Neighborhood-Preserving Graph Sparsification

Abd Errahmane Kiouche
Universite Claude Bernard Lyon 1,
CNRS, INSA Lyon, LIRIS, UMR5205
69622, Villeurbanne, France
abderrahmane.kiouche@gmail.com

Julien Baste
Univ. Lille, CNRS, Centrale Lille, UMR
9189 CRIStAL
59000, Lille, France
julien.baste@univ-lille.fr

Mohammed Haddad
Universite Claude Bernard Lyon 1,
CNRS, INSA Lyon, LIRIS, UMR5205
69622, Villeurbanne, France
mohammed.haddad@univ-lyon1.fr

Hamida Seba
Universite Claude Bernard Lyon 1,
CNRS, INSA Lyon, LIRIS, UMR5205
69622, Villeurbanne, France
hamida.seba@univ-lyon1.fr

Angela Bonifati
Universite Claude Bernard Lyon 1,
CNRS Liris, IUF
69622, Villeurbanne, France
angela.bonifati@univ-lyon1.fr

ABSTRACT

We introduce a new graph sparsification method that targets the neighborhood information available for each node. Our approach is motivated by the fact that neighborhood information is used by several mining and learning tasks on graphs as well as reachability queries. The result of our sparsification technique is a sparsified graph that can be used instead of the original graph in the above tasks while still ensuring fairly good approximations for the results. Moreover, our sparsification method allows users to control the size of the resulting sparsified graph by adjusting the amount of information loss tolerated by the targeted applications. Our extensive experiments conducted on various real and synthetic graphs show that our sparsification considerably reduces the size of the graphs by achieving 40% sparsification rate on average on several input graphs. Furthermore, in the experimental study we show the utility and efficiency of our sparsification algorithm for notable data-driven tasks, such as node classification, graph classification and shortest path approximations.

PVLDB Reference Format:
Abd Errahmane Kiouche, Julien Baste, Mohammed Haddad, Hamida Seba, and Angela Bonifati. Neighborhood-Preserving Graph Sparsification. PVLDB, 17(13): 4853 - 4866, 2024.
doi:10.14778/3704965.3704988

PVLDB Artifact Availability:
The source code, data, and/or other artifacts have been made available at https://gitlab.liris.cnrs.fr/coregraphie/ptspar.

1 INTRODUCTION

Graphs are data modeling abstractions consisting of a set of vertices, also called nodes, and a set of edges connecting the vertices. Vertices represent objects, while edges represent relationships between them. Graphs are widely used in data modeling because of their ability to represent, in a simple and intuitive way, complex processes in both nature and technology, such as social interactions, protein-protein interactions, chemical molecules, transport networks and fraud detection networks, to name a few [38]. However, most of these graphs are very large or grow exponentially as new data arrives. This makes graph querying and analysis a very challenging task.

To tackle scalability and performance issues when dealing with large graph data, plenty of algorithms are devised to simplify graphs in several domains and applications related to graph analysis [30]. The aim is to construct simpler or smaller representations for large graphs mainly to save storage space but also to use the obtained representations, instead of the original graphs, in applications where using the entire large original graphs is not possible or is time consuming [27]. Sparsification is one of these approaches aiming to construct a subgraph of the original graph by removing insignificant edges. The resulting graph is called a skeleton or a backbone. Sparsification is generally application-dependent because the significance of an edge may vary from one application to another. The main idea is obtaining a smaller graph while preserving some properties, even approximately, of the original graph such as the results of distance, and reachability queries [16, 29]. We note that graph sparsification is a lossy graph simplification technique that differs from lossless approaches such as graph compression [5] or graph summarisation and contraction methods [7, 8, 13, 15, 25, 27], where the output structure is not always a graph. These methods also require partial or total decompression to query the output summary. Several sparsification methods are proposed in the literature [31] but there is no generic approach that can target several applications at once. An attempt to build a fairly general approach is provided in [44], relying on reinforcement learning. However, this method does not exempt us from computing the application task on the original graph; on the contrary, this step is mandatory during the training process that needs to be achieved on each graph to be sparsified. This makes it less useful for real-world applications and hard to use for graphs that the algorithm has not used during training.

In this paper, we fill the gap by proposing a general sparsification method, called (p, t)−Sparsification, that focuses on preserving the neighborhood information available around each vertex. This kind of local information is a key property for many graph algorithms but it has never been exploited for the graph sparsification problem.

Proceedings of the VLDB Endowment, Vol. 17, No. 13 ISSN 2150-8097.
doi:10.14778/3704965.3704988

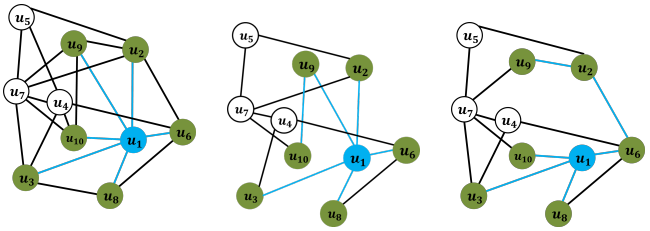

(a) Input Graph **(b)** $p(1) = 50\%$ $t = 2$ **(c)** $p(1) = 50\%$ $t = 3$

Figure 1: (p, t)−sparsification and neighborhood information.

For instance, in graph learning tasks such as node classification, knowing a node's immediate connections is crucial for accurately categorizing it based on the characteristics of its neighbors. Similarly, for reachability queries, where the goal is to find whether there is a path between two nodes, neighborhood information speeds up the process and makes it more efficient. Computing shortest paths in a graph also needs the information about each node's local surroundings to identify the most efficient and least costly paths.

Each node in a graph may have several connections, and handling all of them is time and space consuming. Hence, our approach allows us to finely control how much locality we keep in the obtained sparsified graph. The key idea is to remove some edges while ensuring that, for each node u, a certain amount of its neighbors, determined by a function $p()$, is kept within at most t-hops from u in the resulting sparsified graph, with $t \geq 1$ and $p()$ being the input parameters.

Figure 1 shows how the neighborhood information is preserved with (p, t)-sparsification. In the original graph (cf. Figure 1 (a)), node u_1 has 6 direct neighbors (highlighted in blue): nodes u_2, u_3, u_6, u_8, u_9 and u_{10}. Figure 1 (b) shows a (p, t)−sparsification of G where $p(1) = 50\%$ of the neighbors of node u_1 remain reachable at $t = 2$ hops. We can observe that node u_2 is no longer directly connected to node u_1 in the sparsified graph but remains reachable within 2 hops from it. Figure 1 (c) shows another sparsification of G this time with $p(3) = 100\%$ of the neighbors of node u_1 remaining reachable within $t = 3$ hops. We can also see that nodes u_2 and u_9 are no longer directly connected to node u_1 in the sparsified graph but remain reachable within 3 hops from it. By varying the (p, t) parameters, specified as input, one can see that the obtained sparsification can be accordingly tuned, thus explaining its generalizability to any input graph.

Figure 1 also provides an intuition of how the preserved neighborhood information in the sparsified graph can be used when computing shortest paths. If we consider the case where for each vertex v, all the neighborhood of v can be found within t-hops of v in the sparsified graph, one can easily check that for each path of size k in the original graph, there is a corresponding path of size at most $k \cdot t$ in the sparsified graph. For instance, in Figure 1(a) the shortest path $u_1u_2u_5$ from u_1 to u_5 is of size 2 thus, we know that in Figure 1(b), the shortest path from u_1 to u_5 is of size at most $2 \cdot 2$, which is $u_1u_{10}u_7u_5$ of size 3. The example of Figure 1 also shows that parameters p and t allow to have a generic approach that can produce various sparsifications by using different values for p and t and therefore adapt to multiple applications and needs.

To show this, we use our resulting sparsified graphs in four main tasks: reachability queries, shortest paths computation, node classification and graph classification. The results we obtained on these four different tasks have proven to be effective approximations of the exact results of these tasks obtained on the full original graphs. This validates our method as producing useful sparsified graphs reducing the size of the input graph, and allowing to leverage the sparsified graphs, instead of the original ones, without sacrificing the accuracy of the outcomes. This efficient approximation is crucial in computational environments where resource savings are imperative, yet the integrity of the results must be maintained. It is also interesting in several applications where a quickly delivered approximate result is more useful than an accurate result that takes a long time to obtain.

Summarizing, the main contributions of this paper are as follows:

- We introduce a new general graph sparsification method that targets the preservation of the neighborhood information available for each node in the graph. The key idea is to create a graph skeleton that can replace the original graph in various applications and for arbitrary graph datasets.
- Our method allows the fine-tuning of the size of the resulting sparsified graph by adjusting the amount of preserved neighborhood information.
- Our method allows efficient graph algorithm approximations. Our approach improves the efficiency of various graph algorithms, particularly those that depend on node neighborhood information, such as node/graph classification and shortest path computation. This results in faster and more robust approximations for larger graphs.
- Our techniques allow us to strike a user-driven balance between sparsification ratio and information loss through its parameters p and t. The user-driven customization makes our technique versatile for different applications with varying needs.
- We conduct extensive experiments using real-world datasets in different application domains to assess the effectiveness and efficiency of the proposed method. The source code and the data used in our paper are publicly available at https://gitlab.liris.cnrs.fr/coregraphie/ptspar.

The remainder of this paper is organized as follows: Section 2 defines the basic concepts and notation used in the paper and reviews related work on graph sparsification methods. Section 3 formally defines the problem of neighborhood-preserving graph sparsification and studies its complexity. Then, Section 4 provides a description of the algorithms that we propose to compute this sparsification and analyses their time complexity. Section 5 presents the results obtained through the extensive experiments we undertook to evaluate our sparsification approach, as well as the usefulness of the obtained sparsified graphs. Finally, Section 6 concludes the paper and points out some research perspectives.

2 PRELIMINARIES AND RELATED WORK

2.1 Preliminary

A graph $G = (V, E)$ is a 2-component structure comprising a set V of vertices and a set $E \subseteq V \times V$ of edges connecting the vertices. Edges can be directed, and both vertices and edges can have attributes or weights. In this paper, we consider unweighted and undirected

Table 1: Notation.

Symbol	Description
$G(V, E)$	an undirected unweighted graph with V the set of vertices and E the set of edges.
$N_G^h(v)$	the set of all q-hop neighbors, $q \in \{1, \ldots, h\}$, of vertex v in G
$N(v)$	the set of direct neighbors of v, i.e., $N_G^1(v)$
$deg(v)$	number of direct neighbors of v
d	average node degree of G
$G_s = (V_s, E_s)$	a sparsified graph of G
π	an ordering function defined on E
E^π	the edges of G in the order defined by π
$\mathcal{W}$	the set of all paths in G
$\mathcal{W}_{uv}$	the set of all paths in $\mathcal{W}$ from u to v
$\mathcal{W}_{uv}^i$	the set of all paths of $\mathcal{W}_{uv}$ of length at most i that pass through the edge e
Sr	the sparsification ratio given by $\frac{\lvert E\rvert - \lvert E_S\rvert}{\lvert E\rvert}$

graphs. Thus, an edge connecting vertices u and v is interchangeably denoted by uv or vu.

Two vertices connected by an edge are said to be adjacent. A vertex adjacent to v is called a direct neighbor of v or a 1-hop neighbor of v. The set of all the direct neighbors of a vertex v in G is denoted as $N_G^1(v)$ or simply $N(v)$ when there is no ambiguity. The degree of a vertex v, denoted $deg(v)$, is the number of its 1-hop neighbors, i.e., $deg(v) = |N(v)|$.

A path in a graph is a sequence of edges which joins a sequence of vertices which are all distinct. A vertex v is reachable from a vertex u if there exists a path from u to v. A *q-hop neighbor* of a vertex v is a vertex that can be reached from v with a path of exactly q edges. We will denote by $N_G^h(v)$ the set of all q-hop neighbors, $q \in \{1, \ldots, h\}$, of vertex v in G.

The distance between two vertices u and v is the length, i.e., number of edges, of the shortest path connecting them. We will denote by $\mathcal{W}$ the set of all paths in G, by $\mathcal{W}uv$ the set of all paths from node u to node v, and by $\mathcal{W}uv^i$, the set of paths from node u to node v of length at most i.

Graph sparsification stands for the methods that compute a sparse subgraph of the input graph. Given a graph $G = (V, E)$, a sparsified graph of G is a graph $G_s = (V_s, E_s)$ defined generally on the same set of vertices as G but with less edges, i.e., $V_s = V$ and $E_s \subset E$. In practice, we want the sparsified graph G_s to retain certain properties of G such as the distance between vertices, reachability queries, etc.

When sparsifying a graph, the order in which the edges are processed is generally important. We will denote this order with a bijective function, i.e., permutation, $\pi : E \rightarrow \{1, \ldots, |E|\}$ that associates to each edge $e \in E$ its processing rank $\pi(e)$. We will denote by E^π the edges of G in the order defined by π.

Let $G(V, E)$ be the input graph and $G_s(V_s, E_s)$ be the sparsified graph, the sparsification ratio measures how well G_s reduces the graph G and is given by the ratio of the number of deleted edges over the total number of edges:

$$Sr = \frac{|E| - |E_s|}{|E|} \tag{1}$$

Note that the higher is the sparsification ratio the better is the storage space gain ensured by the sparsification.

Table 1 summarizes these notations.

2.2 Related Work

Graph sparsification is a lossy graph simplification technique that allows to compute a subgraph of the original graph preserving some of its properties. As such, it differs from lossless approaches such as graph contraction methods [7, 8, 13, 15, 25, 27]. The latter methods allow exact computations on the output structure but require partial or total decompression to achieve these computations. Therefore, the two approaches are inherently different. In the remainder of this section, we focus on discussing graph sparsification approaches. Graph sparsification methods can be classified into two main categories: statistical methods and structural methods.

Statistical methods: These methods extract the backbone of graphs by removing edges based on various properties of graph vertices and edges such as degree distribution or betweenness centrality distribution [23, 48]. They rely mainly on edge weights to sparsify the graph. Hence, they perform poorly on unweighted graphs. As an example, we can cite methods that apply a filter on the edges such as the Noise Corrected filter [11] that keeps only the edges with weight greater than a given threshold. High salience backbone filter [16] extracts the graph backbone based on the link (i.e., edge) salience property. The salience of an edge e is a score $s(e)$ that represents a consensus estimate from all nodes of the importance of the edge e. An edge e having a salience score equal to 1.0 is an essential edge for all nodes. If $s(e) = 0$, the edge e has no role and if, $s(e) = 0.5$ then it is important for only half of the nodes [16]. The salient backbone extracts the skeleton by keeping only the edges that have a salience score greater than a certain threshold.

Structural methods: These methods differ based on the structural properties of the input graph they aim to preserve in the constructed backbone. In [36], a backbone, called a spanner, preserves distances between vertices, within a multiplicative or an additive factor. The method in [3] sparsify graphs to preserve cuts, which are partitions of the vertices of a graph into two disjoint subsets. Given any weighted undirected graph $G = (V, E)$, the authors show that one could construct a new graph $G_\varepsilon = (V, E_\varepsilon \subseteq E)$, $0 < \varepsilon < 1$, with $|E_\varepsilon| = O(n \log n/\varepsilon^2)$ edges such that the value of every cut in G is within a multiplicative factor of $1 \pm \varepsilon$ of its value in G_ε. Similarly, we can find sparsification algorithms that preserve graph Laplacian [41], determinant of matrices [14], etc.

In [33] and [6], the authors tackle the problem of sparsifying a graph, while maintaining the connectivity recorded in a given set of observed activity traces represented by a set of trees with specified roots. In [44], the authors attempt to address several sparsification objectives. They propose SparRL, a graph sparsification approach based on graph neural networks (GNNs) and reinforcement learning. However, SparRL requires executing the downstream task algorithm on the original graph to calculate rewards for the reinforcement learning part. This raises questions about the practical utility of this sparsification method and limits its generalization to unseen graphs. Also, the method applies a task-specific optimization which is difficult to apply to some tasks such as learning and classification.

By targeting neighborhood information and allowing to control the amount of information loss in the computed sparsified graph, we aim to be able to use our skeletons in a variety of graph applications. In fact, several graph algorithms, such as node embedding,

node classification, shortest paths, etc. are based on the availability of node neighborhood information. In the remainder of the paper, we show that controlling the amount of this information in the computed skeleton allows a good trade-off between algorithm speed-up and precision loss when using the skeleton as input instead of the original graph in the targeted applications.

3 A NEIGHBORHOOD-PRESERVING GRAPH SPARSIFICATION

In this section, we introduce a new graph sparsification method that targets the amount of neighborhood information available for each node in the graph. The main idea is to sparsify the input graph by removing edges, while ensuring that, for all $1 \leq i \leq t$, a proportion $p(i)$ of the neighbors of each node v is included in the set of the i-hops neighbors of v in the resulting sparsified graph, where $t \geq 1$. We denote such sparsification by (p, t)-sparsification where:

- $p : \mathbb{N}^* \rightarrow [0, 1]$ is a monotonically increasing function, which represents the proportion of each node's input neighbors that must be available in its i-hops neighborhood in the sparsified output graph.
- t : is the minimum integer value for which p reaches its maximal value *i.e.*, $p(i) = p(t), \forall i \geq t$.

More formally, given an undirected graph $G = (V, E)$, a (p, t)-sparsification of G is defined as follows:

Definition 1. *Given a positive integer t and a monotonically increasing function $p : \mathbb{N}^* \rightarrow [0, 1]$ satisfying $p(i) = p(t)$ for all $i > t$, a (p, t)-*sparsification *of a graph $G = (V, E)$ involves finding a subgraph $G_s = (V_s, E_s)$ of G. G_s must have the same set of vertices $V_s = V$, a subset of edges $E_s \subseteq E$, and must satisfy the condition that for each integer $0 < i \leq t$ and each vertex $v \in V$, the set $N^i_{G_s}(v)$ includes at least a proportion $p(i)$ of the set $N^1_G(v)$ of immediate neighbors of v in G.*

The definition implies that the subgraph G_s retains fewer edges than the original graph G, but still captures a specified proportion of the original neighborhood structure.

With (p, t)-sparsification, the function p aims to control the loss of neighborhood information at varying depths. Naturally, a smaller value of p results in a higher sparsification ratio and vice versa.

Lemma 1. *For any (p, t)-sparsification, the number of edges $|E_s|$ of the sparsified graph satisfies the inequality $|E|p(1) \leq |E_s|$.*

Proof. The proof is straightforward and follows from the handshaking lemma [19] which states that in any graph, the sum of the degrees of all the vertices is twice the number of edges. □

Theorem 2. *Finding an optimal (smallest) graph satisfying the (p, t)-sparsification constraints for $t \geq 2$ is an NP-Hard problem.*

Proof Sketch. [1]

The proof of the theorem follows directly from the hardness of finding k-spanners which is known to be NP-complete [37]. □

[1]The detailed proofs are provided in the supplementary material [24].

4 COMPUTING (p, t)-SPARSIFIERS

In this section, we present two main algorithms for finding (p, t)-sparsifiers of an input graph G. The first algorithm is an exact algorithm based on an integer linear programming (ILP) formulation of (p, t)-sparsification. The second algorithm is an approximation whose result depends on the order on which the edges are processed, thus we provide several solutions to this ordering problem.

4.1 Exact Algorithm

Our exact algorithm is obtained by solving an Integer Linear Programming (ILP) formulation of (p, t)-sparsification. This formulation is aimed at finding an optimal (smallest) sparsified graph that meets the (p, t)-sparsification definition. It consists of a set of linear inequalities that constrains the minimization of an objective function. In our case, the objective function counts the number of edges of the sparsified graph (cf. Equation 2) and the constraints are expressed by Inequalities 3 to 5 and domain definition of our variables (cf. Equation 6).

Given a graph $G = (E, V)$ and a (p, t)-sparsification, the following ILP formulation computes a smallest sparsified graph as follows:

$$\text{minimize} \sum_{e \in E} x_e \tag{2}$$

subject to

$$\sum_{v \in N(u)} \sum_{w \in \mathcal{W}^i_{uv}} x_w \geq p(i) \cdot |N(u)| \quad \forall u \in V, i \leq t \tag{3}$$

$$\sum_{w \in \mathcal{W}_{uv}} x_w \leq 1 \quad \forall uv \in E \tag{4}$$

$$x_w \leq x_e \quad \forall w \in \mathcal{W}, e \in w \tag{5}$$

$$x_e, x_w \in \{0, 1\} \quad e \in E, w \in \mathcal{W} \tag{6}$$

where a binary variable x_e is defined for every edge $e \in E$ such that $x_e = 1$ if and only if the edge e is selected to be part of the sparsified graph, otherwise $x_e = 0$. Thus, the objective function $\sum_{e \in E} x_e$ aims at minimizing the number of selected edges in the final solution. The first constraint (cf. Equation 3) ensures that for every vertex u of the graph, the property of (p, t)-sparsification is satisfied *i.e.,* for every distance $i \leq t$, the number of neighbors still connected to u via a path of length at most i is at least $p(i) \cdot |N(u)|$. This is enforced by the binary variables x_w such that for every neighbor v of u, the set of all paths w between u and v is denoted by $\mathcal{W}_{uv}$ and by $\mathcal{W}^i_{uv}$ when considering paths of length i, hence $\sum_{v \in N(u)} \sum_{w \in \mathcal{W}^i_{uv}} x_w \geq p(i) \cdot |N(u)| \forall u \in V, i \leq t$. The second constraint (cf. Equation 4) makes sure that the obtained sparsified graph has no cycles *i.e.,* there is at most one path between any pair of vertices. This is enforced by setting the x_w to 1 to at most one path w among all possible paths $\mathcal{W}_{uv}$ between two adjacent vertices u and v in G. The constraint given by Equation 5 ensures that all edges belonging to a selected path w are selected in the sparsified graph. The last constraint (cf. Equation 6) sets the definition domain of x_e, x_w variables which are defined as binary variables.

Algorithm 1: *ptSpar*($G = (V, E)$, p, t, E^{π})

Input : $G = (V, E)$ a simple Graph, t an integer, $p : \mathbb{N} \rightarrow [0, 1]$, E^{π} an ordering of E
Output : $G_s = (V_s, E_s)$ a sparsified graph

1 $G_s = (V_s, E_s) \leftarrow (V, \emptyset)$;
2 $G' = (V', E') \leftarrow (V, \emptyset)$;
3 **for** $e = uv \in E^{\pi}$ **do**
4 $\quad E' \leftarrow E' \cup \{uv\}$;
5 $\quad$ insert $\leftarrow False$;
6 $\quad N^1_{G'}(u) \leftarrow$ direct neighbors of node u in G';
7 $\quad N^1_{G'}(v) \leftarrow$ direct neighbors of node v in G';
8 $\quad$ **for** $i = 1$ *to* t **do**
9 $\quad\quad N^i_{G_s}(u) \leftarrow$ neighbors of node u in G_s within at most i-hops;
10 $\quad\quad N^i_{G_s}(v) \leftarrow$ neighbors of node v in graph G_s within at most i-hops;
11 $\quad\quad$ **if** $|N^i_{G_s}(u) \cap N^1_{G'}(u)| < p(i)|N^1_{G'}(u)|$ *or* $|N^i_{G_s}(v) \cap N^1_{G'}(v)| < p(i)|N^1_{G'}(v)|$ **then**
12 $\quad\quad\quad$ insert $\leftarrow True$;
13 $\quad\quad\quad$ **Break**;
14 $\quad\quad$ **end**
15 $\quad$ **end**
16 $\quad$ **if** *insert* **then**
17 $\quad\quad E_s \leftarrow E_s \cup \{uv\}$;
18 $\quad$ **end**
19 **end**

4.2 Approximation Algorithm: ptSpar

We propose *ptSpar* (see Algorithm 1) an approximation algorithm that implements (p, t)-sparsification. It takes as input a graph $G = (V, E)$ to sparsify, the sparsification parameters p and t and an ordering E^{π} for processing the edges of the input graph.

The algorithm starts with an empty sparsified graph G_s (see line 1) and grows it incrementally by going through all the edges of the input graph G, in the order E^{π} (see the loop on lines 3 to 18). G' is a working variable initialized to the empty graph and serves to check that an edge inserted in G_s verifies the neighborhood conditions of the (p, t)- sparsification. Each iteration of the loop (lines 3 to 18) corresponds to the processing of a new edge e in the ordering E^{π}. A processed edge e, is first inserted into G' (line 4) but its inclusion in the sparsified graph G_s depends on whether it verifies the condition of (p, t)-sparsification. This is done by setting the variable *insert* to false (line 5).

To see whether edge $e = uv$ needs to be included in G_s, we simply check if G_s without the edge e remains a (p, t)-sparsification for G'. To do so, we check the neighborhood preservation constraints for nodes u and v as they are the only nodes whose neighborhood set is impacted by the arrival of edge e. For such purpose, we need to compute the set of all neighbors of u and v located in a radius <= i, i.e., $N^i_{G_s}(u)$ and $N^i_{G_s}(v)$ (see lines 9-10). To compute $N^i_{G_s}(u)$, respectively $N^i_{G_s}(v)$, we traverse the graph within radius i starting from u, respectively v. Then, we check if the non-insertion of e in G_s violates the neighborhood preservation constraints (line 11), if this is the case e must be inserted in G_s (lines 12-16).

Theorem 3. *The subgraph $G_s = (V_s, E_s)$ output of Algorithm ptSpar is a (p, t)-sparsification, of the input graph $G = (V, E)$.*

Proof Sketch. We proceed by induction and show that if $G_s(k)$ is a (p, t)-sparsification of $G'(k)$, then $G_s(k + 1)$ must also be a (p, t)-sparsification of $G'(k+1)$, where k denotes the iteration step in the algorithm, representing the stage at which the edges are processed. This is demonstrated by showing that an assumption of the contrary leads to a contradiction with the induction hypothesis, thereby confirming the theorem's claim through induction. □

The performance of the *ptSpar* algorithm is significantly impacted by the order in which edges are processed. Different edge orderings can lead to varying efficiencies of (p, t)-sparsification. Optimizing the (p, t)-sparsification for a graph G fundamentally involves finding the most effective edge ordering, $E^{\pi *}$, of the edges as stated in Theorem 4.

Theorem 4. *Let $G = (V, E)$ be a graph. There exists a permutation function π^* of the edge set E for which algorithm ptSpar, gives an optimal (i.e., a minimum size (p, t)-sparsification) of G.*

Proof Sketch. The proof begins by asserting that if at any iteration k, the current output $G_s(k)$ is a (p, t)-sparsification of G, subsequent edges processed will be rejected, maintaining $G_s(k)$ unchanged until the final iteration. This is based on the observation that once a graph meets the (p, t)-sparsification criteria, all its vertices have their neighborhood constraints satisfied, making any additional edge unnecessary. To demonstrate the theorem, we consider G^*_s, a minimum size (p, t)-sparsification of G, and construct π^* such that edges in E^*_s are processed before those in $E - E^*_s$. Under this ordering, once *ptSpar* processes all edges in E^*_s, it will reject any remaining edges, resulting in G^*_s as the output, proving the theorem. □

To improve the sparsification performance of the *ptSpar* algorithm, we propose three sub-optimal orders in the following subsections. Our aim is to approximate the ideal edge processing order thereby enhancing the sparsification effectiveness of the *ptSpar* algorithm. By exploring various edge ordering strategies, we seek to balance computational efficiency with the quality of the resultant sparsified graph. The first sub-optimal order is a random order, the second one is based on edge centrality and the third one relies on a meta-heuristic to find the best order.

Random edge order: This algorithm takes a graph as input and outputs an ordered edge set as illustrated by Algorithm 2. The order is determined by randomly selecting edges (cf. line 4).

Edge centrality based order: In this ordering, we first process the edges with a high centrality value. Centrality is a common measure for the importance of a node or an edge in a graph. The centrality we consider here is a relaxation of local edge betweenness defined in [17]. An edge with a high edge betweenness centrality represents a bridge-like connector between two parts of a graph, the removal of which may affect the shortest paths between these parts. The local edge betweenness of an edge e is the number of shortest paths running along e, the length of which is less than or equal to some constant t. In our proposed metric, We consider all paths with a length at most t, not just the shortest paths. Furthermore, we focus only on paths directly associated with an edge in the edge set E. For

Algorithm 2: Random order

Input : $G = (V, E)$ input Graph
Output : A randomly ordered edge set E^{rd}

1 $E^{rd} \leftarrow \emptyset$;
2 $E' \leftarrow E$;
3 **while** $E' \neq \emptyset$ **do**
4 pick randomly an edge e from E' ;
5 $E' \leftarrow E' - \{e\}$;
6 $E^{rd} \leftarrow E^{rd} \cup \{e\}$;
7 **end**
8 return E^{rd};

Algorithm 3: Edge centrality based order

Input : $G = (V, E)$ input Graph
Output : Centrality based edge order E^c

1 **for** $e \in E$ **do**
2 compute the score s(e) using Equation 7;
3 **end**
4 $E^c \leftarrow$ sort the edges of G in descending order according to $s(e)$;

every edge e, we calculate a centrality score $s(e)$ using Equation 7. In this equation, $\sigma_t(u, v|e)$ denotes the number of paths of length at most t that traverse edge e, and connect two nodes u and v that are directly linked by the edge uv. This approach ensures that the centrality score, we propose, accurately reflects the significance of each edge in connecting directly adjacent nodes within the network.

Once all scores are computed, we sort the edges in descending order according to their score $s(e)$ to obtain an edge ordering. Algorithm 3 formalizes the computation of this ordering.

$$s(e) = \sum_{uv \in E} \sigma_t(u, v|e) \; \forall uv \in E \tag{7}$$

Simulated-Annealing based order: We use Simulated Annealing (SA) [42] as a meta-heuristic to explore the different possible edge orderings. The SA-based ordering is detailed in Algorithm 4. Basically, the algorithm tries several edge orderings and keeps the best one, i.e., the one that produces the smallest sparsified graph. The simulated annealing process starts with a temperature value T_{max} (cf. line 1) that will keep decreasing at each iteration. This temperature is crucial in the SA technique, as it allows the acceptance of suboptimal solutions, particularly in the early stages when the temperature is high, to ensure a comprehensive exploration of the solution space (cf. lines 17-19). Then, the algorithm begins by a random order of the edges O_{init} of the input graph (line 2). It constructs an initial solution, i.e., a temporary sparsified graph G_{init} using the *ptSpar* algorithm (line 3) with this initial order. The number of edges in this initial solution $|E_{init}|$ is recorded in $COST_{best}$ (cf. line 4), tracking the best (i.e., smallest) edge count found in the sparsified graph so far. In the main iterative loop of the algorithm (lines 7-23), running for N iterations, each iteration modifies the current solution slightly by creating a new edge order O_{new} from O_{init} by swapping two randomly selected edges (cf. lines 6 and 7). Then, it evaluates this new edge ordering (cf. lines 8 to 16) by calling the *ptSpar* algorithm again to decide whether to accept this new order based on the size of the resulting sparsified graph. If the

Algorithm 4: Computing the best order with simulated annealing

Input : $G = (V, E)$ a Graph, t an integer, $p : \mathbb{N} \rightarrow [0, 1]$, N an integer (Number of iterations), T_{max} a double (Initial temperature), α a double (decreasing factor of the temperature)
Output : O_{best} the best edge ordering with N iterations

1 $T \leftarrow T_{max}$;
2 $O_{init} \leftarrow$ Random order of E;
3 $G_{init}(V, E_{init}) \leftarrow ptSpar(G, t, p, O_{init})$;
4 $COST_{init} \leftarrow |E_{init}|$;
5 $COST_{best} \leftarrow COST_{init}$;
6 $O_{best} \leftarrow O_{init}$;
7 **for** $i = 1$ *to* N **do**
8 $O_{new} \leftarrow$ Perturbing O_{init} by swapping the order of two random edges;
9 $G_{new}(V, E_{new}) \leftarrow ptSpar(G, t, p, O_{new})$;
10 $COST_{new} \leftarrow |E_{new}|$;
11 **if** $COST_{new} < COST_{best}$ **then**
12 $O_{best} \leftarrow O_{new}$;
13 $COST_{best} \leftarrow COST_{new}$;
14 **end**
15 **else**
16 $r \leftarrow$ random number between 0 and 1;
17 **if** $\exp(\frac{COST_{best} - COST_{new}}{T}) > r$ **then**
18 $O_{best} \leftarrow O_{new}$;
19 $COST_{best} \leftarrow COST_{new}$;
20 **end**
21 **end**
22 $T \leftarrow \alpha * T$;
23 **end**
24 return O_{best};

new ordering does not produce a better sparsified graph, it can still be accepted based on the current temperature (cf. lines 16-19). This will allow a broad exploration of the search space especially in the first stages where the value of T is high. The current temperature T is decreased at each iteration with a factor α, gradually lowering the likelihood of accepting worse solutions as the algorithm progresses (cf. line 17). Finally, the algorithm returns O_{best}, the best edge order found within the given number of iterations (cf. line 18).

Figure 2 illustrates the results of the (p, t)-sparsification algorithms on our running example (cf. Figure 2 (a)) with parameters $t = 2$, $p(1) = 50\%$ and $p(2) = 100\%$. For each algorithm, the retained edges are in blue and the removed ones in dashed grey. The exact algorithm produces the smallest possible sparsified graph with exactly 13 edges (cf. Figure 2 (b)). The *ptSpar* algorithm used with the random edge ordering (cf. Figure 2 (c)) and with the centrality-based ordering (cf. Figure 2 (d)) produces a near-optimal sparsification with 14 edges, which is very close to the theoretical optimal of 13 edges. It is important to note that although both orderings yield sparsifications with the same number of edges, the actual edges retained in each method are not always the same resulting in different sparsified graphs. Additionally, see that the SA algorithm succeeds to produce the optimal sparsification with 13 edges (cf. Figure 2 (e) and Figure 2 (f)) if it is used with a sufficient number of

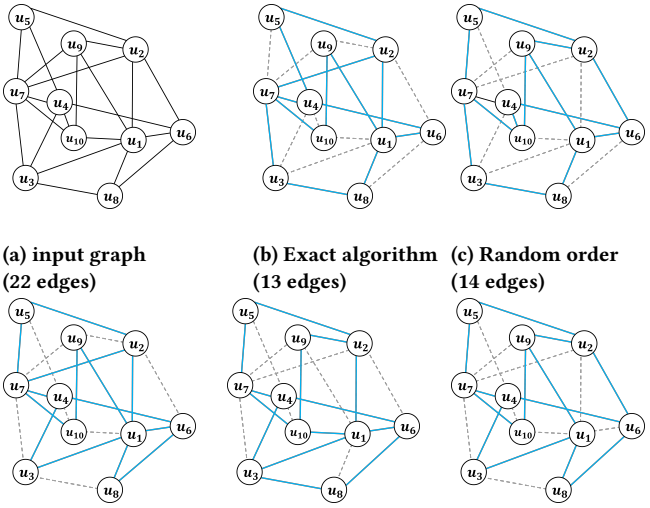

(a) input graph (22 edges) **(b) Exact algorithm (13 edges)** **(c) Random order (14 edges)**

(d) Centrality order (14 edges) **(e) SA 500 iterations (14 edges)** **(f) SA 1000 iterations (13 edges)**

Figure 2: Results of the (p, t)-sparsification algorithms on the running example.

iterations. Therefore, the SA algorithm is a promising approach for obtaining high-quality results that are close to the exact solution, although it does not scale well, as demonstrated in the following sections.

4.3 Complexity analysis

In this section, we analyze the time complexity of all the algorithms presented in the previous section. The theorems and their detailed proofs are provided in the supplementary material [24].

The time complexity of the *ptSpar* algorithm is $O(|E|d^t)$, where d represents the average degree in the graph G. In scenarios where every node is connected to every other node, forming a complete graph, this complexity escalates significantly, reaching $O(|E||V|^t)$.

The time complexity of the random edge ordering algorithm is $O(|E|)$. This complexity involves parsing the edge set E.

The average time complexity of the centrality based-ordering is $O(|E|(d^t + \log(|E|)))$. This complexity involves computing the score $s(e)$ for each edge e and sorting them. The complexity of computing the edge scores is equal to the complexity of listing all paths of length $\leq t$ for each pair of connected nodes, i.e., for each edge in E. The average number of paths of length $\leq t$ between two connected nodes u and v and starting from u is of order $O(d^t)$, where d is the average degree. The number of edges is $|E|$. Therefore, the time complexity of computing all the scores $s(e)$ is $O(|E|d^t)$. In addition, the complexity of sorting all the scores costs $O(|E| \log |E|)$. In the worst case (complete graph), the time complexity would be $O(|E|(V^t + \log(|E|)))$.

The time complexity of the Simulated Annealing (SA)-based ordering is influenced by the number of iterations N, the initial temperature T_0, and the temperature decreasing factor α. Each iteration involves a perturbation of the edge order and a re-evaluation of the sparsification, leading to a complexity that is also dependent on the efficiency of the *ptSpar* algorithm used within it.

This analysis clearly shows that the random order algorithm is the fastest. This explains its scalability when compared to the other edge orderings.

5 EXPERIMENTAL ANALYSIS

In this section, we present an experimental analysis of our sparsification approach. First, we evaluate the performance of the *ptSpar* algorithm provided to compute the sparsification. Then, we provide an analysis of the sensitivity of this sparsification to parameters p and t. Finally, we evaluate its effectiveness on several tasks such as shortest paths and reachability queries computation, node embedding and whole graph embedding. We also compare our sparsification with several baselines and state of the art methods to show its effectiveness. All experiments are carried out on an AMD 32 cores CPU with 768GB of memory. The source code of our algorithms is available at https://gitlab.liris.cnrs.fr/coregraphie/ptspar.

Algorithms. We used the following baselines in our comparative study. To guarantee a fair comparison with all the baselines, we have adopted the default configuration for each of them.

- Random Edge (RE): RE randomly eliminates a given percentage of edges.
- Local Degree (LD) [20]: LD retains the top $deg(v)^\alpha$ edges for each node $v \in V$, where $\alpha \in [0, 1]$
- Edge Forest Fire (EFF) [20]: EFF is based on the Forest Fire node sampling algorithm [28]. It initiates a fire at a random node and burns approximately $p/(1-p)$ neighbors, where p represents the probability threshold for burning a neighbor. Burnt neighbors are enqueued for subsequent fire initiation. EFF prunes edges based on the frequency of edge visits.
- Algebraic Distance (AD) [10]: AD uses random walk distance to compute the algebraic distance $\alpha(u, v)$ between two nodes. A low algebraic distance implies a high likelihood that a random walk starting from u will reach v within a small number of steps. It assigns an edge score of $1 - \alpha(u, v)$ to prioritize short-range edges.
- L-Spar (LS) [39]: LS employs the Jaccard similarity function on the adjacency lists of nodes u and v to determine the edge score of (u, v). It ranks edges locally (with respect to each node) and prunes them based on their ranks.
- Simmelian Backbone (SB) [34] : SB calculates weights by counting how many triangles each edge is part of, and then retains only those edges that form the most triangles, indicating strong and interconnected relationships in the graph. During sparsification, SB removes the lower-ranked edges of each node using a specified edge-prune ratio.
- Quadrilateral Simmelian Backbone (QSB) [35]: QSB measures the Simmelianness weight of each edge (u, v) by taking into account the shared quadrangles of u and v. It follows the same pruning strategy as SB.
- Salient backbone (SLB) [16] : SLB sparsifies a graph using the disparity filter which consists in calculating a statistical significance (p-value) for each edge based on its weight and the total weight of all edges connected to the same node. Edges with p-values below

Table 2: Characteristics of datasets used in our experiments.

Name	#graphs	$\|V\|$	$\|E\|$	Use case
BLOG-CATALOG	1	10.31K	333.98K	MLNC/SP/EO/EL
CA-ASTROPH	1	18.77K	198.11K	SP/EO/EL
CA-HEPTH	1	9.8K	25.9K	SP/EO/EL
CITESEER	1	3.2K	4.5K	NC/SP/EO/EL
COLLAB*	5000	372.5K	49.1M	GC
CORA	1	2.7K	5.4K	NC/SP/EL/EO
ENZYMES*	600	19.5K	74.6K	GC
FLICKR	1	89K	899K	NC/SP/EO
FLICKR-Large	1	80.51K	5.89M	MLNC
FRIENDSTER	1	65.6M	1.8B	AC/SP/EL
GSH-HOST	1	68.6M	1.8B	EO/SP/EL
IMDB-BINARY*	1000	19.77K	96.53K	GC/EL
MSRC-21C*	209	8.4K	20.2K	GC/EL
LIVEJOURNAL	1	3.99M	34.68M	EO/SP/EL
PROTEINS*	1113	43.5K	162.1K	GC/EL
PUBMED	1	19.7K	44.3K	EO/NC/SP/EL
SYNTH1	30	20	60	EO
SYNTH2	30	50	350	EO
SYNTH3	30	100	1.4K	EO
TWITTER	1	41.6M	1.4B	EO/SP/EL

MLNC: Multi-label Node Classification, SP: Shortest Paths, NC: Node Classification, GC: Graph Classification, EO: Edge ordering choice, EL: Entropy loss

*This dataset contains several labeled graphs for graph classification use cases.

a certain threshold are retained, forming the "backbone" of the graph. The rest, considered less important, are discarded.

- SparRL [44], a deep reinforcement learning-based method, sparsifies a graph by formulating the process as a Partially Observable Markov Decision Process (POMDP). It starts with a graph and at each step, chooses an edge to prune based on a policy learned from a Double DQN network. The policy is trained to maximize a reward function that encourages the preservation of certain graph properties.

Datasets. Table 2 summarizes the properties of the various datasets that we use in our extensive experiments. We use 20 datasets from various domains. In fact, for each application of our sparsification, we use the most used datasets for its evaluation. For scalability issues, we used two social networks TWITTER [1] and FRIENDSTER [46], and a Web graph GSH-HOST[2]. The COLLAB, ENZYMES, IMDB-BINARY, MSRC-21C, and PROTEINS datasets [22] contain several graphs and are labeled for graph classification use cases. The CORA, CITESEER, PUBMED, and FLICKR datasets [4, 47] have ground truth for node classification use cases. The synthetic graphs are used mainly to study the different ordering solutions we provided for the *ptSpar* algorithm.

Metrics. We use the following metrics:

- Sparsification runtime measured in seconds,
- Sparsification ratio that represents the ratio of the number of deleted edges over the total number of edges (see Equation 1), and
- Entropy loss to measure the information loss after the sparsification. The graph entropy is a measure of the structural information of a graph and serves as a complexity measure [12]. Given a graph $G(V, E)$, the Shanon entropy of G (i.e., $I(G)$) is computed as follows [12]:

$$I(G) = -\sum_{u \in V} \frac{deg(u)}{\sum_{u \in V} deg(u)} \log\left(\frac{deg(u)}{\sum_{u \in V} deg(u)}\right) \quad (8)$$

The entropy loss is the normalized difference between the entropy of the original graph and the entropy of the sparsified graph. Let G be the original graph and G_s be the sparsified graph, we compute the entropy loss as follows:

$$E_{loss} = \frac{|I(G) - I(G_s)|}{I(G)} \quad (9)$$

Note that the lower is the entropy loss the better is the sparsification.

For all the experiments, **TO** means that the algorithm has been timed out after a total time spent on the first terminating algorithm plus 3 hours.

5.1 Evaluating the edge ordering methods

In this subsection, we present a comparative experimental study of the edge orderings we considered for optimizing the *ptSpar* algorithm. These orderings are: random ordering, centrality based ordering and the Simulated-annealing (SA) ordering. The aim of these experiments is to show that the performance in term of sparsification ratio can be improved by considering different edge orderings. For this experiment, we use all the datasets except the labeled ones dedicated to classification only (indicated with a * in Table 2), and the following sparsification parameters $t = 2$, $p(1) = 0.0$ and $p(2) = 0.5$. For a reliable and accurate comparison, we carried-out around thirty tests on each dataset for each edge ordering solution. The results of the comparison are depicted in Table 3.

Note that the user configuration of the SA is $T_0 = 10$, $N = 1000$ and $\alpha = 0.99$. We notice that the centrality and the SA orderings outperform the random ordering of edges in terms of sparsification performance. The results clearly show that the centrality and the SA orderings are the best algorithms compared to the exact algorithm, which only finished to run on the two smallest datasets and timed out for all the remaining datasets. The centrality ordering seems really interesting and offers the best trade-off between sparsification performance and runtime. However, we can see that the *ptSpar* algorithm with a random order of edges is much faster than with the other orderings methods and can also be used on large datasets. Therefore, we will be using it in the rest of the experiments.

5.2 Evaluating the impact of the sparsification parameters p and t

In this series of experiments, we study the effect of parameters p and t on the sparsification performance.

As mentioned before, our sparsification allows users to control the trade-off between information loss and sparsification ratio (i.e., space gain). To do so, the user varies the parameters p and t according to its needs (available memory and the targeted use case) to find the configuration that suits him. The ideal scenario is to minimize the information loss (entropy loss) while maximizing the sparsification ratio. To this end, we define an utility function, which allows to compute the above trade-off as follows: $Tr = e^{\frac{-E_{loss}}{Sr}}$. This function allows us to select the best p values for a fixed t value. Once different t values have been computed, the final t value is chosen by ranking based on the utility function. Table 4 gives the sparsification ratio and entropy loss obtained by our sparsification on the CA-AstroPh dataset, while varying the neighborhood preservation proportion p.

Table 3: Evaluation of the *ptSpar* algorithm with different edge orderings.

Dataset	Random		Edge centrality		Simulated annealing		Exact	
	$\|E_s\|$	Time	$\|E_s\|$	Time	$\|E_s\|$	Time	$\|E_s\|$	Time
SYNTH1	28	**0.001s**	23.55	0.01s	**21.56**	0.5s	20.15	2m10s
SYNTH2	121.6	**0.008s**	**105.66**	0.02s	**105.9**	5.2s	100.14	4h41m
SYNTH3	367	**0.05s**	**323.2**	0.09s	340.4	40s	TO	TO
CITESEER	3180	**0.014s**	3277	0.027s	3070	8.15s	TO	TO
CORA	3358	**0.019s**	3441	0.038s	3237	12.86s	TO	TO
PUBMED	28094	**0.283s**	28977	0.501s	25167	1m50s	TO	TO
BLOG	169897	**4min8s**	171104	5m6s	TO	TO	TO	TO
CA-CATALOG	15623	**0.124s**	15976	0.239s	15416	2m29s	TO	TO
CA-HEPTH	104249	**9.257s**	104785	12.46s	104785	2h39m	TO	TO
LIVE JOURNAL	25.95 M	**2h17min**	24.64 M	22h10m	TO	TO	TO	TO
FRIENDSTER	1.38 B	**23h14min**	TO	TO	TO	TO	TO	TO
FLICKR	251525	**16.14s**	255524	35.92s	251003	4h33m	TO	TO
GSH-HOST	1.17 B	**21h9min**	TO	TO	TO	TO	TO	TO
TWITTER	816.3M	18h37min	TO	TO	TO	TO	TO	TO

We set $p(t) = 1$ in all experiments, which means that the whole initial neighborhood of each node can be retrieved in a neighborhood of radius $r = t$ at maximum. This ensures that reachability queries are fully preserved for all vertices. As expected, the sparsification ratio decreases (and the entropy loss increases) as the preserved proportion of neighborhood increases and vice-versa. Some of the values of the sparsification ratio obtained with the various combinations of parameters are very satisfactory. The same holds for the entropy loss (max value $< 5\%$). In addition, we remark that the sparsification ratio range is wide (from 7% to 75%) which confirms the possibility of effectively controlling the trade-off information loss/sparsification ratio using parameters p and t. The choice of the best configuration of parameters depends essentially on the nature of the graph to be sparsified and the user needs. Particularly, for this example, the configurations ($t = 2,\ p = (0.5, 1)$) and ($t = 3,\ p = (0.5, 0.7, 1)$) seem interesting and are a good trade-off between sparsification ratio and entropy loss with a sparsification ratio $> 45\%$ and an entropy loss $< 1\%$.

5.3 Evaluation the distribution of the shortest path lengths with (p, t)-Sparsification

In this experiment, we show that (p, t)-sparsification allows to approximate distances (shortest path lengths) between nodes. To do so, we sparsify three unweighted undirected graphs with the following combination of parameters $t = 2$, $p(1) = 0.5$, and $p(2) = 1.0$. Then, we compute all shortest paths between all nodes.

Figure 3 shows the distribution of the shortest path lengths in the original and sparsified graphs for the three datasets. We note that the two curves have almost the same pace. This shows that our sparsification preserves the distribution of the lengths of the shortest paths on the 3 datasets. However, the curves of the sparsified graphs are slightly stretched and shifted from the original curves. This is due to the stretching of the paths as a result of sparsification. This stretch is not really considerable because of the preservation of 50% of the direct neighbors of each vertex in the graph.
It is important to note here that we cannot draw a similar distribution for the shortest paths of other baseline methods, as they do not preserve the connectivity of the graph. However, we have evaluated the increase in the length of the shortest paths and the loss of reachability queries for these methods in Section 5.5.

5.4 Evaluating the information loss

In this series of experiments, we focus on evaluating the quality of the obtained sparsified graphs by measuring the loss of entropy for each method. Table 5 gives the information loss rates measured by the loss in entropy for all the sparsification methods on all datasets. For a rigorous comparison, we have selected datasets that have different densities and contain hundreds of graphs. It's worth noting that SparRL, which is a deep reinforcement learning-based method, was not applied in this context due to two main reasons: (1) It is computationally intensive and requires individual training for each graph. This is inefficient for datasets containing hundreds of graphs. (2) It has a reward function with a limited scope that does not consider entropy – a key factor in this case.

We can observe that none of the baselines excels on all datasets. However, on average (p,t)-sparsification outperforms all the baselines. The reasons for that are due to compressibility properties of networks [32], such as transitivity and degree heterogeneity, that are leveraged by our sparsification technique. Graphs with higher transitivity and degree heterogeneity are generally more compressible. This explains why the results are favorable for our sparsification method compared to baselines that do not affect those properties. The (p, t)-sparsification method proves to be effective and high-quality across multiple datasets, including large-scale ones. Indeed, it shows, on average, the least entropy loss compared to the baselines. This indicates that the (p, t)-sparsification method is capable of preserving most of the graph's original information, proving it to be the best choice for various use cases.

5.5 Evaluating the usefulness of the sparsified graphs

We have applied many graph algorithms on the sparsified graphs. Our first motivation is to be able to use these algorithms directly on the sparsified graphs with acceptable performance So, the purpose of the following experiments is to show the effectiveness of our sparsification in terms of handling large graphs and providing good approximations of the original results of the considered tasks. For this, and for all the following experiments, we compute a new

Table 4: Sparsification ratio vs entropy loss with different combinations of parameters p and t.

				CA-ASTROPH			PUBMED			FLICKER		
t	$p(1)$	$p(2)$	$p(3)$	Sr	E_{loss}	Tr	Sr	E_{loss}	Tr	Sr	E_{loss}	Tr
2	0.2	1.0	-	58.13%	1.71%	0.971	11%	0.73%	**0.935**	5.3%	0.23%	**0.958**
	0.5	1.0	-	45.82%	0.90%	**0.980**	10.0%	0.70%	**0.935**	5.2%	0.25%	0.953
	0.7	1.0	-	26.39%	0.66%	0.975	8.2%	0.70%	0.918	4.5%	0.23%	0.950
	0.9	1.0	-	7.43%	0.31%	0.959	2.2%	0.26%	0.888	1.4%	0.17%	0.885
3	0.0	0.2	1.0	75.00%	4.61%	0.940	29.0%	1.70%	**0.943**	30.2%	1.64%	**0.947**
	0.2	0.5	1.0	71.50%	2.57%	0.964	27.9%	1.68%	0.941	25.7%	1.44%	0.945
	0.5	0.7	1.0	46.73%	0.85%	**0.981**	19.5%	1.15%	0.942	19.1%	1.05%	0.946
	0.7	0.9	1.0	26.43%	0.66%	0.975	9.3%	0.85%	0.913	6.5%	0.47%	0.930

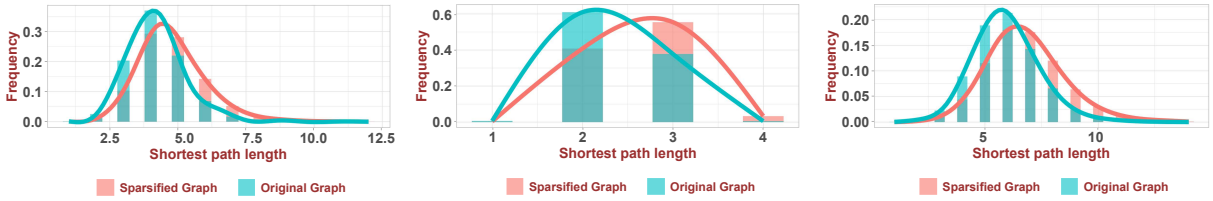

(a) CA-ASTROPH dataset **(d) BLOG-CATALOG dataset** **(c) CA-HEPTH dataset**

Figure 3: Distribution of shortest path lengths of the original and sparsified graphs.

Table 5: Effect of the sparsification on the entropy loss

Dataset	ptSpar	SLB	AD	LS	QSB	SB	EFF	LD	RE
COLLAB	**1.00%**	20.80%	11.80%	7.70%	2.90%	1.70%	7.40%	21.30%	6.80%
IMDB-BINARY	1.50%	22.70%	3.70%	6.40%	1.20%	**1.10%**	6.60%	19.40%	5.90%
MSRC_21C	**0.60%**	4.40%	1.50%	**0.60%**	1.70%	1.80%	2.90%	6.40%	2.70%
PROTEINS	1.50%	4.10%	4.90%	**1.30%**	5.30%	4.40%	3.80%	4.80%	2.50%
PUBMED	0.71%	0.29%	0.52%	**0.07%**	0.60%	0.63%	0.78%	0.57%	0.88%
CITESEER	0.65%	0.77%	0.27%	**0.17%**	1.00%	1.08%	1.87%	0.56%	1.50%
CA-HEPTH	**0.62%**	4.28%	1.43%	0.63%	1.75%	1.77%	3.07%	6.57%	2.66%
CORA	0.40%	0.51%	**0.14%**	0.21%	1.12%	1.18%	1.48%	1.30%	1.58%
FLICKR	0.26%	0.48%	**0.07%**	0.24%	**0.07%**	**0.07%**	0.47%	0.44%	0.72%
LIVE JOURNAL	0.78%	OT	2.70%	1.09%	**0.31%**	0.64%	2.48%	0.46%	1.73%
CA-ASTROPH	0.90%	0.94%	2.89%	1.56%	**0.75%**	0.98%	3.15%	2.84%	1.85%
BLOG-CATALOG	1.67%	TO	**0.10%**	1.28%	1.60%	4.45%	3.76%	6.74%	7.26%
ENZYMES	1.58%	21.80%	4.00%	6.70%	1.50%	**1.45%**	8.30%	17.93%	5.87%
FRIENDSTER	0.71%	TO	0.62%	**0.48%**	0.60%	0.63%	0.78%	1.26%	0.68%
GSH-HOST	0.69%	TO	0.78%	**0.61%**	0.84%	0.92%	0.99%	1.13%	0.79%
TWITTER	1.13%	TO	0.90%	**0.83%**	0.87%	1.1%	1.37%	1.85%	1.65%
Average	**0.91%**	-	2.48%	1.95%	1.56%	1.57%	3.11%	5.68%	2.9%

metric in addition to the sparsification ratio namely:

Performance Preservation: This metric measures the degree to which the performance of the considered task on the sparsified graph approaches the one obtained on the original graph. In contexts like classification, it gauges the relative preservation of accuracy from the original graph to its sparsified counterpart. It is computed as the ratio between the accuracy on the sparsified graph and the accuracy on the original graph. A higher value of this metric signifies that the sparsified graph has effectively retained, or closely approximated, the properties of the original graph.

For the graph kernel task, we also observed a speed-up of the algorithm. We measure this by the ratio between the algorithm run-time on the original graph and its run-time on the sparsified graph. The higher the speed-up factor, the faster the graph kernel algorithm on the sparsified graph. For the remaining tasks, the speed-up is always equal to 1.

Shortest Paths and Reachability Queries: To assess the impact of sparsification on shortest paths, we calculated two metrics: (1) the average increase in shortest path length between 10,000 randomly selected node pairs. A lower value indicates that the sparsification method more effectively preserves the original graph's shortest path lengths. (2) the failure rate, which represents the percentage of node pairs that became disconnected (i.e., unreachable from each other) in the sparsified graph. This metric reflects the success rate of reachability queries in the sparsified graph, with lower values indicating better performance.

The results are presented in Table 6. For each baseline, the first given value shows the average increase in shortest path lengths,

while the value in parentheses represents the failure rate of reachability queries. We can clearly observe that the (p, t)-sparsification technique demonstrates consistent and effective performance across all datasets, including large-scale ones. Indeed, it shows, on average, the best result in terms of preserving the shortest paths between randomly chosen pairs of nodes, thus maintaining the structural integrity of the original graphs to a great extent. Furthermore, it achieves a zero percent failure rate on all datasets, which shows that the sparsified graphs generated using (p, t)-sparsification are highly connected, preserving the reachability between nodes effectively. The LD method also performs competitively on shortest path preservation but fails to maintain the connectivity of the obtained sparsified graphs leading to a high failure rates in reachability queries. The SLB method, which has the lowest increase on shortest paths for small datasets, does not scale on large ones. The other baselines exhibit a higher increase in path length and a higher failure rate across most datasets, implying a greater degree of distortion in the sparsified graphs.

Graph kernels: Graph kernels predominantly rely on local neighborhood information of nodes. Such methods derive graph representations by delving deep into node neighborhoods and extracting pertinent features, encompassing walks, shortest paths, and other local substructures. Given that our (p, t)-sparsification meticulously retains the local neighborhood up to a radius t, a pertinent question arises: can graph kernels algorithms accelerate on sparsified graphs without significant compromise on performance? To answer this inquiry, we run a graph classification task on graph classification datasets, namely COLLAB, IMDB-BINARY, MSRC-21C, and PROTEINS. Here, we set $t = 3$, $p(1) = 0.0$, $p(2) = 0.5$ and $p(3) = 1.0$. On these sparsified datasets, we executed various graph embedding algorithms, including the Shortest Path graph kernel (SP) [9], Weisfeiler-Lehman Optimal Assignment WL-OA graph kernel [26, 40], The Neighborhood Hash NH graph kernel [21] and deep Renyi entropy graph kernel (REK) [45]. We gauged the efficacy of these algorithms on both input graph and sparsified graphs. We used the SVM algorithm as a classifier, with a 10-fold cross-validation, serving as our performance metric. For the sake of fairness, all baseline methods maintained an identical sparsification ratio. The SparRL method was omitted from this evaluation due to its innate latency and the requisite training for each graph, proving inefficient given the multiplicity of graphs in our datasets.

Table 7 shows the performance of graph kernels on the sparsified graphs. We notice that all kernels run faster on sparsified graphs in most cases. This Kernel computation speed-up is more noticeable on denser datasets such as COLLAB. Since the sparsified graphs produced by all methods are of the same size (same sparsification ratio for fair comparison), the speed up factors are the same for all methods. However, the performance preservation of graph kernel methods on sparsified graphs provides pivotal insights into the robustness and efficacy of different sparsification approaches, particularly emphasizing our (p, t)-sparsification method. Across a diverse range of datasets, our method's performance, in many instances, either leads the cohort or remains competitively in line with the best-performing methods. For example, in the COLLAB dataset with the Shortest Path (SP) graph kernel, our method reaches a performance preservation of 100%, a performance matched only by SLB and LD, while outperforming other benchmarks such as SB, EFF, and LS. This level of consistency in preserving the integrity of the original graph structure continues across various kernels like WL, NH, and REK. What is particularly notable is the general outperformance of our method when contrasted against SB in datasets like MSRC_21C using the REK kernel, where our method achieves a perfect score of 100% versus SB's 20%. Yet, our method proves its mettle even in situations where it doesn't lead but exhibits comparable performance, such as in the case of the WL kernel in the same dataset. While LD achieved the top score of 92%, our method's 88% was closely aligned, demonstrating its competitiveness. However, it's worth noting that while our method isn't always the definitive leader across all datasets and kernels, it consistently ranks among the top contenders, rarely deviating far from the highest scores. The associated speed-up rates also underscore the computational advantages of our sparsification approach. To sum up, the (p, t)-sparsification method we propose serves as a powerful tool, often leading in performance preservation and, when not, still staying well ranked within the top performing methods across a wide range of datasets and graph kernels.

Node embedding: In this series of experiments, we use sparsified graphs to compute node embedding. Then, to see if the obtained embeddings are as relevant as the ones computed on the full graph, we evaluate their efficacy in two tasks: node classification and multi-label classification. We leveraged two predominant algorithms for this endeavor: the Graph Attention Network (GAT) [43] for node classification, and Node2vec [18] for multi-label classification. The experiments are conducted on graphs with distinct sparsification ratios: 45% for the multi-label task and 20% for the graph classification task. It's worth noting that, for multi-label classification on the Flickr-large dataset, we have not included the results of the Salient Backbone (SLB) method because it failed to sparsify this large graph within the time limit of 24 hours. Additionally, the SparRL method is not present in evaluations for both tasks because it is not possible to express classification within its objective function.

Table 8 presents the outcomes of the different sparsification methods. Across all datasets, we can see that (p, t)sparsification consistently excelled. On datasets such as 'PROTEINS', (p, t)sparsification achieves an almost impeccable accuracy preservation rate of 99.68%. Similarly, in the 'Cora' and 'Flickr' datasets, it achieves the impressive rates of 96.97% and 99.16%. These figures attest to the technique's proficiency in preserving critical graph structures essential for GAT. LD and LS also delivered interesting outcomes in certain datasets but QSB and SB reported suboptimal results, further underscoring the significance of (p, t)-sparsification's results.

Table 9 presents the results of multi-label node classification using node2vec embedding on the sparsified graphs. We can clearly see the out-performance of (p, t)-Sparsification. In the Blog Catalog dataset, (p, t)-Sparsification achieves 93.03% for Micro F1 and 90.75% for Macro F1 metrics. Its excellent performance is further underscored in the Flickr dataset, with a perfect 100% in both metrics. This clearly shows the robustness of (p, t)-Sparsification in retaining crucial graph properties vital for node2vec.The other methods have much less effective results.

Table 6: Performance on shortest paths and reachability queries.

Dataset	AD	LS	QSB	SB	EFF	LD	RD	ptSpar	SLB
CORA	14.1% (55.6%)	46.3% (19.6%)	15.1% (41.4%)	15.9% (41.1%)	8.0% (22.8%)	7.0% (2.7%)	14.1% (18.1%)	2.7% **(0.0%)**	**2.3%** (4.1%)
CITESEER	2.4% (92.8%)	44.5% (51.2%)	15.1% (46.9%)	12.4% (46.0%)	1.6% (22.0%)	13.9% (20.0%)	12.0% (28.9%)	1.7% **(0.0%)**	**1.5%** (5.9%)
PUBMED	19.6% (61.0%)	23.5% (5.0%)	7.8% (39.6%)	7.3% (39.9%)	5.5% (41.7%)	4.7% (0.4%)	8.3% (25.6%)	2.2% **(0.0%)**	**1.3%** (5.2%)
FLICKR	4.7% (36.0%)	6.3% (0.0%)	9.2% (0.3%)	9.0% (0.3%)	19.3% (6.0%)	4.4% (0.0%)	10.7% (0.0%)	2.8% **(0.0%)**	**1.9%** (0.5%)
CA-HEPTH	13.8% (51.6%)	44.9% (19.0%)	15.7% (43.1%)	16.0% (42.3%)	8.2% (21.7%)	7.6% (3.1%)	15.6% (29.1%)	**2.6% (0.0%)**	**2.6%** (5.0%)
LIVEJOURNAL	46.56% (54.94%)	22.99% (2.56%)	39.77% (52.43%)	41.97% (52.35%)	17.52% (4.69%)	**4.94%** (4.86%)	10.24% (17.87%)	6.8% **(0.0%)**	TO
CA-ASTROPH	14.67% (69.85%)	46.52% (11.345%)	52.9% (50.68%)	52.7% (52.78%)	20.62% (3.16%)	**8.77%** (6.1%)	15.25% (13.23%)	15.09% **(0.0%)**	21% (15.3%)
BLOG-CATALOG	44.52% (11.41%)	25.97% (0.0%)	44.03% (33.25%)	36.31% (60.78%)	27.21% (6.97%)	**7.34%** (5.4%)	13.30% (8.54%)	11.14% **(0.0%)**	TO
FRIENDSTER	2.2% (0.5%)	1.25% (0.8%)	1.65% (0.83%)	1.74% (1.76%)	1.95% (0.72%)	1.43% (0.63%)	4.6% (0.92%)	**1.2% (0.0%)**	TO
GSH-HOST	1.42% (0.05%)	1.35% (0.25%)	1.67% (0.80%)	1.80% (1.70%)	2.2% (1.65%)	**0.90%** (0.90%)	1.65% (0.85%)	1.25% **(0.0%)**	TO
TWITTER	3.15%(0.1%)	2.78%(0.48%)	2.95%(2.1%)	2.64%(3.4%)	**1.95%**(1.2%)	3.2%(0.68%)	4.28%(0.75%)	2.24%**(0.0%)**	TO
Average	15.19% (39.44%)	24.21% (10.02%)	18.72 (28.31%)	17.98% (31.12%)	10.37% (12.05%)	5.83% (4.07%)	10.00%(13.08%)	**4.52% (0.00%)**	-

Table 7: Graph kernel performance on the sparsified graphs.

Dataset Sr	Kernel	Speed up	Performance Preservation					
			EFF	LD	LS	ptSpar	SB	SLB
COLLAB 91.4%	SP	2.75	97%	**100%**	90%	**100%**	95%	**100%**
	WL	1.23	86%	**92%**	87%	88%	88%	85%
	NH	1.54	83%	**88%**	84%	87%	84%	83%
	REK	1.86	83%	78%	86%	**88%**	87%	65%
IMDB-BINARY 72.2%	SP	1.14	96%	**100%**	81%	100%	93%	99%
	WL	1	93%	**96%**	90%	95%	89%	89%
	NH	1.11	91%	92%	83%	**94%**	89%	87%
	REK	1.46	91%	81%	93%	95%	**97%**	68%
MSRC_21C 46.8%	SP	1.04	97%	97%	89%	**100%**	89%	99%
	WL	1.24	98%	**100%**	93%	**100%**	95%	24%
	NH	1.34	97%	95%	91%	**100%**	94%	24%
	REK	1.15	95%	96%	99%	**100%**	99%	20%
PROTEINS 36.1%	SP	1.39	98%	**100%**	92%	**100%**	96%	**100%**
	WL	1.2	96%	96%	89%	**97%**	95%	94%
	NH	1.12	98%	**99%**	94%	**99%**	98%	95%
	REK	1.86	97%	**96%**	99%	**99%**	97%	80%

Table 8: Performance of Node classification on sparsified graphs.

Method	CORA	CITESEER	PUBMED	FLICKR
AD	90.47%	87.07%	89.32%	92.15%
EFF	71.17%	55.51%	81.87%	99.01%
LD	96.80%	**96.75**%	99.65%	99.58%
LS	94.81%	96.26%	97.94%	97.42%
ptSpar	**96.97%**	94.53%	**99.68%**	**99.16%**
QSB	46.66%	36.90%	52.59%	93.08%
RE	90.47%	85.10%	88.31%	94.02%
SB	47.27%	37.34%	53.76%	93.56%
SLB	87.3%	78.40%	88.83%	92.72%

6 CONCLUSION AND FUTURE WORK

In this paper, we presented a graph sparsification approach designed to produce a graph skeleton that can be used instead of the original large graph as input in many graph analysis algorithms. To do so, our sparsification controls the amount of neighborhood information preserved in the resulting sparsified graph with two parameters: a function p that gives the proportion of each node's original neighbors to be preserved in its i-hops neighborhood in the sparsified graph, and a threshold t for which p reaches its maximal value. We also presented several algorithms to compute this sparsification with the minimum cost, and showed their effectiveness in sparsifying input graphs through an extensive experimental evaluation on multiple real-life as well as synthetic graph datasets. Furthermore, we showed that the skeletons computed by the proposed approach can be used without any addition or de-sparsification as input to multiple graph applications, such as node embedding, graph classification, and shortest path approximations.As for future work, we consider a more thorough analysis of (p, t)-sparsification impact on walk based graph learning algorithms such as Node2vec and DeepWalk. In fact, we observed some situations where the learning accuracy increased when the graph was sparsified. This was a quite unexpected observation. While we guess that walks are biased in the right direction by removing edges, characterizing such edges remains an open question. Another important open question is to find an efficient method to order graph edges. This would allow us to significantly improve the time complexity of the approach. In addition, we aim to design an incremental version of our sparsification to deal with dynamic graphs or graph streams.

We note also that our approach can be used on both directed and undirected graphs. However, our sparsification does not consider the labels of the edges. To sparsify or compress edge-labeled graphs, a new model need to be defined so as to take into account these labels for example when defining the edge ordering.

Table 9: Performance of multi-label node classification on sparsified graphs.

Method	BLOG-CATALOG		FLICKR-LARGE	
	Micro F1 %	Macro F1 %	Micro F1 %	Macro F1 %
AD	35.84 %	15.23 %	43.11 %	28.6 %
EFF	36.48 %	16.61 %	44.01 %	29.6%
LD	37.74 %	14.56 %	44.22 %	28.2%
LS	35.46 %	16.69 %	44.43 %	29.6 %
ptSpar	**93.03 %**	**90.75** %	**100%**	**100%**
QSB	36.54 %	14.31 %	45.86 %	31.5 %
RE	35.37 %	13.05%	44.70 %	45%
SB	38.35 %	14.35 %	44.49%	24.9%

ACKNOWLEDGMENTS

This work was supported by Agence Nationale de la Rechetche (ANR) under grant number ANR-20-CE23-0002. The authors thank Walid Megherbi for his help.

REFERENCES

[1] 2010. *Twitter 2010 dataset.* https://snap.stanford.edu/data/twitter-2010.html Accessed: 2024-10-26.

[2] 2015. *Gsh host dataset.* http://law.di.unimi.it/webdata/gsh-2015-host Accessed: 2024-10-26.

[3] András A. Benczúr and David R. Karger. 1996. Approximating s-t Minimum Cuts in $O(n^2)$ Time. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (Philadelphia, Pennsylvania, USA) *(STOC '96)*. Association for Computing Machinery, New York, NY, USA, 47–55. https://doi.org/10.1145/237814.237827

[4] Aleksandar Bojchevski and Stephan Günnemann. 2017. Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking. *arXiv preprint arXiv:1707.03815* (2017).

[5] P. Boldi and S. Vigna. 2004. The Webgraph Framework I: Compression Techniques. In *Proceedings of the 13th International Conference on World Wide Web* (New York, NY, USA) *(WWW '04)*. Association for Computing Machinery, New York, NY, USA, 595–602. https://doi.org/10.1145/988672.988752

[6] Francesco Bonchi, Gianmarco De Francisci Morales, Aristides Gionis, and Antti Ukkonen. 2013. Activity preserving graph simplification. *Data Mining and Knowledge Discovery* 27, 3 (Nov. 2013), 321–343. https://doi.org/10.1007/s10618-013-0328-8

[7] Angela Bonifati, Stefania Dumbrava, and Haridimos Kondylakis. 2020. Graph Summarization. *CoRR* abs/2004.14794 (2020). arXiv:2004.14794 https://arxiv.org/abs/2004.14794

[8] Angela Bonifati, Stefania Dumbrava, Haridimos Kondylakis, Georgia Troullinou, and Giannis Vassiliou. 2025. Progressive Querying on Knowledge Graphs. In *Proceedings 28th International Conference on Extending Database Technology, EDBT 2025, Barcelona, Spain, March 25-28, 2025*. 106–118.

[9] Karsten M Borgwardt and Hans-Peter Kriegel. 2005. Shortest-path kernels on graphs. In *Fifth IEEE international conference on data mining (ICDM'05) (ICDM '05)*. IEEE, 8–pp.

[10] Jie Chen and Ilya Safro. 2011. Algebraic distance on graphs. *SIAM Journal on Scientific Computing* 33, 6 (2011), 3468–3490.

[11] Michele Coscia and Frank MH Neffke. 2017. Network backboning with noisy data. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*. IEEE, 425–436.

[12] Matthias Dehmer and Abbe Mowshowitz. 2011. A history of graph entropy measures. *Information Sciences* 181, 1 (2011), 57–78.

[13] Stefania Dumbrava, Angela Bonifati, Amaia Nazabal Ruiz Diaz, and Romain Vuillemot. 2019. Approximate Querying on Property Graphs. In *Scalable Uncertainty Management*, Nahla Ben Amor, Benjamin Quost, and Martin Theobald (Eds.). Springer International Publishing, Cham, 250–265.

[14] David Durfee, John Peebles, Richard Peng, and B.Rao Anup. 2020. Determinant-Preserving Sparsification of SDDM Matrices. *SIAM J. Comput.* 49, 4 (2020), FOCS17–350–FOCS17–408. https://doi.org/10.1137/18M1165979

[15] Wenfei Fan, Yuanhao Li, Muyang Liu, and Can Lu. 2021. Making Graphs Compact by Lossless Contraction. In *Proceedings of the 2021 International Conference on Management of Data* (Virtual Event, China) *(SIGMOD '21)*. Association for Computing Machinery, New York, NY, USA, 472–484. https://doi.org/10.1145/3448016.3452797

[16] Daniel Grady, Christian Thiemann, and Dirk Brockmann. 2012. Robust classification of salient links in complex networks. *Nature communications* 3, 1 (2012), 1–10.

[17] Steve Gregory. 2008. A Fast Algorithm to Find Overlapping Communities in Networks. In *Machine Learning and Knowledge Discovery in Databases*, Walter Daelemans, Bart Goethals, and Katharina Morik (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 408–423.

[18] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining (KDD '16)*. 855–864.

[19] David S. Gunderson. 2014. *Handbook of Mathematical Induction: Theory and Applications.* CRC Press.

[20] Michael Hamann, Gerd Lindner, Henning Meyerhenke, Christian L Staudt, and Dorothea Wagner. 2016. Structure-preserving sparsification methods for social networks. *Social Network Analysis and Mining* 6 (2016), 1–22.

[21] Shohei Hido and Hisashi Kashima. 2009. A linear-time graph kernel. In *2009 Ninth IEEE International Conference on Data Mining (ICDM '09)*. IEEE, 179–188.

[22] Kristian Kersting, Nils M. Kriege, Christopher Morris, Petra Mutzel, and Marion Neumann. 2016. Benchmark Data Sets for Graph Kernels. http://graphkernels.cs.tu-dortmund.de Accessed: 2024-10-26.

[23] Dong-Hee Kim, Jae Dong Noh, and Hawoong Jeong. 2004. Scale-free trees: The skeletons of complex networks. *Physical Review E* 70, 4 (2004), 046126.

[24] Abd Errahmane Kiouche, Julien Baste, Mohammed Haddad, Hamida Seba, and Angela Bonifati. 2024. http://gitlab.liris.cnrs.fr/coregraphie/ptspar/-/blob/main/Supplementary_material.pdf. Accessed: 2024-10-26.

[25] Danai Koutra, U Kang, Jilles Vreeken, and Christos Faloutsos. 2015. Summarizing and understanding large graphs. *Statistical Analysis and Data Mining: The ASA Data Science Journal* 8, 3 (2015), 183–202. https://doi.org/10.1002/sam.11267 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/sam.11267

[26] Nils M Kriege, Pierre-Louis Giscard, and Richard Wilson. 2016. On valid optimal assignment kernels and applications to graph classification. *Advances in neural information processing systems* 29 (2016).

[27] Sofiane Lagraa, Hamida Seba, Riadh Khennoufa, Abir MBaya, and Hamamache Kheddouci. 2014. A distance measure for large graphs based on prime graphs. *Pattern Recognition* 47, 9 (2014), 2993–3005. https://doi.org/10.1016/j.patcog.2014.03.014

[28] Jure Leskovec and Christos Faloutsos. 2006. Sampling from large graphs. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD '06)*. 631–636.

[29] Yuzhi Liang, Chen chen, Yukun Wang, Kai Lei, Min Yang, and Ziyu Lyu. 2020. Reachability preserving compression for dynamic graph. *Information Sciences* 520 (2020), 232–249. https://doi.org/10.1016/j.ins.2020.02.028

[30] Y. Liu, T. Safavi, A. Dighe, and D. Koutra. 2018. Graph Summarization Methods and Applications: A Survey. *ACM Comput. Surv.* 51, 3, Article 62 (June 2018), 34 pages. https://doi.org/10.1145/3186727

[31] Yike Liu, Tara Safavi, Abhilash Dighe, and Danai Koutra. 2018. Graph summarization methods and applications: A survey. *ACM computing surveys (CSUR)* 51, 3 (2018), 1–34.

[32] Christopher W. Lynn and Danielle S. Bassett. 2021. Quantifying the compressibility of complex networks. *Proceedings of the National Academy of Sciences* 118, 32 (2021), e2023473118. https://doi.org/10.1073/pnas.2023473118 arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.2023473118

[33] Michael Mathioudakis, Francesco Bonchi, Carlos Castillo, Aristides Gionis, and Antti Ukkonen. 2011. Sparsification of Influence Networks. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Diego, California, USA) *(KDD '11)*. Association for Computing Machinery, New York, NY, USA, 529–537. https://doi.org/10.1145/2020408.2020492

[34] Bobo Nick, Conrad Lee, Pádraig Cunningham, and Ulrik Brandes. 2013. Simmelian backbones: Amplifying hidden homophily in facebook networks. In *Proceedings of the 2013 IEEE/ACM international conference on advances in social networks analysis and mining*. 525–532.

[35] Arlind Nocaj, Mark Ortmann, and Ulrik Brandes. 2014. Untangling hairballs: From 3 to 14 degrees of separation. In *International symposium on graph drawing*. Springer, 101–112.

[36] L. Paul Chew. 1989. There are planar graphs almost as good as the complete graph. *J. Comput. System Sci.* 39, 2 (1989), 205–219. https://doi.org/10.1016/0022-0000(89)90044-5

[37] David Peleg and Alejandro A. Schäffer. 1989. Graph spanners. *Journal of Graph Theory* 13, 1 (1989), 99–116. https://doi.org/10.1002/jgt.3190130114 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/jgt.3190130114

[38] Sherif Sakr, Angela Bonifati, Hannes Voigt, Alexandru Iosup, Khaled Ammar, Renzo Angles, Walid G. Aref, Marcelo Arenas, Maciej Besta, Peter A. Boncz, Khuzaima Daudjee, Emanuele Della Valle, Stefania Dumbrava, Olaf Hartig, Bernhard Haslhofer, Tim Hegeman, Jan Hidders, Katja Hose, Adriana Iamnitchi, Vasiliki Kalavri, Hugo Kapp, Wim Martens, M. Tamer Özsu, Eric Peukert, Stefan Plantikow, Mohamed Ragab, Matei Ripeanu, Semih Salihoglu, Christian Schulz, Petra Selmer, Juan F. Sequeda, Joshua Shinavier, Gábor Szárnyas, Riccardo Tommasini, Antonino Tumeo, Alexandru Uta, Ana Lucia Varbanescu, Hsiang-Yun Wu, Nikolay Yakovets, Da Yan, and Eiko Yoneki. 2021. The future is big graphs: a community view on graph processing systems. *Commun. ACM* 64, 9 (2021), 62–71. https://doi.org/10.1145/3434642

[39] Venu Satuluri, Srinivasan Parthasarathy, and Yiye Ruan. 2011. Local graph sparsification for scalable clustering. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data (KDD '11)*. 721–732.

[40] Nino Shervashidze, Pascal Schweitzer, Erik Jan Van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. 2011. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research* 12, 9 (2011).

[41] Daniel A. Spielman and Shang-Hua Teng. 2011. Spectral Sparsification of Graphs. *SIAM J. Comput.* 40, 4 (2011), 981–1025. https://doi.org/10.1137/08074489X

[42] Peter J. M. van Laarhoven and Emile H. L. Aarts. 1987. *Simulated annealing.* Springer Netherlands, Dordrecht, 7–15. https://doi.org/10.1007/978-94-015-7744-1_2

[43] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, Yoshua Bengio, et al. 2017. Graph attention networks. *stat* 1050, 20 (2017), 10–48550.

[44] Ryan Wickman, Xiaofei Zhang, and Weizi Li. 2022. A Generic Graph Sparsification Framework using Deep Reinforcement Learning. In *2022 IEEE International Conference on Data Mining (ICDM)*. 1221–1226. https://doi.org/10.1109/ICDM54844.2022.00158

[45] Lixiang Xu, Lu Bai, Xiaoyi Jiang, Ming Tan, Daoqiang Zhang, and Bin Luo. 2021. Deep Rényi entropy graph kernel. *Pattern Recognition* 111 (2021), 107668.

[46] Jaewon Yang and Jure Leskovec. 2015. Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems* 42, 1 (01 Jan 2015), 181–213. https://doi.org/10.1007/s10115-013-0693-z

[47] Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. 2019. Graphsaint: Graph sampling based inductive learning method. *arXiv preprint arXiv:1907.04931* (2019).

[48] Ronda J Zhang, H Eugene Stanley, and Fred Y Ye. 2018. Extracting h-backbone as a core structure in weighted networks. *Scientific reports* 8, 1 (2018), 1–7.