhood, the length of the shortest path from v to u is at most ϵ , i.e. $d_G(u, v) \leq \epsilon$.

Theorem 1. *The Hausdorff distance between (V, d_G) and its ϵ -net induced subspace (L, d_L) is at most ϵ .*

Proof. For any $u \in L \subseteq V$, there exists a vertex $v \in V$ such that $d_G(u, v) \leq \epsilon$, by definition of ϵ -neighbourhood. Hence, $\max_L \min_V d_G(u, v) \leq \epsilon$. By Lemma 1, $\max_V \min_L d_G(u, v) \leq \epsilon$. Since the Hausdorff distance $d_H(V, L)$ is defined as the maximum of $\max_L \min_V d_G(u, v)$ and $\max_V \min_L d_G(u, v)$, therefore $d_H(V, L)$ is upper bounded by ϵ .

Topological Approximation Guarantee of an ϵ -net induced Lazy witness complex on graphs. In addition to an ϵ -net being an ϵ -approximation of the space (V, d_G) , we prove that the lazy witness complex induced by the ϵ -net, as landmarks, is a good approximation (Theorem 2) to the Vietoris-Rips complex on the same set of vertices. This approximation ratio is independent of the algorithm constructing the ϵ -net. As a step towards Theorem 2, we state Lemma 2. Lemma 2 is implied by 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 of (V, d_G) , the distance $d_G(u, v)$ from any vertex $u \in L$ to its 1-nearest neighbour $v \in V$ is at most ϵ .*

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

Theorem 2. *If L is an ϵ -net of the point cloud V for $\epsilon \in \mathbf{R}^+$, $LW_\alpha(V, L, \nu = 1)$ 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 , $R_{\alpha/3}(L) \subseteq LW_\alpha(V, L, 1) \subseteq R_{3\alpha}(L)$ for $\alpha \geq 2\epsilon$.*

Proof. In order to prove the first inclusion, consider a k -simplex $\sigma_k = [x_0 \cdots x_k] \in R_{\alpha/3}(L)$. For any edge $[x_i x_j] \in \sigma_k$, let w_t be the point in V that is nearest to the vertices of $[x_i x_j]$. Without loss of generality, let that vertex be x_j . Since w_t is the nearest neighbour of x_j , by Lemma 2, $d_G(w_t, x_j) \leq \epsilon \leq \frac{\alpha}{2}$. Since $[x_i x_j] \in R_{\alpha/3}$, $d_G(x_i, x_j) \leq \frac{\alpha}{3} < \frac{\alpha}{2}$. By triangle inequality, $d_G(w_t, x_i) \leq \frac{\alpha}{2} + \frac{\alpha}{2} \leq \alpha$. 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 $LW_\alpha(V, L, 1)$. Since the argument is true for any $x_i, x_j \in \sigma_k$, the k -simplex $\sigma_k \in LW_\alpha(V, L, 1)$.

In order to prove the second inclusion, consider a k -simplex $\sigma_k = [x_0 x_1 \cdots x_k] \in LW_\alpha(V, L, 1)$. Therefore, by definition of lazy witness complex, for any edge $[x_i x_j]$ of σ_k there is a witness $w \in V$ such that, the $(d^1(w) + \alpha)$ -neighbourhood of w contains both x_i and x_j . Hence, $d_G(w, x_i) \leq d^1(w) + \alpha \leq \epsilon + \alpha$ (by Lemma 2) $\leq 3\alpha/2$. Similarly, $d_G(w, x_j) \leq 3\alpha/2$. By triangle inequality, $d_G(x_i, x_j) \leq 3\alpha$. 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, their corresponding partial persistence diagrams $Dgm_{>2\epsilon}(LW)$ and $Dgm_{>2\epsilon}(R)$ are $3 \log 3$ -approximations of each other in log-scale, by the persistence approximation lemma [26]. In Section 5, we empirically validate this bound for the lazy witness complex induced by the ϵ -net landmarks.

Size of an ϵ -net. We prove an upper bound on the size of an ϵ -net of a connected unweighted graph using the *doubling dimension*.

The doubling dimension [16] of a metric space $M = (X, d)$ is the smallest positive number D such that any ϵ -neighbourhood in M can be covered by 2^D number of $\frac{\epsilon}{2}$ -neighbourhoods. A metric space is called *doubling* if its doubling dimension is bounded. The space (V, d_G) is a doubling metric space.

Gupta et. al. [16] showed that the doubling dimension $D(G)$ of an unweighted connected graph G is related to its local density. The local density of an unweighted connected graph G , denoted $\beta(G)$, is the smallest value β such that $|\mathcal{N}_\epsilon(v)| \leq \beta\epsilon$, for all $v \in V$ and $\epsilon \in \mathbf{N}$. To be precise the doubling dimension $D(G) \leq 4 \log(2\beta(G))$ [16]. We use this result along with the following lemma to prove bound on the size of an *epsilon*-net of an unweighted graph in Theorem 3.

Lemma 3. *For any connected unweighted graph G of diameter Δ and doubling dimension $D(G)$, there exists an ϵ -net of size at most $(\frac{\Delta}{\epsilon})^{D(G)}$ where $\Delta \geq \epsilon \geq 1$*

Proof. Let Δ be the diameter of an unweighted graph G of doubling dimension $D(G)$. Thus, V (a Δ -neighbourhood) can be covered by $2^{D(G)}$ number of $\Delta/2$ -neighbourhoods. Each $\Delta/2$ -neighbourhood can be covered by $2^{D(G)}$ number of $\Delta/4$ -neighbourhoods. Thus V can be covered by $2^{2D(G)}$ number of $\Delta/4$ -neighbourhoods. Repeating this $\log_2(\frac{\Delta}{\epsilon})$ -times, we get that V can be covered by $2^{D(G) \log_2(\frac{\Delta}{\epsilon})}$ number of ϵ -neighbourhoods. Each of the ϵ -neighbourhoods contain at most one ϵ -net-landmark. Hence, there exists an ϵ -net of size $(\frac{\Delta}{\epsilon})^{D(G)}$ for any connected, unweighted graph G .

Theorem 3. *For any connected unweighted graph $G = (V, E)$ of diameter Δ , there exists an ϵ -net of size at most $(\frac{\Delta}{\epsilon})^{O(\log(\frac{|V|}{\epsilon}))}$*

Proof. In a connected unweighted graph G , the size $|\mathcal{N}_\epsilon(v)|$ of ϵ -neighbourhood of a vertex v is greater than ϵ for $\Delta \geq \epsilon \geq 1$. Thus, local density $\beta(G) > 1$. We observe that $\max_v |\mathcal{N}_\epsilon(v)| < |V|$. Thus, $\beta(G) = \max_{v, \epsilon} \frac{|\mathcal{N}_\epsilon(v)|}{\epsilon}$ is at most $\frac{|V|}{\epsilon}$. Applying the result from Gupta et.al. [16], the doubling dimension $D(G) \leq 4 \log(\frac{|V|}{\epsilon})$. The rest follows from Lemma 3.

4 Algorithms for Computing ϵ -Net

In this section, we propose and elaborate three algorithms, namely Greedy- ϵ -Net, Iterative- ϵ -Net, and SPTpruning- ϵ -net, for computing ϵ -net on graphs.

4.1 Greedy- ϵ -Net Algorithm

We propose a greedy algorithm, namely Greedy- ϵ -Net, to compute a minimal-cardinality ϵ -net of a graph. Greedy- ϵ -Net (Algorithm 1) maintains a hash table with vertices as keys and the number of vertices in their ϵ -cover as values. At each step, Greedy- ϵ -Net selects a vertex with the largest ϵ -cover, marks the covered vertices, and updates the ϵ -cover of other vertices until all the vertices are marked as covered.

Algorithm 1 Greedy- ϵ -Net

Input: Graph $G = (V, E, W)$, parameter ϵ **Output:** Set of Landmarks L

```

1: Initialize  $L = \phi$ 
2: Let  $nc$  be the hash table with vertices as keys, number of vertices in their  $\epsilon$ -cover
   as value.
3: for all  $u \in V$  do
4:    $nc[u] = \epsilon\text{-BFS}(G, u, \epsilon)$ 
5: end for
6: Initialize all vertex  $u \in V$  as marked.
7: repeat
8:   Sort  $nc$  in descending order of its value.
9:   if  $marked[u] = \text{False}$  then
10:     $L.insert(u)$ 
11:    for all vertex  $u'$  in  $u$ 's  $\epsilon$ -cover do
12:      Mark  $u'$  as True.
13:      Delete key  $u'$  from  $nc$ .
14:      for all  $v \in V$  do
15:        if  $u'$  is in  $\epsilon$ -cover of  $v$  then
16:          Decrease  $nc[v]$  by 1.
17:        end if
18:      end for
19:    end for
20:   end if
21:   Delete key  $u$  from  $nc$ .
22: until all vertices are marked

```

Algorithm 2 ϵ -BFS

Input: Graph $G = (V, E, W)$, vertex u , parameter ϵ .**Output:** Set of vertices in u 's ϵ -cover, C_ϵ

```

1: Initialize Queue  $Q = \{u\}$ 
2: initialize  $C_\epsilon = \phi$ 
3: while  $Q \neq \phi$  do
4:    $v = \text{DEQUEUE}(Q)$ 
5:    $v.marked = \text{True}$ 
6:   for all  $v' \in G.Adj[v]$  do
7:     if  $v'.marked = \text{False}$  then
8:        $v'.d = v.d + W[v, v']$ 
9:       if  $v'.d \leq \epsilon$  then
10:         $C_\epsilon = C_\epsilon \cup \{v'\}$ 
11:       end if
12:     end if
13:   end for
14: end while

```

4.2 Iterative- ϵ -Net Algorithm

Iterative- ϵ -Net (Algorithm 3) is a diffusive algorithm that maintains a set $C_{(\epsilon, 2\epsilon)}$ that contains the set of unmarked vertices that are within a ring of $(\epsilon, 2\epsilon]$ distance from the current set of landmarks. We call them the **ring vertices**. Iterative-

Algorithm 3 Iterative- ϵ -Net

Input: Graph $G = (V, E, W)$, parameter ϵ **Output:** Set of landmarks L

```

1: Initialise  $L = \phi$ 
2:  $i = 1$ 
3: Select initial landmark  $l_i = u$  uniformly at random from  $V$ .
4:  $L = L \cup \{u\}$ 
5: Let  $C_{(\epsilon, 2\epsilon)}^u$  be the set of unmarked vertices  $v$  such that  $\epsilon < d_G(u, v) \leq 2\epsilon$ .
6: Let  $C_{2\epsilon}^u$  be the set of vertices  $v$  such that  $d_G(u, v) > 2\epsilon$  and are the closest.
7: repeat
8:    $C_{\epsilon, 2\epsilon}^u = \text{PartialBFS}(G, l_i, C_{(\epsilon, 2\epsilon)}^u, C_{2\epsilon}^u)$ 
9:   if  $C_{2\epsilon}^u$  is empty then
10:     select  $l_{i+1}$  uniformly at random from  $C_{(\epsilon, 2\epsilon)}^u$ .
11:   else
12:     select  $l_{i+1}$  uniformly at random from  $C_{2\epsilon}^u$ .
13:   end if
14:    $L = L \cup \{l_{i+1}\}$ .
15:    $i = i + 1$ 
16:    $u = l_{i+1}$ 
17: until all vertices are marked

```

ϵ -Net also maintains another set $C_{2\epsilon}$ that contains the set of unmarked vertices that are at distance at least 2ϵ but are adjacent to the bordering vertices of the cover. We call them the **enveloping vertices**.

Iterative- ϵ -net, at each iteration, uniformly at random selects a vertex u from the current set of enveloping vertices as next landmark, and run Partial-BFS (Algorithm 4) starting at u to mark the vertices in its ϵ -cover, as well as to update the enveloping and ring vertex sets. If enveloping vertex set is empty it selects a ring vertex uniformly at random as next landmark.

Iterative ϵ -net algorithm has the property that some vertex in the ϵ -cover of landmark l_{i+1} is always adjacent to some vertex in the ϵ -cover of landmark l_i . Thus the two covers are adjacent as sets.

4.3 Sparsity based Pruning Algorithm on Shortest Path Tree

We propose SPTpruning- ϵ -Net algorithm (Algorithm 5) that constructs a shortest path tree of the graph and uses the tree to compute an ϵ -net. Algorithm 5 computes a shortest path tree rooted at a vertex chosen uniformly at random. The algorithm uses ϵ -BFS (Algorithm 2) to construct a preliminary BFS spanning tree of the graph (line 2). Then the algorithm constructs an ϵ -net of the BFS-tree (line 4-28). It does so by traversing the tree level-order starting from root, running ϵ -BFS in the tree to mark covered vertices, and add the set of unmarked vertices at level $\epsilon + 1$ as candidates for landmarks.

An ϵ -net L_{SPT} of a BFS-tree $SPT(G)$ of a graph G has the property that, any vertex $v \in SPT(G) \subset G$ that is covered by some vertex $u \in L_{SPT}$, is also covered by the $u \in V$ in the graph as well. This property follows from the fact

Algorithm 4 PartialBFS

Input: Graph G , vertex u , set $C_{(\epsilon, 2\epsilon)}^u$, set $C_{2\epsilon}^u$.

- 1: Initialise Queue $Q = \{u\}$
- 2: $u.marked = \text{True}$
- 3: **while** $Q \neq \phi$ **do**
- 4: $v = \text{DEQUEUE}(Q)$
- 5: **for all** $v' \in G.Adj[v]$ **do**
- 6: **if** $v'.marked = \text{False}$ **then**
- 7: $v'.d = v.d + W[v, v']$
- 8: **if** $v'.d \leq \epsilon$ **then**
- 9: $v'.marked = \text{True}$
- 10: Remove v' from $C_{(\epsilon, 2\epsilon)}^u$ and $C_{2\epsilon}^u$ if exists.
- 11: $\text{ENQUEUE}(Q, v')$
- 12: **else if** $\epsilon < v'.d \leq 2\epsilon$ **then**
- 13: $C_{(\epsilon, 2\epsilon)}^u = C_{(\epsilon, 2\epsilon)}^u \cup \{v'\}$
- 14: $\text{ENQUEUE}(Q, v')$
- 15: **else**
- 16: $C_{2\epsilon}^u = C_{2\epsilon}^u \cup \{v'\}$
- 17: **end if**
- 18: **end if**
- 19: **end for**
- 20: **end while**

that, the distance between any vertex pair $d_G(u, v)$ in the graph can only be shorter than their distance $d_{SPT}(u, v)$ in the BFS-tree.

Unless one of the vertex is a root, the distance between any pair of vertices in the tree is not guaranteed to be the shortest in the graph. Thus an ϵ -net of a BFS-tree does not have ϵ -sparsity in the graph containing the BFS-tree as a subgraph. As a remedy, Algorithm 6 prunes vertices from the candidate landmarks that are covered by other candidate landmark.

5 Performance Analysis

In this section, we experimentally and comparatively analyse the performance of the three proposed algorithms to compute an ϵ -net of graphs. We discuss the effectiveness and efficiency of the algorithms and also validate that the ϵ -net computed by any of these algorithms satisfy being 3-approximation of the Vietoris-Rips complex.

5.1 Datasets

We evaluate the performance of our algorithms using two real-world datasets. The dataset Power [30] is an unweighted graph of US Power-grid (4941 vertices, 6594 edges, diameter 46). Celegans [2] is a weighted graph of Celegans worm's frontal neural network (297 vertices, 2148 edges, diameter 1.333).

5.2 Experimental Setup

We implement the experimental workflow in C++. We use Snap library [22] for graph processing and ϵ -net computations. We modify the Ripser³ library

³ <https://github.com/Ripser/riper>

Algorithm 5 SPTpruning- ϵ -net

Input: Graph G , Diameter Δ , parameter ϵ **Output:** Set of landmarks C

```

1: Select a vertex  $u$  uniformly at random from  $V$ .
2: Run  $\epsilon$ -BFS( $G, u, \Delta$ ) to construct a BFS spanning tree rooted at  $u$ 
3: Let  $SPT$  be the tree.
4: Initialise Queue  $Q = \{u\}$ 
5:  $u.marked = \text{True}$ 
6: Let  $\epsilon$ -net of the SPT  $C_\epsilon = \phi$ 
7: repeat
8:   while  $Q \neq \phi$  do
9:      $u' = \text{DEQUEUE}(Q)$ 
10:     $u'.marked = \text{True}$ 
11:     $u'.d = 0$ 
12:    Initialise Queue  $Q' = \{u'\}$ 
13:    while  $Q' \neq \phi$  do
14:       $v = \text{DEQUEUE}(Q')$ 
15:      for all  $v' \in SPT.Adj[v]$  do
16:        if  $v'.marked = \text{False}$  then
17:           $v'.d = v.d + W[v, v']$ 
18:          if  $v'.d \leq \epsilon$  then
19:             $v'.marked = \text{True}$ 
20:             $\text{ENQUEUE}(Q', v')$ 
21:          else
22:             $C_\epsilon = C_\epsilon \cup \{v'\}$ 
23:             $\text{ENQUEUE}(Q, v')$ 
24:          end if
25:        end if
26:      end for
27:    end while
28:  end while
29: until All vertex in  $SPT$  are marked
30:  $C = \text{Prune}(G, C_\epsilon, \epsilon)$ 

```

Algorithm 6 Prune

Input: Graph G , A set of landmarks C_ϵ , parameter ϵ **Output:** Set of landmarks C

```

1:  $C = C_\epsilon$ 
2: for all vertex  $v \in C_\epsilon$  do
3:   if  $v.marked = \text{False}$  then
4:     Run  $\epsilon$ -BFS( $G, v, \epsilon$ )
5:   else
6:      $C = C \setminus \{v\}$ 
7:   end if
8: end for

```

to compute lazy witness complexes and their persistent intervals. We use R-TDA package [14] to compute bottleneck distances. All experiments are run on a machine with an Intel(R) Xeon(R)@2.20GHz CPU and 80 GB memory




Fig. 2: Performance of the algorithms on Celegans dataset.




Fig. 3: Performance of the algorithms on Power dataset.

limit. 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 corresponding dataset. We compute persistent intervals at dimension 0 and 1.

5.3 Efficiency and Effectiveness

We measure the efficiency of our algorithms in terms of CPU time (in ms) required to select ϵ -net landmarks for a given ϵ , and the overall computation of persistent homology of each graph. The overall total computation time includes the time an algorithm spends constructing ϵ -net and time spent on computing persistent intervals at dimension 0 and 1. We measure the effectiveness of the algorithms using the number of landmarks they select by corresponding ϵ -net construction.

Figures 2 and 3 illustrate the experimental results for the Celegans and the Power dataset respectively. We observe that the number of landmarks selected by all the algorithms decrease as the ϵ -increases. The landmark selection time of the Greedy- ϵ -Net increases as the ϵ increases independent of the dataset. For other two algorithms, the landmark selection time varies depending on the density of the graph. The landmark computation time of Iterative- ϵ -net and SPTpruning- ϵ -net are almost invariant with ϵ .

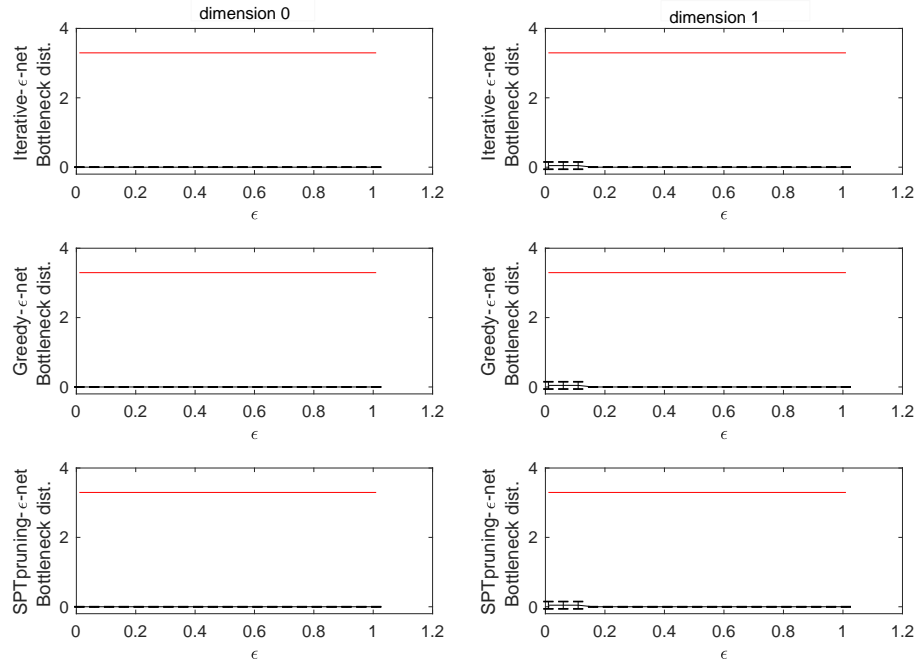


Fig. 4: Validation of the approximation guarantee of the ϵ -net induced lazy witness complexes for Celegans dataset. The bottleneck distance between the partial persistence diagram of the corresponding Vietoris-Rips and lazy witness filtrations is less than $3 \log 3$ in both dimensions 0 and 1.

We observe that Iterative- ϵ -Net takes longer time (> 10 ms) to select landmarks compared to the SPTpruning- ϵ -Net algorithm but it selects less number of landmarks than SPTpruning- ϵ -Net algorithm. Thus, the overall runtime of lazy witness complex construction using Iterative- ϵ -Net is smaller than that of the SPTpruning- ϵ -Net. The empirical performance analysis instantiates Iterative- ϵ -Net as the practical and efficient choice to construct ϵ -net induced lazy witness complex for graphs.

5.4 Validating the Approximation Guarantee

In order to validate the approximation guarantee in Theorem 2, we construct the Vietoris-Rips complex and lazy witness complexes using ϵ -nets for different values of ϵ and different algorithms. We compare the corresponding complexes by computing bottleneck distances between the persistence diagrams at dimension 0 and 1. We retain only the partial diagram with points in the diagram born after 2ϵ for this purpose. Figure 4 validates the guarantee on Celegans dataset. We omit the validation on Power dataset for the sake of brevity.

6 Conclusion and Future Work

We investigate the computation of persistent homologies on weighted graphs. We extend the notion of ϵ -net for point clouds to weighted graphs. We further

propose an ϵ -net induced lazy witness complex that leverages the ϵ -net and their geodesic distances to select landmarks.

We prove that an ϵ -net of a connected graph is an ϵ -approximation of its set of vertices in Hausdorff distance. We also prove that the lazy witness complex induced by an ϵ -net is a 3-approximation of the Vietoris-Rips complex induced by the same set. We prove the existence of an ϵ -net (of a graph) of size at most $(\frac{\Delta}{\epsilon})^{O(\log \frac{|V|}{\epsilon})}$, where Δ is the diameter of the graph.

We present three algorithms for constructing an ϵ -net of a graph. We comparatively and empirically evaluate the efficiency and effectiveness of the choice of landmarks that they induce for the topological data analysis of several real world graphs. The empirical performance analysis instantiates Iterative- ϵ -Net as the practical and efficient choice to construct ϵ -net induced lazy witness complex for graphs.

An interesting future work would be to leverage the notion of ϵ -net induced lazy witness complex for faster and scalable computation of machine learning problems, such as clustering [8], deep learning [19], kernel density estimation [6], for both graphs and point clouds.

Acknowledgement

This work is partially supported by the National University of Singapore Institute for Data Science project WATCHA and by Singapore Ministry of Education project Janus.

References

1. Arafat, N.A., Basu, D., Bressan, S.: Topological data analysis with ϵ -net induced lazy witness complex. arXiv preprint arXiv:1906.06122 (2019)
2. Badhwar, R., Bagler, G.: Control of neuronal network in caenorhabditis elegans. PloS one 10(9), e0139204 (2015)
3. Berge, C.: Graphs and hypergraphs. North-Holland, New York (1976)
4. Bubenik, P.: Statistical topological data analysis using persistence landscapes. The Journal of Machine Learning Research 16(1), 77–102 (2015)
5. Carlsson, G.: Topology and data. Bulletin of The American Mathematical Society, Vol. 46(2), pp. 255–308 (2009)
6. Carriere, M., Cuturi, M., Oudot, S.: Sliced wasserstein kernel for persistence diagrams. In: Proceedings of the 34th International Conference on Machine Learning-Volume 70. pp. 664–673. JMLR. org (2017)
7. Carstens, C.J., Horadam, K.J.: Persistent Homology of Collaboration Networks. Mathematical Problems in Engineering 2013(6), 1–7 (jun 2013)
8. Chazal, F., Guibas, L.J., Oudot, S.Y., Skraba, P.: Persistence-based clustering in riemannian manifolds. Journal of the ACM (JACM) 60(6), 41 (2013)
9. 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)
10. 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)
11. Dey, T.K., Fan, F., Wang, Y.: Graph induced complex on point data. Computational Geometry 48(8), 575–588 (2015)

12. Duman, A.N., Pirim, H.: Gene coexpression network comparison via persistent homology. *International journal of genomics* 2018 (2018)
13. Edelsbrunner, H., Harer, J.: *Computational topology: an introduction*. American Mathematical Soc. (2010)
14. Fasy, B.T., Kim, J., Lecci, F., Maria, C.: Introduction to the r package tda. arXiv preprint arXiv:1411.1830 (2014)
15. Guibas, L.J., Oudot, S.Y.: Reconstruction using witness complexes. *Discrete & computational geometry* 40(3), 325–356 (2008)
16. Gupta, A., Krauthgamer, R., Lee, J.R.: Bounded geometries, fractals, and low-distortion embeddings. In: *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings*. pp. 534–543. IEEE (2003)
17. 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)
18. Heinonen, J.: *Lectures on analysis on metric spaces*. Springer Science & Business Media (2012)
19. Hofer, C., Kwitt, R., Niethammer, M., Dixit, M.: Connectivity-optimized representation learning via persistent homology. In: *Proceedings of the 36th International Conference on Machine Learning*. vol. 97, pp. 2751–2760. PMLR (2019)
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. Leskovec, J., Sosič, R.: Snap: A general-purpose network analysis and graph-mining library. *ACM Transactions on Intelligent Systems and Technology (TIST)* 8(1), 1 (2016)
23. Newman, M.E.: Analysis of weighted networks. *Physical review E* 70(5) (2004)
24. 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)
25. Petri, G., Scalamiero, 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. Silva, J., Marques, J., Lemos, J.: Selecting landmark points for sparse manifold learning. In: *Advances in neural information processing systems*. pp. 1241–1248 (2006)
28. Vandaele, R., De Bie, T., Saeys, Y.: Local topological data analysis to uncover the global structure of data approaching graph-structured topologies. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. pp. 19–36. Springer (2018)
29. Vietoris, L.: Über den höheren zusammenhang kompakter räume und eine klasse von zusammenhangstreuen abbildungen. *Mathematische Annalen* 97(1), 454–472 (1927)
30. Watts, D.J., Strogatz, S.H.: Collective dynamics of ‘small-world’ networks. *Nature* 393(1), 440–442 (1998)
31. Zomorodian, A.: Fast construction of the vietoris-rips complex. *Computers & Graphics* 34(3), 263–271 (2010)