

# Multi-relation Graph Summarization

XIANGYU KE, Nanyang Technological University, Singapore

ARIJIT KHAN, Nanyang Technological University, Singapore

FRANCESCO BONCHI, ISI Foundation, Italy and Eurecat, Spain

Graph summarization is beneficial in a wide range of applications, such as visualization, interactive and exploratory analysis, approximate query processing, reducing the on-disk storage footprint, and graph processing in modern hardware. However, the bulk of the literature on graph summarization surprisingly overlooks the possibility of having edges of different types. In this paper, we study the novel problem of producing summaries of multi-relation networks, i.e., graphs where multiple edges of different types may exist between any pair of nodes. Multi-relation graphs are an expressive model of real-world activities, in which a relation can be a topic in social networks, an interaction type in genetic networks, or a snapshot in temporal graphs.

The first approach that we consider for multi-relation graph summarization is a two-step method based on summarizing each relation in isolation, and then aggregating the resulting summaries in some clever way to produce a final unique summary. In doing this, as a side contribution, we provide the first polynomial-time approximation algorithm based on the  $k$ -Median clustering for the classic problem of lossless single-relation graph summarization.

Then, we demonstrate the shortcomings of these two-step methods, and propose holistic approaches, both approximate and heuristic algorithms, to compute a summary directly for multi-relation graphs. In particular, we prove that the approximation bound of  $k$ -Median clustering for the single relation solution can be maintained in a multi-relation graph with proper aggregation operation over adjacency matrices corresponding to its multiple relations. Experimental results and case studies (on co-authorship networks and brain networks) validate the effectiveness and efficiency of the proposed algorithms.

CCS Concepts: • **Information systems** → *Data mining*; **Network data models**.

Additional Key Words and Phrases: graph summarization, multi-relation graph, approximation, k-median

## ACM Reference Format:

Xiangyu Ke, Arijit Khan, and Francesco Bonchi. 2021. Multi-relation Graph Summarization. *ACM Trans. Knowl. Discov. Data.* 1, 1, Article 1 (January 2021), 29 pages. <https://doi.org/10.1145/3494561>

## 1 INTRODUCTION

Fueled by the unprecedented growth rate of knowledge graphs, social networks, and Internet-of-Things<sup>1</sup>, the problems of storing, managing, and mining very large graph data have received an enormous deal of attention in the data mining research community in recent years. At the current

<sup>1</sup>The Knowledge and Action Graph of Microsoft has 21 billion facts, 18 billion action links, and over five billion relationships between more than one billion people, places, and things [35]. Facebook has 800 million active users [49]. Graph-of-Things (GoT), which is a live knowledge graph system for Internet-of-Things, has been adding millions of records per hour, and roughly about 10 billion RDF triples per month [59].

Authors' addresses: Xiangyu Ke, [xiangyu001@ntu.edu.sg](mailto:xiangyu001@ntu.edu.sg), Nanyang Technological University, Singapore; Arijit Khan, [arijit.khan@ntu.edu.sg](mailto:arijit.khan@ntu.edu.sg), Nanyang Technological University, Singapore; Francesco Bonchi, [francesco.bonchi@isi.it](mailto:francesco.bonchi@isi.it), ISI Foundation, Italy and Eurecat, Spain.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1556-4681/2021/1-ART1 \$15.00  
<https://doi.org/10.1145/3494561>

rate of data volume increase, in fact, it is becoming highly impractical to store, manage, process, and visualize such big graphs. Graph summarization alleviates these issues by producing a concise graph representation (i.e., summary) that still can be meaningfully explored and queried. Graph summarization has shown to be beneficial in a wide range of applications, such as visualization, interactive and exploratory analysis, approximate query processing, reducing the on-disk storage footprint, graph embedding, and graph processing in modern hardware [7, 44, 48, 53, 83].

Surprisingly, little attention has been paid to the problem of summarizing multi-relation graphs. In real-world, entities are often correlated in multiple ways, either explicitly or implicitly. For instance, BioGRID (thebiogrid.org) describes seven different types of genetic interactions between genes in Homo Sapiens. STRING (string-db.org) models protein-to-protein interactions with six types of correlations statistically learned from existing protein databases, revealing that most protein interactions are associated with at least two types of correlations. Other applications where multiple relations exist between entities include social and financial networks [28], urban and transportation systems [14], ecology research [72], and recommender systems [36, 54, 81].

When multiple relations exist between entities, data is modeled as *multi-relation networks* (also known as *multi-layer*, *multiplex*, or *multi-dimensional networks*) [25]. This graph model has been attracting increasing research interest in graph analytics, such as in shortest path finding [85], core decomposition and densest subgraph discovery [28, 29], node clustering [8], frequent subgraphs mining [82], and in social networks analysis [23], just to mention a few.

In this paper, we study, for the first time, the problem of multi-relation graph summarization. Before discussing the contributions of our work and how they collocate in the state of the art, we need to provide some background notions and formally define the problems.

## 1.1 Background and related work

We consider graph summarization obtained by aggregating nodes into *supernodes*. In particular, we adopt *lossless* summarization as introduced by Navlakha et al. [57], for simple undirected single-relation graphs. Given a graph  $G = (V, E)$ , a lossless summary  $\mathcal{S} = \langle G_S, C_S \rangle$  consists of a *summary graph*  $G_S = (V_S, E_S)$  and a set of *edge corrections*  $C_S = \langle C_S^+, C_S^- \rangle$ , where:

- $V_S = \{S_1, \dots, S_k\}$  is the set of *supernodes* inducing a partition of  $V$ , i.e.,  $\bigcup_{i=1}^k S_i = V$  and  $\bigcap_{i=1}^k S_i = \emptyset$ ;
- $E_S \subseteq V_S \times V_S$  is a set of *superedges* between supernodes (possibly including self-loops);
- $C_S^+ \subseteq E$  is the set of edges to be inserted to reconstruct  $G$ , while  $C_S^- \subseteq (V \times V) \setminus E$  is the set of edges to be deleted.

The summarization is lossless because given the summary  $\mathcal{S}$ , we can reconstruct the original graph *exactly*, by (1) “exploding” each superedge  $(U, W) \in E_S$ , i.e., creating an edge  $(u, w)$  for each pair of nodes  $u, w \in U \times W$ , (2) adding each edge in  $C_S^+$ , and (3) removing the edges in  $C_S^-$ . More formally:  $E = \{(u, w) \mid u \in U, w \in W, (U, W) \in E_S\} \cup C_S^+ \setminus C_S^-$ . An example of lossless summary is provided in Figure 1.

The problem studied by Navlakha et al. [57] is to find the smallest possible summary of an input graph.

**PROBLEM 1 (LOSSLESS-SUM [57]).** *Given a simple, undirected, graph  $G = (V, E)$ , find its smallest lossless summary  $\mathcal{S} = \langle G_S, C_S \rangle$ , i.e., the one that minimizes  $|E_S| + |C_S|$ .*

The cost of storing the mapping from nodes in  $V$  to supernodes in  $V_S$  is disregarded in the objective function [57], since it remains constant across different summaries, i.e.,  $\mathcal{O}(|V|)$ . The problem can be seen through the lenses of the *Minimum Description Length* (MDL) principle [65], which states that the best theory to infer from a set of data is the one which minimizes the sum of




Fig. 1. A simple graph and its summary.

the size of the theory and the size of the data when encoded through the theory. Here, the data is the input graph  $G$ , the theory is the summary graph represented by supernodes  $V_S$  and superedges  $E_S$ , and the correction list  $C_S$  is the encoding of the data with regards to the theory. For instance, in the example in Figure 1 the cost of the summary is 2 (1 superedge + 1 correction), against a cost of 5 of the original graph (i.e., 5 edges).

As observed in [57], a summary is entirely defined by the partitioning of nodes into supernodes. In fact, once provided such partitioning, superedges can be deterministically decided by simply checking whether they induces advantages w.r.t. the summary cost or not. Finally, once decided the superedges, the corrections are straightforwardly identified. Therefore, Problem 1 is essentially a graph partitioning problem: given a single-relation graph  $G$  with  $m$  nodes and  $n$  edges, there are  $\sum_{k=1}^n \binom{n}{k} = \mathcal{O}(2^n)$  possibilities to partition these  $n$  nodes into supernodes  $V_S$ . To solve it, Navlakha et al. [57] proposed a simple greedy agglomerative heuristic *without quality guarantee*. They also studied a lossy version of the problem. To the best of our knowledge, the computational complexity of Problem 1 has not been addressed in the literature [57, 71]. We keep the problem of investigating the computational complexity of Problem 1 open here, and consider it as an interesting future work.

A similar approach is followed by LeFevre and Terzi [50] who study summaries obtained by aggregating nodes into supernodes. However, they keep on each superedge the information about how many original edges it represents (but not which ones): this is clearly a lossy summarization. Their objective is then to find the summary minimizing the loss for a given number  $k$  of allowed supernodes (which implicitly controls the compression rate). The loss is represented by the *reconstruction error*, i.e., the difference between the original graph and the probabilistic graph that one can reconstruct from the summary. Similar to [57], [50] propose a simple greedy agglomerative heuristic *with no quality guarantee*. Later Riondato et al. [63, 64] propose *the first polynomial-time approximation algorithm* for the problem of [50]. In this paper we show that, following a similar intuition, we can achieve *the first polynomial-time approximation algorithm* for the classic lossless summarization problem of [57].

## 1.2 Multi-relation graph summarization

In this work, we extend the notion of lossless summarization over multi-relation graphs. An undirected, multi-relation graph  $G$  is a triplet  $(V, E, R)$ , where  $V$  is a set of  $n$  nodes,  $R$  is a set of  $q$  relations, and  $E \subseteq V \times V \times R$  is a set of  $m$  undirected edges. Therefore, in a multi-relation graph, each edge is a triplet: e.g., an edge between nodes  $u$  and  $v$  in relation  $r \in R$  is represented by  $(u, v, r)$ . A summary  $\mathcal{S} = \langle G_S, C_S \rangle$  is defined as in the single-relation case, the only difference is that the correction edges in  $C_S$  are triplets and also superedges are now triplets, i.e.,  $E_S \subseteq V_S \times V_S \times R$ .

**EXAMPLE 1.** Figure 2 provides an example of a multi-relation graph and its summary. The graph on the left-hand side is defined over 3 relations, contains 5 nodes and 16 edges. The summary on the right-hand side is obtained by grouping  $\{a, c\}$  and  $\{b, d\}$  as two supernodes and keeping  $\{e\}$  as a




Fig. 2. A multi-relation graph and its summary.




Fig. 3. The individual summary for each relation in Figure 2(a).

supernode over all 3 relations. No correction is required here. The cost of such summary is thus 6 (given by 6 superedges + 0 correction), while the cost of the original graph was 16 (i.e., 16 edges).

The problem we study in this article is formally defined as follows.

**PROBLEM 2 (LOSSLESS-SUM-MULTI).** Given a multi-relation graph  $G = (V, E, R)$ , find its smallest lossless summary  $\mathcal{S} = \langle G_S, C_S \rangle$ , i.e., the one that minimizes  $|E_S| + |C_S|$ .

In many applications one might be interested in a summary with a predefined number  $k$  of supernodes. Moreover, one might solve Problem 2 by using an algorithm which takes in input the number of supernodes  $k$ , with a wrapper for selecting the optimal value of  $k$ . Therefore, in this article we also tackle the following problem.

**PROBLEM 3 ( $k$ -LOSSLESS-SUM-MULTI).** Given a multi-relation graph  $G = (V, E, R)$  and  $k \in \mathbb{N}$ , find the lossless summary  $\mathcal{S} = \langle G_S, C_S \rangle$ , such that  $G_S = (V_S, E_S, R)$  and  $|V_S| = k$ , that minimizes  $|E_S| + |C_S|$ .

### 1.3 Why not keeping an individual summary for each relation?

Figure 3 shows the optimal summary for each relation of the multi-relation graph in Figure 2(a). One could argue that storing these three individual summaries would also serve as a lossless summarization for the given multi-relation graph. However, this is not a good option due to two reasons: (1) The optimal multi-relation summary in Figure 2(b) provides us more insights about the input network. For example, by only looking at the individual summaries in Figure 3, we cannot easily determine the fact that the nodes in set  $\{a, c\}$  are fully connected with the nodes in set  $\{b, d\}$  via all relations. However, the nodes within each of these sets interact with themselves in different ways. Thus, it is better to characterize them as two strongly connected supernodes with different self-loops as in the multi-relation summary in Figure 2(b). (2) More storage is required for maintaining all individual summaries, because each individual summary might require storing of a different node mapping (i.e., the mapping from nodes in  $V$  to supernodes in  $V_S$ ). This is demonstrated in our experiments (§ 8.2).

## 1.4 Contributions and roadmap

The considerations above also suggest a natural two-step approach to multi-relation graph summarization: first summarize the input graph one relation at a time, then aggregate the various summary graphs into a single summary. We follow this intuition and, as a first solution, we develop such a two-step approach. For both steps we compare different methods, producing several baseline two-step algorithms.

Among the algorithms we design for the first step, we also consider a  $k$ -median algorithm to produce single-relation summary, inspired by [63]. For this algorithm we prove approximation guarantees: this is, to the best of our knowledge, the *first polynomial-time approximation algorithm for the classic lossless graph summarization problem* of [57].

Finally we show, by means of an example, an effectiveness limitation suffered by the two-step approach. Therefore, following the intuition behind the example, we move on to design *holistic* approaches, which are experimentally shown to be faster and more accurate than the two-step approaches. Moreover, our holistic  $k$ -Median<sup>+</sup> algorithm maintains the same approximation guarantee for multi-relation graph summarization (§ 5). Our final Hybrid algorithm combines the Greedy method [57] and the approximate solution  $k$ -Median, to provide the most compact summary in practice (§7).

Our main contribution is to initiate investigation into multi-relation graph summarization. Besides, this paper achieves the following contributions:

- We revise the classic single-relation graph lossless summarization problem, and provide the first polynomial-time approximation algorithm (§3).
- We design basic two-step algorithms, which first generate a lossless summary for each relation, then properly aggregate them to obtain one uniform summary. We also highlight the limits of this approach (§4).
- We propose holistic algorithms for more compact and efficient summary generation. Our holistic  $k$ -Median<sup>+</sup> algorithm maintains the same approximation guarantee for multi-relation graph summarization (§5).
- We combine the traditional Greedy method [57] and the approximation solution  $k$ -Median as the final proposed algorithm Hybrid, which empirically produces the most compact summary (§7).
- Our empirical evaluation on four real-world networks confirms that our holistic algorithms can produce more compact summaries and are faster than the two-step approaches (§8).
- Real-world applications on visualization, classification, and query processing demonstrate the effectiveness and efficiency of our proposal (§9).

Next section cover additional related literature. Section 10 concludes the paper and discusses future work.

## 2 OTHER RELATED WORK

Graph summarization has been used for a wider range of problems related to static plain graphs [9–11, 19, 20, 39, 47, 62], static attributed graphs [18, 68, 75, 84], dynamic and stream graphs [1–3, 22, 27, 32, 43, 67, 69, 73, 78, 86], probabilistic and distributed graph summarization [33, 52]. We refer the reader to excellent surveys and tutorials [7, 17, 44, 48, 51, 53, 76]. Regardless this wide literature, *no prior work has tackled graph summarization in the multi-relation graph setting*. Therefore, the rest of this section covers the literature about single-relation graph summarization. In Section 1.1, we already reviewed the most important related papers which constitute the background for our work. Table 1 collocates our contribution within the most important related work.

Table 1. Characterization of the most related papers.

Paper	Lossless	Multi-Relation	Approx. Guarantees
[57]	✓		
[50]			
[63]			✓
[5]			✓
[71]	✓		
this work	✓	✓	✓

**Aggregation-based graph summary.** Notable techniques under this category are pattern mining and community based summarization [13, 46, 66], supernode and edge-correction (thus lossless) [57, 71], supernode and reconstruction-error (thus lossy) [5, 50, 63]. Supernode based aggregation methods [5, 50, 57, 63, 71] are most similar to ours and are summarised in Table 1.

Very recently, SSumM [38] proposes lossy graph summarization to minimize the reconstruction error, however the constraint is on size of the summary graph in bits (and not on the number of supernodes as in [50, 63]). Notice that our focus in this work is lossless summarization over multi-relation graphs, which is different from [38]. We experimentally demonstrate the summary cost in bits based on various storage formats in § 8.

**Web graph and social networks compression.** Boldi and Vigna [10] show that web graphs are compressible down to almost two bits per edge. Chierichetti et al. [19] use shingle ordering instead of lexicographical ordering of web pages, in order to tackle social networks. Finding an order of nodes, which captures the “regularity” of the network, is a challenging problem. Boldi et al. [9] introduce a layered label propagation algorithm for reordering very large graphs. Other interesting works include [11, 20, 39, 47, 62]. These methods focus on reducing the number of bits needed to encode an edge, and none compute graph summaries.

**Attribute-based graph summary.** Nodes and edges of many real-world graphs are annotated with attributes. Hence, there exist graph summarization works considering both topology and semantics of the node and edge attributes [45]. FUSE [68] is a functional summarization technique for protein interaction networks, and this helps comprehending high-level functional relationships in disease-related PPI networks such as Alzheimer’s disease network. SNAP [75] and OLAP [18] allow interactive summarization at various resolutions over heterogeneous networks. Topology and attribute-based summarization of a large collection of small graphs (e.g., chemical compounds) and its application in constructing data-driven visual graph query interfaces are discussed in [84]. These methods are not directly comparable to ours, since our summarization deals with the graph structure. For example, a superedge in [75] exists between a pair of supernodes if any node in a supernode has *at least one* edge to the nodes in the other supernodes. However, in our problem formulation, we consider the exact number of edges between them.

**Application-oriented graph summary.** These are graph summarization techniques for efficient query answering and pattern mining, such as reachability, shortest path, and pattern matching queries [26, 77, 87], eigenvector centrality, degree, and adjacency queries [50], neighborhood query [55], keyword search [80], distributed graph computation [40], graph mining [16, 21, 49, 56], information cascade and influential node discovery [60, 61, 70]. We demonstrate applications of our multi-relation graph summary in efficient query processing in §9.

**Other related graph computation.** Related graph analytics problems include sampling [34], sparsification [6], clustering and community detection [4, 58, 79], graph embedding [37], partitioning

[41], and dense subgraph mining [31]. As discussed in [37, 44, 51], these problems are different from graph summarization.

### 3 SINGLE-RELATION GRAPH SUMMARIZATION: $K$ -MEDIAN CLUSTERING

In this section, we provide an approximation algorithm for single-relation graph summarization problem based on  $k$ -Median clustering. The  $k$ -Median clustering is performed on the rows of the adjacency matrix  $A_G$  of the input graph  $G$ , to create  $k$  supernodes. In particular, the goal of  $k$ -Median clustering is to find a set of  $k$  centers  $\mathbf{x} = \{x_1, x_2, \dots, x_k\}$  that minimizes the  $k$ -Median cost for the node set  $V \subseteq \mathbb{R}^n$ . The  $k$ -median cost is defined as  $\sum_{v \in V} d(v, \mathbf{x})$ , where  $d(v, \mathbf{x}) = \min_{x \in \mathbf{x}} d(v, x)$ . Here,  $d(v, x)$  denotes the Euclidean distance between two points (i.e., nodes)  $v, x \in \mathbb{R}^n$ . The nodes are then grouped into  $k$  supernodes based on their nearest cluster center. Notice that the  $k$ -summary is a graph summary with exactly  $k$  supernodes. After obtaining the supernode partitioning, we include superedges and correction list as discussed in § 1.1. The time complexity of  $k$ -Median is  $O(m + nk \log n)$  [63].

We prove that the  $k$ -Median algorithm guarantees  $16$ -approximation to the optimal solution for the Lossless-Sum problem with  $k$  supernodes. In the previous study [63], the  $k$ -Median clustering based technique was applied to generate approximated lossy summary. We bridge the gap between the reconstruction error (Equation 1) in lossy summary and the correction list size (Equation 3) in lossless summary, which is our problem. This enables reusing the same technique in a different problem setting (i.e., our problem), and it achieves a different approximation factor from that in [63]. Here, we only consider the number of correction edges  $|C_S|$ , since the number of superedges  $|E_S|$  can be bounded with  $O(k^2)$  for a  $k$ -summary. Our proof is built on top of a theoretical result by Riondato et al. [63], that establishes 8-approximation guarantee for a similar  $k$ -Median algorithm, with respect to the quality of a *lossy summarization*, known as the  $l_p$  reconstruction error as below. For a lossy summarization [5, 50, 63], only the graph summary  $G_S = (V_S, E_S)$  is created; no additional correction list is stored. The summary  $G_S$  is a complete graph in this case, including all self-loops, i.e.,  $E_S = V_S \times V_S$ . Given a summary, the graph is approximately reconstructed by the expected adjacency matrix,  $A_{G_S}^\uparrow$ , which is an  $(n \times n)$  matrix with  $A_{G_S}^\uparrow(u, w) = \frac{|E_{UW}|}{|U||W|}$ . Here,  $U, W$  are supernodes in  $V_S$  such that  $u \in U$  and  $w \in W$ , and  $E_{UW}$  is the set of edges that actually exist between  $U$  and  $W$  in the original graph  $G$ . The quality of the summary, called the  $l_p$  reconstruction error, is measured by a norm of difference between the input adjacency matrix  $A_G$  and the reconstructed adjacency matrix  $A_{G_S}^\uparrow$ .

The  $l_p$  reconstruction error ( $RE_p$ ) of a summary  $G_S$  for a graph  $G$  is:

$$RE_p(G, G_S) = \sqrt[p]{\sum_{u=1}^{|V|} \sum_{w=1}^{|V|} (|A_G(u, w) - A_{G_S}^\uparrow(u, w)|)^p} \quad (1)$$

In this paper, we use  $p = 1$ , that is, the  $l_1$  reconstruction error. From [63], we have the following theorem.

**THEOREM 1.** *Let  $G_{S^*}$  be the  $k$ -summary induced by the  $k$ -Median partitioning of the rows of  $A_G$ , and let  $G_{S^+}$  be the optimal  $k$ -summary for  $G$  with respect to the  $l_1$ -reconstruction error. The  $l_1$ -reconstruction error of  $G_{S^*}$  is an 8-approximation to the best  $l_1$ -reconstruction error. Formally:  $RE_1(G, G_{S^*}) \leq 8 \cdot RE_1(G, G_{S^+})$ .*

**LEMMA 1.** *Let  $G_{S^+}$  be the optimal  $k$ -summary for  $G$  with respect to the  $l_1$ -reconstruction error, and let  $G_{S^*}$  be the optimal  $k$ -summary for  $G$  with respect to the number of correction edges. The correction list size,  $|C_{S^*}|$  of  $G_{S^*}$  is a 2-approximation to the best size of correction list,  $|C_{S^+}|$ . Formally:  $|C_{S^*}| \leq 2 \cdot |C_{S^+}|$ .*

PROOF. Let us denote by  $\alpha_{UW} = \frac{|E_{UW}|}{|U||W|}$ . From the definition of  $l_1$  reconstruction error in Equation 1, we get:

$$\begin{aligned} RE_1(G, G_S) &= \sum_{u=1}^{|V|} \sum_{w=1}^{|V|} |A_G(u, w) - A_{G_S}^\uparrow(u, w)| \\ &= 2 \cdot \sum_{(U, W) \in V_S \times V_S} |U||W| \alpha_{UW} (1 - \alpha_{UW}) \end{aligned} \quad (2)$$

The intuition behind this derivation is that there are  $|U||W|$  cells corresponding to each supernode pair  $(U, W)$  in both the original adjacency matrix  $A_G$  and the reconstructed adjacency matrix  $A_{G_S}^\uparrow$ . In  $A_{G_S}^\uparrow$ , all such cells are filled with  $\alpha_{UW}$ . In  $A_G$ , there are  $\alpha_{UW}$  proportion of cells having value 1, and the rest  $(1 - \alpha_{UW})$  proportion of cells having value 0. The subtraction results of  $|A_G(u, w) - A_{G_S}^\uparrow(u, w)|$  for the first group are all  $(1 - \alpha_{UW})$ , and those for the second group are all  $\alpha_{UW}$ . Thus, we derive the second line in Equation 2.

Now in the context of lossless summary, we decide whether to keep a superedge between a pair of supernodes  $(U, W)$  by the edge density between them. If  $\alpha_{UW} > 0.5$ , maintaining a superedge can result in less storage overhead of  $C_S^-$ , than that of  $C_S^+$  without this superedge. Suppose  $C_S$  is the correction list, then its size can be calculated as below.

$$|C_S| = \frac{1}{2} \sum_{(U, W) \in V_S \times V_S} |U||W| \begin{cases} (1 - \alpha_{UW}), & \text{if } \alpha_{UW} > 0.5 \\ \alpha_{UW} & , \text{ otherwise} \end{cases} \quad (3)$$

From Equations 2 and 3, we have:

$$\frac{RE_1(G, G_S)}{|C_S|} = \begin{cases} 4\alpha_{UV} & , \text{if } \alpha_{UV} > 0.5 \\ 4(1 - \alpha_{UV}), & \text{otherwise} \end{cases} \quad (4)$$

Clearly,  $2 \leq \frac{RE_1(G, G_S)}{|C_S|} \leq 4$ , since  $\alpha_{UV} \in [0, 1]$ . In other words,

$$\frac{1}{4} RE_1(G, G_S) \leq |C_S| \leq \frac{1}{2} RE_1(G, G_S) \quad (5)$$

Suppose  $G_{S^+}$  be the optimal  $k$ -summary for  $G$  with respect to the  $l_1$ -reconstruction error, and let  $G_{S^*}$  be the optimal  $k$ -summary for  $G$  with respect to the number of correction edges. Thus, we get:

$$\begin{aligned} |C_{S^+}| &\leq \frac{1}{2} RE_1(G, G_{S^+}) && \triangleright \text{ by the r.h.s of Equation 5} \\ &\leq \frac{1}{2} RE_1(G, G_{S^*}) && \triangleright \text{ since } G_{S^+} \text{ is optimal wrt } RE_1 \\ &\leq 2 \cdot |C_{S^*}| && \triangleright \text{ by the l.h.s of Equation 5} \end{aligned} \quad (6)$$

This completes the proof.  $\square$

**THEOREM 2.** *Let  $G_{S^*}$  be the  $k$ -summary induced by the  $k$ -Median partitioning of the rows of  $A_G$ , and let  $G_{S^+}$  be the optimal  $k$ -summary for  $G$  with respect to the number of correction edges. The correction list size,  $|C_{S^*}|$  of  $G_{S^*}$  is a 16-approximation to the best size of correction list,  $|C_{S^+}|$ . Formally:  $|C_{S^*}| \leq 16 \cdot |C_{S^+}|$*



PROOF. We denote by  $G_{S^+}$  the optimal  $k$ -summary for  $G$  with respect to the  $l_1$ -reconstruction error. Next, we derive the follows.

$$\begin{aligned}
|C_{S^\#}| &\leq \frac{1}{2} RE_1(G, G_{S^\#}) && \triangleright \text{ by the r.h.s of Equation 5} \\
&\leq 4 \cdot RE_1(G, G_{S^+}) && \triangleright \text{ by Theorem 1} \\
&\leq 4 \cdot RE_1(G, G_{S^*}) && \triangleright \text{ since } G_{S^+} \text{ is optimal wrt } RE_1 \\
&\leq 16 \cdot |C_{S^*}| && \triangleright \text{ by the l.h.s of Equation 5}
\end{aligned} \tag{7}$$

Hence, the proof is completed.  $\square$

Note that in §6, we discuss several empirical methods for finding a suitable  $k$  for the  $k$ -Median method, which helps it adapt to the general Lossless-Sum problem.

#### 4 MULTI-RELATION GRAPH SUMMARY: BASELINE METHODS

In this section, we first present several straightforward baselines for the lossless summarization of multi-relational graphs, then demonstrate how they suffer from effectiveness issues, which will be instrumental in developing a more accurate and scalable solution in §5.

Our baseline algorithms follow a two-step approach:

(1) We explore the input graph for one relation at a time, and generate a lossless summary for each of them.

(2) The summaries across relations are properly aggregated to obtain one uniform summary.

In the first step, our problem is same as the Lossless-Sum problem. Given a set of summaries  $\{S_1, S_2, \dots, S_q\}$ , each for a specific relation, our next target is to find a single summary, i.e., a partition into supernodes, that agrees as much as possible with the  $q$  individual summaries.

In addition to the  $k$ -Median approach introduced in §3, we briefly revisit some widely-used graph summarization techniques for the Lossless-Sum problem. To the best of our knowledge, all existing algorithms [57, 71] for the Lossless-Sum problem are heuristic in nature, without any theoretical guarantee on the summary size. In §4.2, we tackle the problem of summary aggregation and provide several methods to aggregate the individual summaries across relations. Finally, in §4.3 we discuss potential limitations of these two-step baseline algorithms.

##### 4.1 Single-relation graph summarization algorithms

We first revisit Greedy and Randomized algorithms from [57]. Then, we discuss an advanced algorithm, SWeG, with similar idea in recent literature [71]. An example is provided to demonstrate these algorithms.

**Greedy algorithm.** The Greedy algorithm [57] is a heuristic, bottom-up approach. It first considers every node in the input graph as a supernode, and iteratively merges the best pair  $\{u, w\}$  with the maximum reduction in summary cost. The general workflow of Greedy algorithm is given below: (1) It computes the potential cost reduction for all pairs of nodes in the input graph  $G$  which are 2-hops apart, and records those pairs which are positive. (2) The best pair of nodes  $\{U, W\}$  with highest cost reduction is merged into a new node  $W$ . (3) Delete cost reduction records about  $U$  or  $W$  for all the nodes which is within 2-hops to  $U$  or  $W$ , and compute their cost reduction to the new supernode  $H$ . (4) Update the cost reduction between nodes which are neighbors of  $H$  ( $I$  is a neighbor of  $H$  if there exists any edge  $\{a, b\} \in E$ ,  $a \in H$ ,  $b \in I$ ). (5) Repeat (2)-(4) until no positive cost reduction exists.

Notice that Greedy directly works with the Lossless-Sum problem. It can easily handle the additional input  $k$  for the number of supernodes by (a) terminating earlier when the number of




Fig. 4. Example for single-relation graph summary methods.

supernodes becomes  $k$  (even before satisfying condition (5)), or **(b)** force to merge the pairs with less “sacrifice” in summary quality (i.e., negative cost reduction) after (5), if a smaller  $k$  is required.

Let  $d_{av}$  be the average degree for each node, the time complexity of Greedy is  $O(d_{av}^3 (d_{av} + \log n + \log d_{av}))$  [57].

**Randomized algorithm.** Comparing with the basic Greedy algorithm, the Randomized approach reduces the high computation overhead, by sacrificing the compression quality. It, in fact, has the worst performance in compression based on our experimental results in § 8. In each step, it randomly selects an unexplored supernode  $U$ , and computes cost reduction with all its unexplored neighbors. If no positive reduction exists,  $U$  is marked explored, and the algorithm continues. Otherwise,  $U$  is merged with its best neighbor  $W$  (i.e., having the highest reduction) into a new node  $H$ .  $U$  and  $W$  are removed, and  $H$  is now unexplored. The algorithm stops when all nodes are explored. The complexity of the Randomized algorithm is  $O(d_{av}^3)$  [57].

**SWeG.** The most recent algorithm, SWeG [71], is an advanced version of the Randomized algorithm with further efficiency improvement. It first divides the graph into smaller disjoint groups. Each node group contains supernodes with similar connectivity, based on the concept of the shingle of a node  $u$ , which is defined as:  $f(u) = \min_{w \in N_u \text{ or } w=u} h(w)$ .  $h$  is a bijective hash function  $h : V \rightarrow \{1, \dots, |V|\}$ , and  $N_v$  is the set of neighbors of node  $v$  in the input graph  $G$ . Two nodes have the same shingle with probability equal to the Jaccard similarity of their neighbor sets [12]. The shingle of a supernode  $U$  is extended to be:  $F(U) = \min_{u \in U} \{f(u)\}$ . Two supernode  $U \neq W \in V_S$  are more likely to have the same shingle if the nodes in  $U$  and those in  $W$  share similar connectivity. After the nodes are partitioned into groups, a similar procedure as Randomized is conducted within each group to merge supernodes. Moreover, this allows distributed implementation. The whole pipeline (grouping nodes, and merging nodes within each group) repeats  $T$  times, each time with a different, randomly generated hash function  $h$ .  $h$  can be easily produced by shuffling  $\{1, \dots, |V|\}$ .

The dividing step takes  $O(|E|)$  running time. The cost of the merging step is same as that of the Randomized algorithm:  $O(d_{av}^3)$ , but here  $d_{av}$  may be smaller since it operates on a smaller graph. If we only allow sequential implementation and repeat  $T$  times, it becomes  $O(pT d_{av}^3)$ , where  $p$  is the number of disjoint groups.

**EXAMPLE 2.** Figure 4 provides an example for single-relation graph summarization algorithms. Given the example graph on the left, Greedy produces the summary  $S_{GD}$ . At the beginning, the cost of each node equals to the number of edges incident to this node. Merging node  $a$  with  $c$ , or merging node  $b$  with  $d$  both lead to 50% cost reduction, regardless of order. Then we achieve an intermediate result, which is the same as  $S_{KM}$ . The cost of the superedge between supernodes  $ac$  and  $bd$  are counted twice from both side, which encourages Greedy to continue merging. This demonstrates that Greedy tends to result in bigger-size supernodes. Though, this will not cause any issue in this simple example, we later demonstrate in §4.3 that this may create trouble in aggregated summary finding (i.e., the second

step of baseline approaches).  $S_{KM}$  is the summary for  $k$ -Median when  $k = 2$ , which can be easily obtained with the provided adjacency matrix. If  $k$  is set to be 1,  $k$ -Median will return same result as  $S_{GD}$ . The result of the Randomized algorithm depends on the random order of node selection. In this example, however, the final result of Randomized would also be  $S_{GD}$ . The costs of these two summaries are same:  $S_{GD}$  has 1 superedge and 1 correction edge, while  $S_{KM}$  has 2 superedges.

## 4.2 Summary aggregation

The second phase of our baseline approach is aggregating summaries across individual relations to obtain one uniform summary for the entire multi-relation graph. We formally define this problem following the well-known clustering aggregation problem [30]. Both summary and clustering techniques partition the nodes (objects) into groups, and the aggregation operation further aims at finding a new uniform partition with minimum total disagreements to the current partitions.

A summary  $\mathcal{S}$  corresponds to a partition of  $n$  nodes. The partition size  $k$  can either be given as an input (§ 3), or can be automatically decided by the algorithm (§ 4.1). For any pair of nodes  $(u, v)$ , the indicator function  $I_{\mathcal{S}}(u, v)$  returns 1 if and only if  $u$  and  $v$  are located in the same supernode under summary  $\mathcal{S}$  (0 otherwise). Then, we have  $D_{u,v}(\mathcal{S}_i, \mathcal{S}_j)$  ( $i, j \in [1, q]$ ) to characterize if two summaries  $\mathcal{S}_i$  and  $\mathcal{S}_j$  disagree to each other on the partitioning of  $u$  and  $v$ . Formally,

$$D_{u,v}(\mathcal{S}_i, \mathcal{S}_j) = \begin{cases} 1, & \text{if } I_{\mathcal{S}_i}(u, v) \neq I_{\mathcal{S}_j}(u, v) \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

The total disagreement between summaries  $\mathcal{S}_i$  and  $\mathcal{S}_j$  is:

$$D_V(\mathcal{S}_i, \mathcal{S}_j) = \sum_{(u,v) \in V \times V} D_{u,v}(\mathcal{S}_i, \mathcal{S}_j) \quad (9)$$

In general, this metric counts the number of node pairs on which the two summaries disagree to each other. Therefore, the *summary aggregation* problem is given as follows.

**PROBLEM 4 (SUM-AGG).** *Given a set of  $q$  summaries (i.e., node partitions)  $\{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_q\}$  on a set of nodes  $V$ , compute a new summary  $\mathcal{S}$  that minimizes the total disagreements with all the given summaries, that is, it minimizes  $\sum_{i=1}^q D_V(\mathcal{S}, \mathcal{S}_i)$ .*

We now show that the Sum-Agg problem can be reduced to the well-studied correlation clustering problem [4], thus standard procedures for solving correlation clustering can be employed. For any pair of nodes  $(u, v) \in V \times V$ , we define the distance between them as  $\mathcal{D}(u, v) = \frac{1}{q} \cdot |\{i : 1 \leq i \leq q \text{ and } I_{\mathcal{S}_i}(u, v) = 0\}|$ , which means the fraction of summaries that assign the pair  $(u, v)$  into different supernodes. The correlation clustering objective is to find a node partitioning  $\mathcal{P}$  that minimizes  $\mathcal{D}_{\mathcal{P}} = \sum_{(u,v), I_{\mathcal{P}}(u,v)=1} \mathcal{D}(u, v) + \sum_{(u,v), I_{\mathcal{P}}(u,v)=0} (1 - \mathcal{D}(u, v))$ . If the solution  $\mathcal{P}$  places  $u, v$  in the same group, it will disagree with  $q \cdot \mathcal{D}(u, v)$  of the original partitionings due to individual relations. In contrast, it will disagree with  $q(1 - \mathcal{D}(u, v))$  remaining partitionings if the solution keeps  $u, v$  separate. Thus, for any partitioning  $\mathcal{P}$ , we have  $q \cdot \mathcal{D}_{\mathcal{P}} = \sum_{i=1}^q D_V(\mathcal{P}, \mathcal{S}_i)$ , which is our objective of Sum-Agg (i.e., Problem 4). Due to this reduction, the following algorithms for solving correlation clustering can be employed to solve our problem. Notice that some of them also come with provable approximation guarantees.

**The BEST algorithm.** The most simple algorithm, BEST, is to find one of the input summaries,  $\mathcal{S}_i$ , that minimizes the total number of disagreements to others. It can be computed in time  $O(q^2n)$ , where  $q$  is the number of input summaries (also the number of relations), and  $n$  the number of nodes in the input graph. Though simple, it yields a solution with an approximation ratio at most  $2 \cdot (1 - 1/q)$  [30].




Fig. 5. Empirical comparison among five summary aggregation algorithms over *DBLP\_10*.

**The Balls algorithm.** The Balls algorithm [15] first sorts the nodes in an increasing order of the total distance to all other nodes. Recall that for any pair of nodes  $(u, v) \in V \times V$ , we define the distance between them as  $\mathcal{D}(u, v) = \frac{1}{q} \cdot |\{i : 1 \leq i \leq q \text{ and } I_{S_i}(u, v) = 0\}|$ , that is, the fraction of summaries that assign the pair  $(u, v)$  into different supernodes. The algorithm is defined with an input parameter  $\alpha$ . The intuition of the algorithm is to find a set of nodes that are close to each other, and far from other nodes. In order to find a good cluster, we take all nodes that are close (within a ball) to a node  $u$ . The triangle inequality guarantees that if two nodes are close to  $u$ , then they are also relatively close to each other. Once such a cluster is found, we remove it from the graph, and proceed with the rest of the nodes. The Balls algorithm consumes  $\mathcal{O}(n^2)$  running time, while ensuring  $\max\{\frac{1-\alpha}{\alpha}, \frac{1+2\alpha}{1-2\alpha}, \frac{2-2\alpha}{1-2\alpha}\}$  approximation guarantee [30]. The additional time cost for pre-computing distances between all pairs of nodes is  $\mathcal{O}(qn^2)$ .

**The Agglomerative algorithm.** It follows the standard bottom-up procedure. The distance  $\mathcal{D}(u, v)$  between nodes  $u, v$  is the same as in the Balls algorithm. The agglomerative algorithm first creates a singleton group for each node. Then, in each step, it picks the pair of groups  $A$  and  $B$  with the smallest average distance, which is defined as the average distance of all pair of nodes  $(a, b)$ , such that  $a \in A$  and  $b \in B$ . If it is less than  $1/2$ , the two selected groups are merged into a single one. The algorithm terminates when no merging is possible. The time complexity of this algorithm is  $\mathcal{O}(n^2 \log n)$ .

**The Furthest algorithm.** The Furthest algorithm is a top-down method. At the beginning, all nodes are placed in one group. Then, the pair of nodes which are furthest apart are found, and kept in two different groups. All remaining nodes are assigned to the center that incurs the least cost. In the following steps, each time a new center is found to be furthest from the current centers, and the node assignment changes accordingly. The procedure continues until the new solution induces a worse cost. Suppose that at the end we have  $k$  centers, then total running time will be  $\mathcal{O}(k^2n)$ .

**The LocalSearch algorithm.** The LocalSearch algorithm starts with some partition of nodes. It can be a randomly generated one, or the result of any aforementioned algorithm. The algorithm goes through each node, and decides to keep it still, move it to other group, or make it a singleton, by comparing various costs. An efficient way to compute the cost of assigning a node  $v$  to the cluster  $C_i$  is as follows:

$$\text{cost}(v, C_i) = M(v, C_i) + \sum_{j \neq i} (|C_j| - M(v, C_j)) \quad (10)$$




Fig. 6. (a) The graph with distances between nodes for summary aggregation computation. We omit node  $E$  here since its distance to any other node is 1. (b) The aggregated summary of the individual summaries. All the summary aggregation algorithms produce the same result in this example. (c) The potentially best summary  $S_2$ .

Here,  $M(v, C_i) = \sum_{u \in C_i} \mathcal{D}(u, v)$ . The cost of assigning a node  $v$  as a singleton is  $\sum_j (|C_j| - M(v, C_j))$ . The running time of LocalSearch is  $O(Tn^2)$ , where  $T$  is the number of iterations before no better move can be found.

**Empirical comparison.** Figure 5 presents the experimental comparison of the five aforementioned summary aggregation algorithms. *Relative size* [57, 71] is defined as  $\frac{|E_S| + |C|}{|E|}$ . The numerator is our objective function (Problem 2), i.e., cost of the summary, and the denominator is constant for a given graph, i.e., graph size. Smaller relative size indicates better compactness. All individual summaries are generated by the Greedy algorithm. The *compactness* is evaluated with the relative size of the summary to that of the original graph. For more details on the dataset and experimental setup, we refer to §8. We observe that the compactness of Agglomerative, Furthest, and LocalSearch are comparable, and they all outperform BEST and Ball algorithms. Considering also the running time, the Furthest algorithm is slightly better than others, thus it is selected as the default summary aggregation method in our following experimental section.

### 4.3 Limitation of baselines methods

Consider again the running example multi-relation graph in Figure 2(a) and the summaries built by Greedy for each relation in isolation (Figure 3).

Figure 6(a) reports the distance values for the summary aggregation step. In this setting, all the summary aggregation algorithms would return summary  $S_1$ , which has 3 superedges and 6 correction edges. However, there exists another summary  $S_2$ , which can represent the input graph with 6 superedges and without any correction edge. This example demonstrates that the two-step baselines may return a lower-quality solution.

Minimizing disagreements between the final summary and the individual summaries (obtained via summarizing on individual relations) does not directly optimize our ultimate objective in Lossless-Sum-Multi and  $k$ -Lossless-Sum-Multi problems. Recall that the summary aggregation operation tends to minimize disagreements between the final summary and the input set of summaries. Therefore, if the input has a majority population of low-quality summaries, or inherently meaningless summaries for certain relations, the aggregation will be forced to agree to those useless summaries.

Notice that in Figure 3, each of the individually optimal summaries  $S_{r_1}$  and  $S_{r_2}$  has two supernodes, as opposed to three supernodes in the individually optimal summary  $S_{r_3}$ . This is because Greedy tends to produce larger-size supernodes. Next, when we aggregate  $S_{r_1}$ ,  $S_{r_2}$ , and  $S_{r_3}$  as in Figure 6(a)-(b), the resultant summary  $S_1$  also has two supernodes, exactly same as the two supernodes in  $S_{r_1}$  and

$S_{r_2}$ . Unfortunately,  $S_2$  that has three supernodes (same as in  $S_{r_3}$ ) is the optimal summary (Figure 6(c)). This demonstrates that: **(1)** if the input has a majority population of low-quality summaries, the aggregation will be forced to agree to those useless summaries. **(2)** Individual summaries having larger-size supernodes, as produced by Greedy, may cause trouble in the aggregation procedure.

## 5 MULTI-RELATION GRAPH SUMMARY: HOLISTIC METHODS

We next present holistic algorithms that, unlike the two-step baseline approaches, summarize the graph in a lossless manner and considering all relations at once. In particular, we shall discuss holistic versions of  $k$ -Median (§5.1), Greedy (§5.2), and Randomized (§5.3) algorithms, referred to as  $k$ -Median<sup>+</sup>, Greedy<sup>+</sup>, and Randomized<sup>+</sup>, respectively.

### 5.1 $k$ -Median<sup>+</sup>

For a multi-relation graph, an adjacency matrix exists for each relation. We explore several operations for matrix aggregation, e.g., Sum, Or, and Concatenate, and formally prove that the Concatenate operation maintains the properties for  $k$ -Median based technique to return an approximated solution.

Our intuitive idea is whether it is possible to *aggregate the adjacency matrices due to different relations into one aggregated matrix, and then cluster the rows of this aggregated matrix*. Potential aggregation operations may include Sum, Or, and Concatenate. Let  $\{A_{G_1}, A_{G_2}, \dots, A_{G_q}\}$  be a set of  $q$  ( $n \times n$ ) matrices, the Sum operation produces an aggregated matrix by  $A_G(i, j) = \sum_{1 \leq x \leq r} A_{G_x}(i, j)$ , the Or operation produces an aggregated matrix by  $A_G(i, j) = \vee_{1 \leq x \leq r} A_{G_x}(i, j)$ , and the Concatenate operation produces an aggregated matrix  $A_G = (A_{G_1} | A_{G_2} | \dots | A_{G_q})$  by concatenating the rows. Among them, we show below that concatenation permits  $k$ -Median<sup>+</sup> in achieving 16-approximation guarantee to the optimal summary size.

**THEOREM 3.** *Let  $A_G$  be the concatenated matrix of the adjacency matrices for individual relations, i.e.,  $A_G = (A_{G_1} | A_{G_2} | \dots | A_{G_q})$ , where  $A_{G_i}$  is the adjacency matrix for relation  $i$ ,  $i \in (1, q)$ . Let  $G_{S^*}$  be the  $k$ -summary induced by the  $k$ -Median partitioning of the rows of  $A_G$ , and let  $G_{S^*}$  be the optimal  $k$ -summary for  $G$  with respect to the number of correction edges. The correction list size,  $|C_{S^*}|$  of  $G_{S^*}$  is a 16-approximation to the best size of correction list,  $|C_{S^*}|$ . Formally,*

$$|C_{S^*}| \leq 16 \cdot |C_{S^*}| \quad (11)$$

To prove Theorem 3, we need to justify that Lemma 1 and Theorem 1 (mentioned in § 3 in the context of  $k$ -Median algorithm over single-relation graphs) still hold with our concatenated matrix.

First, we verify the correctness of Lemma 1 in this case. To extend our objective of counting the number of correction edges from single-relation to multi-relation, we only need to include an additional sum operation to Equation 3.

$$|C| = \frac{1}{2} \sum_{r \in R} \sum_{U, W \in V_S \times V_S} |U||W| \begin{cases} \alpha_{UW,r} & , \text{ if } \alpha_{UW,r} \leq 0.5 \\ (1 - \alpha_{UW,r}) & , \text{ otherwise} \end{cases} \quad (12)$$

Since we have a holistic summary corresponding to same set of supernodes across relations, the edge density between supernodes  $U$  and  $W$  for each relation  $r$ , denoted by  $\alpha_{UW,r}$ , can be calculated in the same way. Hence, the  $l_1$  reconstruction error (Equation 2) can be modified as:

$$RE_1(G, G_S) = 2 \sum_{r \in R} \sum_{(U, W) \in V_S \times V_S} |U||W| \alpha_{UW,r} (1 - \alpha_{UW,r}) \quad (13)$$

Based on the above two equations, one can prove an equivalent lemma of our earlier Lemma 1, for the multi-relation case. That is,  $|C_{S^+}| \leq 2 \cdot |C_{S^*}|$ . Here,  $G_{S^+}$  is the optimal  $k$ -summary for  $G$  with

respect to the  $l_1$ -reconstruction error, and  $|C_{S^+}|$  is the correction list size for  $G_{S^+}$ . Clearly, Sum and Or operations violate Equation 12.

Next, we rewrite the  $l_1$  reconstruction error in its original form (equivalent of Equation 1).

$$\begin{aligned} RE_1(G, G_S) &= \sum_{r \in R} \sum_{u=1}^{|V|} \sum_{w=1}^{|V|} \left| A_{G_r}(u, w) - A_{G_{S^+}, r}^\uparrow(u, w) \right| \\ &= \sum_{r \in R} \left\| A_{G_r} - A_{G_{S^+}, r}^\uparrow \right\| \end{aligned} \quad (14)$$

where  $A_{G_{S^+}, r}^\uparrow(u, w)$  is the edge density between the supernodes  $U$  and  $W$  for relation  $r$ , such that  $u \in U, w \in W$ .

To prove Theorem 1, Riondato et al. [63] defined an orthogonal smoothing projection  $P$  for a partitioning  $\mathcal{P}$  of  $n$  nodes (i.e., rows) in the adjacency matrix  $A_G$ . Since in our holistic summary, all relations would share the same partitioning  $\mathcal{P}$ , the same projection  $P$  can be applied to each adjacency matrix  $A_{G_r}$ . The  $l_1$ -reconstruction matrix for each relation can be computed by  $A_{G_{S^+}, r}^\uparrow(u, w) = PA_{G_r}P$ . By definition,  $A_{G_r}P$  is the  $k$ -Means matrix. And we define the  $k$ -Median matrix of relation  $r$  as below:

$$B_{G_r}(u, w) = \text{median}(\{A(x, y) \mid \{x, y\} \in \Pi_{UW}\}) \quad (15)$$

$\Pi_{UW}$  is the set of all possible pairs  $\{x, y\}$ , such that  $x \in U$  and  $y \in W$ . Lemma 1 and Lemma 4 in [63] provide the inequalities bridging the  $l_1$ -reconstruction error  $\|A - PAP\|$ ,  $k$ -means cost  $\|A - AP\|$  and the  $k$ -median cost  $\|A - B\|$ . Thus, we prove Theorem 1 for multi-relation case as follows.

$$\begin{aligned} RE_1(G, G_{S^\#}) &= \sum_{r \in R} \|A_{G_r} - A_{G_{S^\#}, r}^\uparrow\| \quad \triangleright \text{by Equation 14} \\ &= \sum_{r \in R} \|A_{G_r} - P_{G_{S^\#}} A_{G_r} P_{G_{S^\#}}\| \\ &\leq 2 \cdot \sum_{r \in R} \|A_{G_r} - A_{G_r} P_{G_{S^\#}}\| \quad \triangleright \text{Lemma 4, [63]} \\ &\leq 4 \cdot \sum_{r \in R} \|A_{G_r} - B_{G_{S^\#}}\| \quad \triangleright \text{Lemma 1, [63]} \\ &\leq 4 \cdot \sum_{r \in R} \|A_{G_r} - B_{G_{S^+}}\| \quad \triangleright G_{S^\#} \text{ is best for } k\text{-Median} \\ &\leq 4 \cdot \sum_{r \in R} \|A_{G_r} - A_{G_r} P_{G_{S^+}}\| \quad \triangleright \text{Lemma 1, [63]} \\ &\leq 8 \cdot \sum_{r \in R} \|A_{G_r} - P_{G_{S^+}} A_{G_r} P_{G_{S^+}}\| \quad \triangleright \text{Lemma 4, [63]} \\ &= 8 \cdot RE_1(G, G_{S^+}) \end{aligned} \quad (16)$$

Since both Theorem 1 and Lemma 1 hold for the multi-relation case, one can prove the correctness of 16-approximation result due to  $k$ -Median on the concatenated adjacency matrix. Therefore, Theorem 3 follows.

## 5.2 Greedy<sup>+</sup>

The Greedy algorithm can be generalized to a holistic algorithm, Greedy<sup>+</sup> for the Lossless-Sum-Multi problem without changing the workflow. Between any pair of supernodes  $U$  and  $W$ , let  $\Pi_{UW}$  be the set of all possible pairs  $\{a, b\}$ , such that  $a \in U$  and  $b \in W$ .  $E_{UW} \subseteq \Pi_{UW}$  is defined as the set

of edges actually present in the input graph  $G$ , i.e.,  $E_{UW} = \Pi_{UW} \cap E$ . Obviously,  $\Pi_{UW}$  is the same no matter for which relation. Next, we define  $A_{UW,r} \subseteq \Pi_{UW}$  to be the set of edges actually present in the original graph  $G$  for relation  $r$ , i.e.,  $A_{UW,r} = \Pi_{UW} \cap E_r$ . Similarly, the cost of a supernode pair  $(U, W)$  for relation  $r$  can be calculated as:

$$C(U, W, r) = \min\{|\Pi_{UW}| - |A_{UW,r}| + 1, |A_{UW,r}|\} \quad (17)$$

The neighbor set  $N_r(U)$  of  $U$  in relation  $r$  is defined to be the set of supernodes  $W$  that have such edge  $\{a, b\} \in A_{UW,r}$ . Moreover, the cost of maintaining a supernode  $U$  would add up across relations as follows:

$$C(U) = \sum_{r \in R} \sum_{X \in N_r(U)} C(U, X, r) \quad (18)$$

The cost reduction due to merging supernodes  $U$  and  $W$  into a new supernode  $H$  can be calculated as:

$$\Delta C(U, W) = \frac{C(U) + C(W) - C(H)}{C(U) + C(W)} \quad (19)$$

Taking fraction instead of the absolute cost reduction in the above equation is to get rid of the bias towards nodes with higher degree.

**EXAMPLE 3.** We use the example graph in Figure 2(a) for the demonstration of Greedy<sup>+</sup>. At the beginning, the cost of each node equals to the number of edges incident to this node, for all relations. Thus, node  $a$  and  $c$  will be selected for merging in the first round, resulting in  $\frac{7+7-(3+3+1)}{7+7} = 0.5$  cost reduction. Similarly, the nodes  $b$  and  $d$  will be merged in the second round. The cost again reduces by  $\frac{4+4-(3+1)}{4+4} = 0.5$ . Then, it can be easily verified that no further merging can result in positive cost reduction. Thus, it returns the optimal summary,  $S_2$  in Figure 6(c), for this example.

### 5.3 Randomized<sup>+</sup>

Randomized<sup>+</sup> follows the same modification in cost computation as the holistic Greedy<sup>+</sup>. The computation step does not change when comparing with the original Randomized algorithm. For SWeG algorithm, it is non-trivial to extend the node set dividing step to multi-relation case. We, therefore, leave it as a future direction.

## 6 FINDING OPTIMAL NUMBER OF SUPERNODES

Our  $k$ -Median-based approaches (i.e.,  $k$ -Median two-step baseline and holistic  $k$ -Median<sup>+</sup>) require  $k$ , a predefined number of supernodes, as an input. In practice, the user may not explicitly provide the target number of supernodes. In this section, we suggest to use the number of supernodes returned by Greedy<sup>+</sup> algorithm as the optimal value of  $k$  for  $k$ -Median-based approaches, and verify its good performance via the Elbow method [74] as below.

**The Elbow method.** The elbow method empirically verifies the cost of the  $k$ -median clustering when varying  $k$ . Within each cluster, it computes the median distance between each node to its center, and takes the sum of all median distances for all nodes, which is known as the *within-cluster sum-of-squares* error (WSS). The WSS is plotted against the cluster number  $k$ , we select the  $k$  for which WSS first starts to diminish. In the plot of WSS-versus- $k$ , this is visible as an “elbow”. Using the elbow as a cutoff point is a common heuristic<sup>2</sup> in optimization to choose a point where diminishing returns are no longer worth the additional cost. In clustering, this implies that

<sup>2</sup>Theoretically, *within-cluster sum-of-squares* error (WSS) monotonically decreases with larger cluster number  $k$ , thus the optimal  $k$  is trivially equal to the total number of nodes in the input graph. In practice, we would like to find a clustering where the WSS no longer decreases sharply.






Fig. 7. Compactness analysis with #supernodes  $k$ .

one should choose a number of clusters so that adding another cluster will not give much better modeling of the data.

First, we employ Greedy to suggest a preliminary  $k'$ . Then, we vary the  $k$  in a range with  $k'$  as the center, e.g.,  $[k' - 2000, k' + 2000]$ , and apply the Elbow method to verify the performance of  $k'$ . In practice, we can replace the clustering cost with our exact summary cost and select the optimal  $k$  via Equation 20, where  $S(k)$  denotes the  $k$ -Median summary having  $k$  supernodes. The exact summary cost is measured via *relative size* [57, 71], which is defined as  $\frac{|E_S| + |C|}{|E|}$ . The numerator is our objective function (Problem 2), i.e., cost of the summary, and the denominator is constant for a given graph, i.e., graph size. Smaller relative size means better compactness. We show the relative size on  $y$ -axis and  $k$  on  $x$ -axis, the curve becomes “valley” shape instead of “elbow” shape (Figure 7). The vertically dashed line denotes the number of supernodes found with Greedy+, which are always located in the “valley bottom” in all our datasets. This shows that our Greedy+ algorithm can suggest a good  $k$  for our  $k$ -Median+ approach.

$$\operatorname{argmin}_{k \in [k' - 2000, k' + 2000]} \frac{|E_{S(k)}| + |C_{S(k)}|}{|E|} \quad (20)$$

## 7 MULTI-RELATION GRAPH SUMMARY: HYBRID ALGORITHM

In this section, we discuss the shortcomings of the Greedy+ and the  $k$ -Median+ holistic algorithms, and provide a hybrid algorithm, Hybrid, based on properly combining them, to produce the most effective summary as shown empirically in the next section.

**The shortcomings of Greedy+.** The Greedy+ algorithm is a bottom-up iterative approach. The subsequent rounds of Greedy+ highly rely on the results of previous rounds. If the previous few rounds get trapped in some bad results, the follow-up rounds have no way to fix it.

**The shortcomings of  $k$ -Median+.** First, the  $k$ -Median+ algorithm requires an input number of supernodes  $k$ . Thus, it cannot directly solve the Lossless-Sum-Multi problem. Second, our experimental results in § 8.2 present that the compactness of the  $k$ -Median+ summaries are worse than those of Greedy+ summaries in practice.

Based on previous discussion, we can find that the Greedy+ algorithm and the  $k$ -Median+ algorithm are complementary: (1) Greedy+ can suggest a potentially good  $k'$  for  $k$ -Median+; (2)  $k$ -Median+ is able to directly generate a summary for any  $k$  around  $k'$ ; (3) For any  $k$ -Median+ summary, Greedy+ can further improve its quality, if possible.

**Hybrid algorithm.** Based on aforementioned properties, we properly combine the Greedy+ algorithm and the  $k$ -Median+ algorithm as our ultimately proposed Hybrid algorithm, as below.

(1) As discussed in § 6, we determine the optimal  $k$  for  $k$ -Median+, with the help of the Greedy+ method. (2) We generate a summary by  $k$ -Median+ algorithm with the best  $k$  found by Greedy+ (3) We conduct the Greedy+ algorithm again starting from the summary generated in (2), to further improve its compactness, if possible.

Table 2. Properties of datasets.

Dataset	#Nodes	#Edges	#Relations	Domain
<i>Homo</i>	18 222	153 923	7	genetic
<i>Amazon</i>	410 237	8 132 507	4	co-purchasing
<i>DBLP6</i>	892 531	6 045 859	6	co-authorship
<i>DBLP230</i>	892 531	7 207 253	230	co-authorship
<i>Twitter</i>	4 898 247	8 053 440	4	social

Since Hybrid applies  $k$ -Median<sup>+</sup>, followed by Greedy<sup>+</sup> in steps 2-3 above, it also requires  $k$ , a predefined number of supernodes, as an input. Therefore, in step 1, we apply the same preprocessing technique as in § 6 to determine the optimal  $k$  for  $k$ -Median<sup>+</sup>. However, notice that since we further apply Greedy<sup>+</sup> to improve the compactness in the third step, the optimal number of supernodes returned by Hybrid may eventually be reduced, in comparison with  $k$ -Median<sup>+</sup>.

## 8 EXPERIMENTAL RESULTS

We conduct experiments to demonstrate the effectiveness (compactness), efficiency, and scalability of our algorithms (averaged over 10 runs). The code is implemented in C++, and is executed on a single core, 40GB, 2.40GHz Xeon server.

### 8.1 Experimental setup

*8.1.1 Datasets.* We use five real-world, multi-relation networks, whose main characteristics are listed in Table 2.

***Homo*** (<https://comunelab.fbk.eu/data.php>) network describes different types of genetic interactions between genes in Homo Sapiens. Nodes are genes and edges denote their interactions. Seven different relations exist, which are: direct interaction, physical association, suppressive genetic interaction defined by inequality, association, colocalization, additive genetic interaction defined by inequality, and synthetic genetic interaction defined by inequality.

***Amazon*** (<https://snap.stanford.edu/data>) is a co-purchasing temporal network with four snapshots between March and June 2003, each as a relation. Nodes are products and edges are their co-purchasing relationships.

***DBLP*** (<http://dblp.uni-trier.de/xml>) is a well known collaboration network. We downloaded it on Dec 31, 2020. Each node is an author and edges denote their co-authorships. We use two versions of *DBLP* dataset. *DBLP6* has 6 relations for 6 representative sub-areas of computer science: data management, artificial intelligence, computer architecture, computer networks, theory of computing, and systems & software. An edge between a node pair exists for a relation if they have published as co-authors in the top-tier venues under this sub-area. The top-tier (i.e., rank A) conferences and journals for each sub-area are given by the CCF ranking: [https://www.ccf.org.cn/Academic\\_Evaluation/By\\_category](https://www.ccf.org.cn/Academic_Evaluation/By_category). *DBLP230* is generated based on [42] with the latest data. It has 230 relations for 230 keywords extracted from paper titles, based on both their frequency and how well they can represent various sub-areas of computer science, e.g., database systems, neural networks, FGPA, etc.

***Twitter*** (<https://iee-dataport.org/open-access/usa-nov2020-election-20-mil-tweets-sentiment-and-party-name-labels-dataset>) dataset is generated based on 24M US election related tweets from July 1 to November 11, 2020. Nodes are users and edges are their re-tweet relationships. The relations stand for 4 political parties in 2020 US presidential election.

**8.1.2 Competing algorithms.** Our basic two-step algorithms include (1) Greedy [57], (2) Randomized [57], (3)  $k$ -Median (proposed approximation algorithm for the Lossless-Sum problem), and (4) SWeG [71]. For fairness, we only allow sequential execution of the SWeG algorithm. The detailed description of these methods can be found in §3 and §4.1.

The holistic version of Greedy, Randomized, and  $k$ -Median are represented as GD+ (Greedy+), RD+ (Randomized+), and KM+ ( $k$ -Median+), respectively. Our final algorithm, Hybrid, for lossless, multi-relation graph summarization is represented as HY. In addition, we also consider a method, denoted as ALL, that stores individually optimal summaries (following the  $k$ -Median algorithm) for all relations.

The optimal number of supernodes (reported in Table 3) are automatically decided by GD+ and RD+. For  $k$ -Median, KM+, and HY, we decide the optimal number of supernodes empirically by Greedy methods, as discussed in §6.

**8.1.3 Evaluation metrics used.** We adopt the following evaluation metrics:

- **Relative size.** We measure the compactness of the obtained summaries by *relative size* [57, 71], which is defined as  $\frac{|E_S|+|C_S|}{|E|}$ . Recall that  $E$  denotes the set of edges in the input graph,  $E_S$  the set of superedges between supernodes, and  $C_S$  the set of correction edges. The numerator is our objective function (Problem 2), i.e., cost of the summary, and the denominator is constant for a given graph, i.e., graph size. Smaller relative size means better compactness.
- **Running time.** The total running time for generating a graph summary is reported. For two-step methods, it contains both the time of producing summaries for all individual relations and the time of aggregating them.
- **Storage cost.** We report the exact storage cost for original graphs and the corresponding summaries. The mappings from the node set  $V$  to the supernode set  $V_S$  are included in summaries.

## 8.2 Performance analyses

In Figure 8, the  $y$ -axis presents the relative size. Meanwhile, the  $x$ -axis reports the running time for summary construction. We want the summary to be as compact as possible, and the construction to be as fast as possible. Thus, *better solution would be closer to the origin point* in these plots.

The solid markers in Figure 8 stand for our proposed holistic algorithms, while the hollow markers represent the two-step methods. We observe that the solid markers are closer to the origin point on all datasets, which confirms the superiority of our proposed holistic algorithms, based on both summary compactness and its construction efficiency. Our Hybrid algorithm is shown with the shadow marker. In general, the shadow markers are below all others, demonstrating that *the Hybrid algorithm produces the most compact summaries*.

For  $k$ -Median-based methods ( $k$ -Median, KM+, and HY),  $k$  must be given as an input, whereas the optimal number of supernodes are automatically decided by Greedy, Randomized, SWeG, GD+, and RD+. For fairness of comparison, we find the optimal  $k$  for  $k$ -Median, KM+, and HY with an additional preprocessing step (§6), and report our results for  $k$ -Median, KM+, and HY in Figure 8 with this optimal  $k$  as input. This additional preprocessing time for determining the optimal  $k$  is provided in Table 4. The preprocessing time for  $k$ -Median is higher than that of KM+, since we need to identify an optimal  $k$  for each relation in  $k$ -Median, while we only require a single optimal  $k$  for KM+. HY consumes exactly the same preprocessing time as KM+, since its preprocessing step is same as that of KM+. However, in the third step of HY (§7), it further applies GD+ to improve the compactness, thus the final optimal number of supernodes may reduce for HY, in comparison with KM+ (as reported in Table 3).

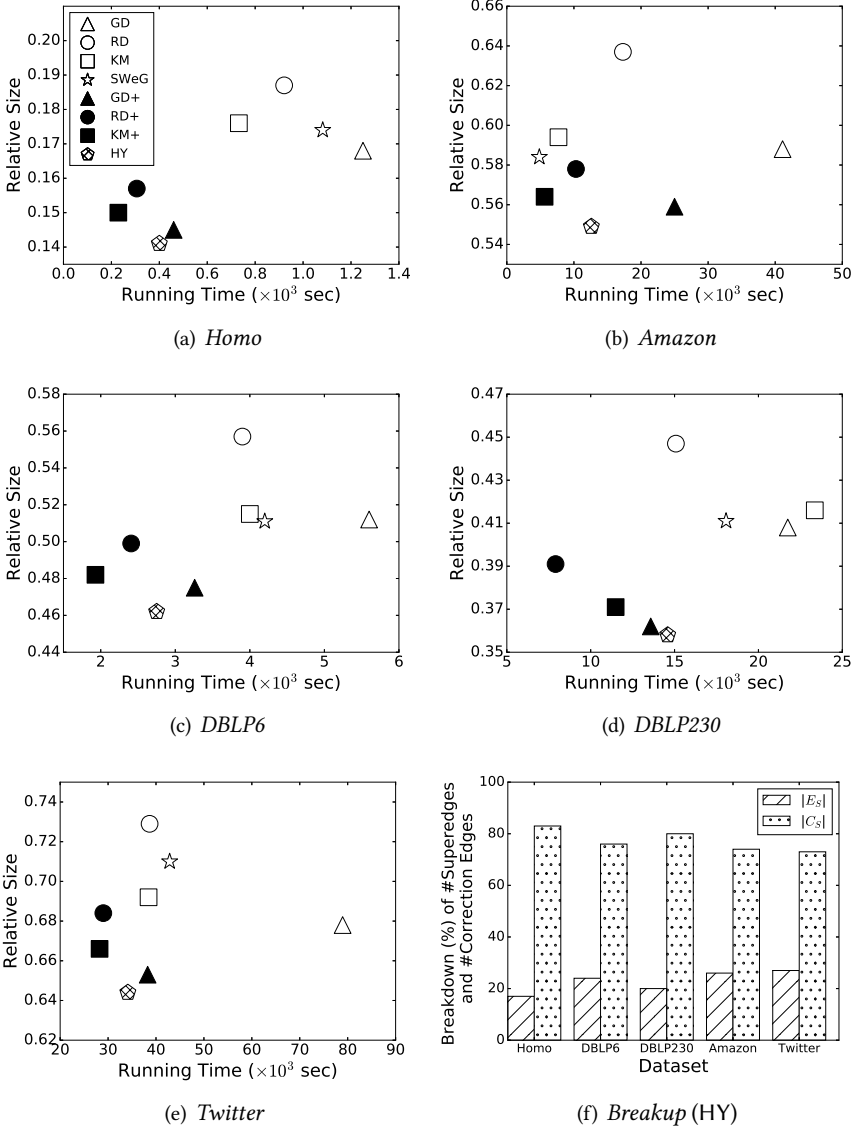


Fig. 8. Trade-off between relative size and running time. Sub-figure (f) presents breakup on #superedges and #correction-edges for summaries.

**Compactness.** For the same kind of algorithm, *the proposed holistic version can return up to 5% more compact summary than the corresponding two-step algorithm*, e.g., the KM+ summary always has at least 3% smaller relative size than the KM summary, over all datasets. Our ultimate HY method can result in about 1% more compact summaries than the best holistic method on all datasets.

Within each algorithm group, we find that (1) Greedy returns the most compact summaries, while Randomized produces the worst results; (2) the compactness of the summaries by SWeG and  $k$ -Median are in the middle range. Usually, they are comparable to the summaries by Greedy. These two observations hold both within the two-step methods and within the holistic algorithm groups.

Table 3. #supernodes, and (#superedges, #correction-edges) in obtained summaries via our holistic methods.

Graph	GD+	KM+	HY
<i>Homo</i>	440, (4.1K,18.2K)	440, (4.2K,18.7K)	428, (4.1K,18.1K)
<i>Amazon</i>	6.1K, (956K,3.6M)	6.1K, (959K,3.6M)	6.0K, (957K,3.6M)
<i>DBLP6</i>	12.2K, (399K,2.6M)	12.2K, (398K,2.8M)	12.1K, (398K,2.6M)
<i>DBLP230</i>	11.3K, (412K,2.3M)	11.3K, (418K,2.4M)	11.3K, (413K,2.2M)
<i>Twitter</i>	20.2K, (1.3M,4.2M)	20.2K, (1.3M,4.2M)	20.1K, (1.3M,4.1M)

Table 4. Additional preprocessing time ( $\times 10^3$  sec) of determining optimal  $k$  for  $k$ -Median-based methods.

Graph	KM	KM+	HY
<i>Homo</i>	1.0	0.4	0.4
<i>Amazon</i>	30	21	21
<i>DBLP6</i>	4.8	2.9	2.9
<i>DBLP230</i>	18	10	10
<i>Twitter</i>	51	32	32

Table 5. Actual storage cost (MB) for graphs and summaries. The actual storage cost for summaries includes the supernode mapping(s). We also report summary storage percentage w.r.t. original graph storage.

Graph	Original	GD+	KM+	HY	ALL
<i>Homo</i>	1.7	0.31 (18%)	0.32 (19%)	<b>0.30</b> (18%)	0.58 (34%)
<i>Amazon</i>	116.9	68.8 (59%)	69.3 (59%)	68.4 (58%)	<b>68.2</b> (58%)
<i>DBLP6</i>	74.3	36.2 (49%)	37.2 (50%)	<b>35.1</b> (47%)	37.2 (50%)
<i>DBLP230</i>	82.5	30.3 (37%)	31.1 (38%)	<b>30.0</b> (37%)	137.8 (167%)
<i>Twitter</i>	237.5	163.9 (69%)	165.0 (69%)	<b>160.7</b> (68%)	182.9 (77%)

The objective function of our summary has two components, the number of superedges  $|E_S|$  and the number of correction edges  $|C_S|$ . Figure 8(f) demonstrates that *the cost of correction list dominates the total summary cost*. It is always about 80% of the total cost.

**Efficiency.** For efficiency, we have the following observations. (1) The holistic algorithms (with solid markers) tend to consume less running time than the two-step methods. Recalling the complexity analysis listed in § 4.1, the two-step methods repeat the summary generation for every relation. This multiplies the time complexity of single summary computation by  $q$ , where  $q$  is the total number of relations. For example, the total time cost for producing  $q$  single-relation summaries will be  $O((m + nk \log n) \cdot q)$  for  $k$ -Median. An additional time for summary aggregation is also required. In contrast, for KM+, the time complexity is only  $O(m' + nk \log nq)$ . The first term  $m'$  is the total number of edges across all relations, which is similar to  $mq$ . However, the second term  $\log nq$  is much smaller than  $q \log n$ . This explains why our proposed holistic algorithms are faster than the two-step ones. (2) Usually, KM+ is the fastest among the proposed holistic methods, and GD+ is the slowest. (3) The running time of HY is always higher than that of KM+, since it requires to run KM+ at first, then applies GD+ to further improve the compactness.

### 8.3 Exact storage cost

We empirically study the exact storage cost of the original graphs and the summaries in the memory. Here, we further compare our uniform summary for all relations, against maintaining all the single-relation optimal summaries, denoted as ALL. For the exact storage, in addition to our objective

Table 6. Actual storage cost (MB) for graphs and summaries with further compression: (super)edges between the same set of (super)nodes over multiple relations are stored as  $\langle node\_i, node\_j, relation\_x, relation\_y, \dots, relation\_z \rangle$ . The actual storage cost for summaries includes the supernode mapping(s).

Graph	Original	GD+	KM+	HY	ALL
<i>Homo</i>	1.3	<b>0.17</b> (13%)	0.18 (14%)	<b>0.17</b> (13%)	0.54 (42%)
<i>Amazon</i>	94.9	50.8 (54%)	51.1 (54%)	<b>48.9</b> (52%)	60.6 (64%)
<i>DBLP6</i>	63.1	27.8 (44%)	28.2 (45%)	<b>26.7</b> (42%)	29.7 (47%)
<i>DBLP230</i>	70.6	24.8 (35%)	25.8 (36%)	<b>24.1</b> (34%)	126.0 (178%)
<i>Twitter</i>	211.5	137.6 (65%)	139.8 (66%)	<b>132.9</b> (63%)	151.4 (72%)

Table 7. Breakup of the actual storage cost (MB) for summaries with further compression (shown in Table 6).  $E_S$  denotes superedges,  $C$  denotes correction edges, and  $M$  denotes the node mapping.

Graph	HY			ALL		
	$ E_S $	$ C $	$ M $	$ E_S $	$ C $	$ M $
<i>Homo</i>	0.02	0.10	0.05	0.05	0.15	0.34
<i>Amazon</i>	10.5	37.0	1.4	15.2	36.6	8.8
<i>DBLP6</i>	5.2	19.9	1.5	5.4	20.0	4.3
<i>DBLP230</i>	4.7	17.8	1.5	6.6	22.9	96.1
<i>Twitter</i>	29.5	128.5	1.7	25.2	114.8	11.4

of superedges and correction list, we need to store the mapping  $M$  from the original node set  $V$  to the supernode set  $V_S$ . In the single-relation graph summary, and in our uniform summary for all relations, this mapping  $M$  simply has the same size as the cardinality of the original node set  $V$ . However, there exist  $q$  mappings for the ALL method, since the supernode partitioning can be different across relations. As shown in Table 5, in practice, the ALL summaries require more storage overhead than our holistic GD+ or HY summary. When the relation number is large, e.g., on *DBLP230*, the ALL summaries have even larger storage cost than the original graphs. Table 7 provides more insight through decomposing the storage cost.

In the storage format so far (reported in Table 5), each entry is represented as  $\langle node\_i, node\_j, relation \rangle$ , following our definition in § 1.1. However, the storage cost can further reduce when there exist (super)edges between the same set of (super)nodes over multiple relations. In such scenarios, a simple way to further reduce the exact storage cost is to keep them as  $\langle node\_i, node\_j, relation\_x, relation\_y, \dots, relation\_z \rangle$ . As shown in Table 6, our proposed summaries, both GD+ and HY, benefit more from this type of storage format, compared to the ALL summary.

## 8.4 Scalability analysis

We analyze the scalability of our methods on the larger datasets, *Twitter* and *DBLP*.

**Graph size.** The *Twitter* dataset has about 4.9 million nodes, we select 1M, 2M, 3M, 4M, and all 4.9M nodes uniformly at random to generate five graphs considering all relations, and apply our algorithms on them. Figure 9(a) demonstrates that *all our proposed holistic and hybrid algorithms scale linearly in graph size (i.e., number of nodes)*.

**Number of relations.** *DBLP\_230* dataset has 230 relations. We randomly choose 50, 100, 150, 200, and all 230 relations, and apply the algorithms on the full graphs of selected relations. Figure 9(b) shows that *the running times of all our algorithms increase linearly with more relations*. In the time

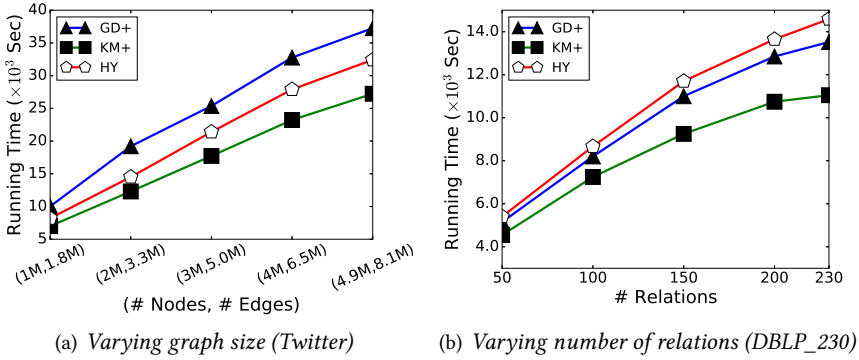


Fig. 9. Scalability analysis.

Table 8. Summary and recommendation.

Method	Summary Compactness	Summary Construction Time	Optimal $k$ Finding (Preprocessing) Time	Approx. Guarantee
GD	★★★	★	not required	✗
RD	★	★★★★	not required	✗
KM	★★	★★★★	★	✓
SWeG	★★	★★★★	not required	✗
GD+	★★★★★	★★	not required	✗
RD+	★★	★★★★★	not required	✗
KM+	★★★★	★★★★★	★★★	✓
HY	★★★★★	★★★	★★★	✓

complexity of KM+:  $O(m' + nk \log nq)$ ,  $m'$  scales about linearly in the number ( $q$ ) of relations, while the second term  $nk \log nq$  keeps nearly the same with increasing  $q$ .

## 8.5 Summary and recommendation

Table 8 summarizes the recommendation level of each method according to different performance metrics. The scale is from 1 to 5 stars, and larger star number stands for higher ranking. Clearly, there is no single winner. The holistic algorithms tend to produce more compact summary and consume less running time than the corresponding two-step version.  $k$ -Median-based approaches (KM, KM+, HY) have approximation guarantees on the summary compactness. According to empirical evaluations, GD+ returns the most compact summary among the holistic methods. However, HY applies GD+ to further improve the practical quality of the approximated solution returned by KM+, and is empirically shown to produce the best quality summaries (§ 8.2). For efficiency, KM+ is the fastest one. RD+ is always faster than GD+, and HY is always slower than KM+. The  $k$ -Median-based methods (KM, KM+, and HY) require an additional preprocessing step to identify the optimal  $k$ , and KM consumes the most amount of preprocessing time.

Based on application requirements, a user can decide to adopt a specific algorithm as per our summary in Table 8. In general case, considering various trade-offs, we recommend Hybrid method (HY) for multi-relation graph summarization. It has good performance in summary compactness from both theoretical and practical perspectives, and its efficiency lies in the middle range.

## 9 APPLICATIONS AND CASE STUDIES

### 9.1 Efficient query processing on graph summaries

Since our graph summary is lossless, we can always answer a graph query using the summary as precisely as in the original graph. Thus, we focus on the efficiency analyses. We present the




Fig. 10. Throughput comparison for neighborhood query answering on original graph and on different graph summaries.

comparison for the *Neighborhood Query* [50, 63]: Given a node  $v$  and a graph  $G = (V, E, R)$ , find the set of nodes  $N_v = \{u | (v, u, r) \in E\}$ , and return the distribution of relations in  $N_v$ . Notice that the *Degree Query* and the *Eigenvector-Centrality Query* in [50, 63] can be answered based on the results of neighborhood query; we do not consider them in the current study.

Figure 10 presents the throughput comparison for neighborhood query on the original graph and on the graph summaries (ALL and HY), using two larger datasets. Query nodes are divided into three groups: low ( $\leq 5$ ), medium ( $(5, 20]$ ), and high ( $> 20$ ), based on their out-degrees. For an hour, we continue to answer neighborhood queries for query nodes selected uniformly at random, and then report the average throughput (per minute). The neighborhood query can be processed more efficiently on graph summary since we explore the superedges and corrections instead of exact edges linked to the query node, and our objective ensures that the former tends to have smaller size than the latter. We have the following observations: **(1)** the throughput on HY summary is about  $2.5\times$  of that on the original graph for low-degree query nodes. It increases to be around  $5\times$  for high-degree nodes, since their edges are more likely to be wrapped within superedges. The benefit of using ALL summary to answer neighborhood query is  $1.5\times$  to  $3\times$  in throughput against using the original graph; **(2)** the efficiency improvement is more significant when the graph summary is more compact, e.g., in *DBLP\_230*. The intuition is that the edges of the query nodes are more likely to be represented via superedges, rather than corrections; and **(3)** when the number of relations is high, e.g., in *DBLP\_230*, the efficiency improvement for answering the query with HY summary is more significant than with ALL summary. This is because we need to repeatedly identify the supernode containing the query node in each relation for ALL summary, and the degree of the query node may be low in each relation, which makes the individual summaries within ALL less beneficial for query answering.

## 9.2 Visualization case study on DBLP

We visualize the summary of *DBLP230* dataset, and present some interesting case studies in Figure 11. In the first case shown in Figure 11(a), we find that all researchers in each supernode actively collaborate with others on the topics such as “data stream” and “approximation”. They have both self-loops and links to each other with such topics. However, a few different topics also appear within each group. For example, *Graham Cormode* also has quite a few “private”-related (i.e., privacy) joint works with *Divesh Srivastava*. Another interesting finding is that both the left-side and the top supernodes contain researchers from database community, since they have more publications in *SIGMOD* and *VLDB*, while the right-side supernode consists of theory community researchers, who publishes more in *STOC*, *SODA*, etc.



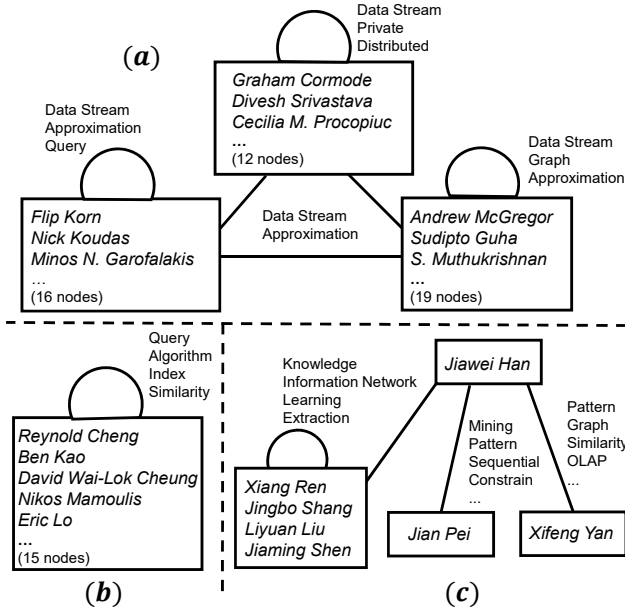


Fig. 11. Case study on *DBLP*.




Fig. 12. Case study on *Austism*. The partitioning of the same set of 14 nodes are quite different for ASD brain networks and for healthy brain networks. The numbers on superedges denote the number of relations having such superedges.

The second case in Figure 11(b) presents a group of people with close collaboration. One can verify that they are all from same geographical location, i.e., *Hong Kong* in this example. This is a frequent pattern in *DBLP*.

Figure 11(c) first shows that very senior researchers tend to be kept as a single node, since they work on quite diverse topics and with many researchers. Here, *Jiawei Han* and his ex-students, *Jian Pei* and *Xifeng Yan*, are all kept as a singleton supernode, and *Jiawei Han* has different collaboration topics with them. His recently graduated student, *Xiang Ren* (in 2018), and some other current students are grouped together, since they collaborate frequently with each other on topics related to knowledge extraction. Such correlations across different relations and node set could not be immediately inferred if one keeps individual summaries for all relations separately.

### 9.3 Visualization and classification case studies on brain networks

For our second case study, we use the *Austism* dataset [24] containing 96 brain networks. 48 of them are collected from ASD (Autism spectrum disorder) patients, and the other 48 brain networks are

of healthy people. In the original dataset, there were 49 networks of ASD patients and 52 networks of healthy people. To maintain a balanced number of networks in each group, and for the ease of our comparison, we only keep the first 48 networks in each group here. All networks share the same set of 116 nodes. The average number of edges in each network is 1336.8 (1336.2 for ASD patients, and 1337.5 for healthy people). If we treat the 48 ASD brain networks as a multi-relation graph, HY returns a summary with 0.514 relative size (0.516 by GD+, 0.517 by KM+). Meanwhile, the relative size of HY multi-relation summary for the 48 healthy people is 0.526 (0.526 by GD+, 0.529 by KM+). However, if we randomly choose 24 networks from each group, and generate a multi-relation summary for them, the relative size will be 0.559 by HY (0.559 by GD+, 0.560 by KM+, both are the average value over 20 attempts). Clearly, it is harder to summarize the mixed group, which implies that *the networks of ASD patients have different structure compared to those of healthy people*, and inspires us to *utilize the multi-relation summaries for ASD patient detection*.

**Classification.** We randomly choose two brain networks, one for ASD patient and one for healthy people. A multi-relation summary  $S_1$  is generated for the rest 47 ASD brain networks, and another multi-relation summary  $S_2$  for the rest 47 healthy ones. Then, we alternatively apply  $S_1$  and  $S_2$  as summaries for the two selected networks, and suggest the label “ASD” to a network if  $S_1$  results in lower cost than  $S_2$ . We repeat the procedure 20 times. Within the 20 true ASD patients, 19 of them are correctly detected. For the 20 healthy people, 3 of them are wrongly labelled as “ASD”. It can be easily calculated that our classification precision is 0.86, recall is 0.95, and F1-score is 0.90. This demonstrates the effectiveness of our multi-relation summaries in this classification task.

**Visualization.** Finally, we visualize some subgraphs of our multi-relation summaries for the two groups of brain networks. As shown in Figure 12, the partitionings of the same set of 14 nodes are quite different for ASD brain networks and for healthy brain networks. Although these supernodes are well-connected to each other, they can not be further merged since they have different superedges.

## 10 CONCLUSIONS

In this paper, we first revisited the classic single-relation graph summarization problem, and provided the first polynomial-time approximation algorithm based on the  $k$ -Median clustering. Then we introduced and investigated the novel problem of multi-relation graph summarization. To solve the problem, we first studied the baseline two-step approaches: first generate a summary for each relation, and then properly aggregate them. We further demonstrated and discussed the limitations of these baselines, and proposed holistic solutions to overcome them. Among them, the holistic  $k$ -Median<sup>+</sup> is able to maintain the approximation guarantee over multi-relation graphs. Finally, we developed Hybrid algorithm as our ultimate solution, by utilizing both the strengths of  $k$ -Median<sup>+</sup> and Greedy<sup>+</sup>. Our experimental results and case studies validated the effectiveness and efficiency of our algorithms.

## ACKNOWLEDGMENTS

Arijit Khan is supported by MOE Tier1 and Tier2 grants RG117/19 and MOE2019T2-2-042.

## REFERENCES

- [1] K. J. Ahn, S. Guha, and A. McGregor. 2012. Analyzing Graph Structure via Linear Measurements. In *SODA*.
- [2] K. J. Ahn, S. Guha, and A. McGregor. 2012. Graph Sketches: Sparsification, Spanners, and Subgraphs. In *PODS*.
- [3] S. Bandyopadhyay, M. Mehta, D. Kuo, M.-K. Sung, R. Chuang, E. J. Jaehng, B. Bodenmiller, K. Licon, W. Copeland, M. Shales, D. Fiedler, J. Dutkowski, A. Guénolé, H. van Attikum, K. M. Shokat, R. D. Kolodner, W.-K. Huh, R. Aebersold, M.-C. Keogh, N. J. Krogan, and T. Ideker. 2010. Rewiring of Genetic Networks in Response to DNA Damage. *Science* 330, 6009 (2010).
- [4] N. Bansal, A. Blum, and Chawla S. 2004. Correlation Clustering. *Machine Learning* 56 (2004), 89–113.

- [5] M. A. Beg, M. Ahmad, A. Zaman, and I. Khan. 2018. Scalable Approximation Algorithm for Graph Summarization. In *PAKDD*.
- [6] A. A. Benczúr and D. R. Karger. 2015. Randomized Approximation Schemes for Cuts and Flows in Capacitated Graphs. *SIAM J. Comput.* 44, 2 (2015), 290–319.
- [7] M. Besta and T. Hoefler. 2018. Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations. *CoRR* abs/1806.01799 (2018).
- [8] B. Boden, S. Günemann, H. Hoffmann, and T. Seidl. 2012. Mining Coherent Subgraphs in Multi-layer Graphs with Edge Labels. In *KDD*.
- [9] P. Boldi, M. Rosa, M. Santini, and S. Vigna. 2011. Layered Label Propagation: A Multiresolution Coordinate-Free Ordering for Compressing Social Networks. In *WWW*.
- [10] P. Boldi and S. Vigna. 2004. The Webgraph Framework I: Compression Techniques. In *WWW*.
- [11] N. R. Brisaboa, S. Ladra, and G. Navarro. 2009. k2-Trees for Compact Web Graph Representation. In *SPIRE*.
- [12] A. Z. Broder, M. Charikar, A. M. Frieze, and M. Mitzenmacher. 2000. Min-wise Independent Permutations. *J. Comput. System Sci.* 60 (2000), 630–659.
- [13] G. Buehrer and K. Chellapilla. 2008. A Scalable Pattern Mining Approach to Web Graph Compression with Communities. In *WSDM*.
- [14] A. Cardillo, J. Gómez-Gardeñes, M. Zanin, M. Romance, D. Papo, F. del Pozo, and S. Boccaletti. 2012. Emergence of Network Features from Multiplexity. *CoRR* abs/1212.2153 (2012).
- [15] M. Charikar, V. Guruswami, and A. Wirth. 2005. Clustering with Qualitative Information. *J. Comput. Syst. Sci.* 71, 3 (2005), 360–383.
- [16] C. Chen, C. X. Lin, M. Fredrikson, M. Christodorescu, X. Yan, and J. Han. 2009. Mining Graph Patterns Efficiently via Randomized Summaries. *PVLDB* 2, 1 (2009), 742–753.
- [17] C. Chen, C. X. Lin, M. Fredrikson, M. Christodorescu, X. Yan, and J. Han. 2010. Mining Large Information Networks by Graph Summarization. In *Link Mining: Models, Algorithms, and Applications*. 475–501.
- [18] C. Chen, X. Yan, F. Zhu, J. Han, and P. S. Yu. 2008. Graph OLAP: Towards Online Analytical Processing on Graphs. In *ICDM*.
- [19] F. Chierichetti, R. Kumar, S. Lattanzi, M. Mitzenmacher, A. Panconesi, and P. Raghavan. 2009. On Compressing Social Networks. In *KDD*.
- [20] Y. Choi and W. Szpankowski. 2012. Compression of Graphical Structures: Fundamental Limits, Algorithms, and Experiments. *IEEE Trans. Information Theory* 58, 2 (2012), 620–638.
- [21] D. J. Cook and L. B. Holder. 1994. Substructure Discovery Using Minimum Description Length and Background Knowledge. *J. Artif. Intell. Res.* 1 (1994), 231–255.
- [22] G. Cormode and S. Muthukrishnan. 2005. Space efficient mining of multigraph streams. In *PODS*.
- [23] M. Coscia, G. Rossetti, D. Pennacchioli, D. Ceccarelli, and F. Giannotti. 2013. "You Know because I Know": A Multidimensional Network Approach to Human Resources Problem. In *ASONAM*.
- [24] C. Craddock, Y. Benhajali, C. Chu, F. Chouinard, A. Evans, A. Jakab, B. S. Khundrakpam, J. D. Lewis, Q. Li, M. Miham, C. Yan, and P. Bellec. 2013. The Neuro Bureau Preprocessing Initiative: Open Sharing of Preprocessed Neuroimaging Data and Derivatives. *Frontiers in Neuroinformatics* 41 (2013).
- [25] M. E. Dickison, M. Magnani, and L. Rossi. 2016. *Multilayer Social Networks*. Cambridge University Press.
- [26] W. Fan, J. Li, X. Wang, and Y. Wu. 2012. Query Preserving Graph Compression. In *SIGMOD*.
- [27] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, and J. Zhang. 2008. Graph Distances in the Data-Stream Model. *SIAM J. Comput.* 38, 5 (2008), 1709–1727.
- [28] E. Galimberti, F. Bonchi, and F. Gullo. 2017. Core Decomposition and Densest Subgraph in Multilayer Networks. In *CIKM*.
- [29] Edoardo Galimberti, Francesco Bonchi, Francesco Gullo, and Tommaso Lanciano. 2020. Core Decomposition in Multilayer Networks: Theory, Algorithms, and Applications. *ACM Trans. Knowl. Discov. Data* 14, 1 (2020), 11:1–11:40.
- [30] A. Gionis, H. Mannila, and P. Tsaparas. 2007. Clustering Aggregation. *TKDD* 1 (2007).
- [31] A. Gionis and C. E. Tsourakakis. 2015. Dense Subgraph Discovery: KDD 2015 Tutorial. In *KDD*.
- [32] X. Gou, L. Zou, C. Zhao, and T. Yang. 2019. Fast and Accurate Graph Stream Summarization. In *ICDE*.
- [33] N. Hassanlou, M. Shoaran, and A. Thomo. 2013. Probabilistic Graph Summarization. In *WAIM (Lecture Notes in Computer Science, Vol. 7923)*. Springer.
- [34] P. Hu and W. C. Lau. 2013. A Survey and Taxonomy of Graph Sampling. *CoRR* abs/1308.5865 (2013).
- [35] N. Jayaram, A. Khan, C. Li, X. Yan, and R. Elmasri. 2015. Querying Knowledge Graphs by Example Entity Tuples. *IEEE Trans. Knowl. Data Eng.* 27, 10 (2015), 2797–2811.
- [36] B. Jin, C. Gao, X. He, D. Jin, and Y. Li. 2020. Multi-behavior Recommendation with Graph Convolutional Networks. In *SIGIR*.

- [37] D. Jin, R. A. Rossi, E. Koh, S. Kim, A. Rao, and D. Koutra. 2019. Latent Network Summarization: Bridging Network Embedding and Summarization. In *KDD*.
- [38] k. Lee, H. Jo, J. Ko, S. Lim, and K. Shin. 2020. SSumM: Sparse Summarization of Massive Graphs. In *KDD*.
- [39] U. Kang and C. Faloutsos. 2011. Beyond ‘Caveman Communities’: Hubs and Spokes for Graph Compression and Mining. In *ICDM*.
- [40] U. Kang, H. Tong, J. Sun, C.-Y. Lin, and C. Faloutsos. 2011. GBASE: A Scalable and General Graph Management System. In *KDD*.
- [41] G. Karypis and V. Kumar. 1995. Analysis of Multilevel Graph Partitioning. In *Supercomputing*.
- [42] X. Ke, A. Khan, and G. Cong. 2018. Finding Seeds and Relevant Tags Jointly: For Targeted Influence Maximization in Social Networks. In *SIGMOD*.
- [43] A. Khan and C. C. Aggarwal. 2017. Toward Query-Friendly Compression of Rapid Graph Streams. *Social Netw. Analys. Mining* 7, 1 (2017), 23:1–23:19.
- [44] A. Khan, S. S. Bhowmick, and F. Bonchi. 2017. Summarizing Static and Dynamic Big Graphs. *PVLDB* 10, 12 (2017), 1981–1984.
- [45] K. U. Khan, W. Nawaz, and Y.-K. Lee. 2014. Set-Based Unified Approach for Attributed Graph Summarization. In *BDCLOUD*.
- [46] D. Koutra, U. Kang, J. Vreeken, and C. Faloutsos. 2014. VOG: Summarizing and Understanding Large Graphs. In *SDM*.
- [47] D. Koutra, U. Kang, J. Vreeken, and C. Faloutsos. 2015. Summarizing and Understanding Large Graphs. *Statistical Analysis and Data Mining* 8, 3 (2015), 183–202.
- [48] D. Koutra, J. Vreeken, and F. Bonchi. 2018. Summarizing Graphs at Multiple Scales: New Trends. In *ICDM*.
- [49] K. A. Kumar and P. Efstathopoulos. 2018. Utility-Driven Graph Summarization. *PVLDB* 12, 4 (2018), 335–347.
- [50] K. LeFevre and E. Terzi. 2010. GraSS: Graph Structure Summarization. In *SDM*.
- [51] S.-D. Lin, M.-Y. Yeh, and C.-T. Li. 2013. Sampling and Summarization for Social Networks. In *SDM*.
- [52] X. Liu, Y. Tian, Q. He, W.-C. Lee, and J. McPherson. 2014. Distributed Graph Summarization. In *CIKM*.
- [53] Y. Liu, T. Safavi, A. Dighe, and D. Koutra. 2018. Graph Summarization Methods and Applications: A Survey. *ACM Comput. Surv.* 51, 3 (2018), 62:1–62:34.
- [54] M. Mao, J. Lu, G. Zhang, and J. Zhang. 2017. Multirelational Social Recommendations via Multigraph Ranking. *IEEE Trans. Cybernetics* 47, 12 (2017), 4049–4061.
- [55] H. Maserrat and J. Pei. 2010. Neighbor Query Friendly Compression of Social Networks. In *KDD*.
- [56] H. Maserrat and J. Pei. 2012. Community Preserving Lossy Compression of Social Networks. In *ICDM*.
- [57] S. Navlakha, R. Rastogi, and N. Shrivastava. 2008. Graph Summarization with Bounded Error. In *SIGMOD*.
- [58] M. E. J. Newman and M. Girvan. 2004. Finding and Evaluating Community Structure in Networks. *Phys. Rev. E* 69 (2004), 026113. Issue 2.
- [59] D. L. Phuoc, H. N. M. Quoc, H. N. Quoc, T. T. Nhat, and M. Hauswirth. 2016. The Graph of Things: A Step towards the Live Knowledge Graph of Connected Things. *J. Web Semant.* 37-38 (2016), 25–35.
- [60] M. Purohit, B. A. Prakash, C. Kang, Y. Zhang, and V. S. Subrahmanian. 2014. Fast Influence-based Coarsening for Large Networks. In *KDD*.
- [61] Q. Qu, S. Liu, C. S. Jensen, F. Zhu, and C. Faloutsos. 2014. Interestingness-Driven Diffusion Process Summarization in Dynamic Networks. In *ECML PKDD*.
- [62] S. Raghavan and H. Garcia-Molina. 2003. Representing Web Graphs. In *ICDE*.
- [63] M. Riondato, D. Garcia-Soriano, and F. Bonchi. 2014. Graph Summarization with Quality Guarantees. In *ICDM*.
- [64] Matteo Riondato, David Garcia-Soriano, and Francesco Bonchi. 2017. Graph summarization with quality guarantees. *Data Min. Knowl. Discov.* 31, 2 (2017), 314–349.
- [65] J. Rissanen. 1978. Modelling by the Shortest Data Description. *Automatica* 14 (1978).
- [66] R. A. Rossi and R. Zhou. 2018. GraphZIP: A Clique-based Sparse Graph Compression Method. *J. Big Data* 5 (2018), 10.
- [67] B.-S. Seah, S. S. Bhowmick, and C. F. Dewey Jr. 2014. DiffNet: Automatic Differential Functional Summarization of dE-MAP Networks. *Methods* 63, 3 (2014).
- [68] B.-S. Seah, S. S. Bhowmick, C. F. Dewey Jr, and H. Yu. 2012. FUSE: A Profit Maximization Approach for Functional Summarization of Biological Networks. *BMC Bioinformatics* 13 (2012).
- [69] N. Shah, D. Koutra, T. Zou, B. Gallagher, and C. Faloutsos. 2015. TimeCrunch: Interpretable Dynamic Graph Summarization. In *KDD*.
- [70] L. Shi, S. Sun, Y. Xuan, Y. Su, H. Tong, S. Ma, and Y. Chen. 2016. TOPIC: Toward Perfect Influence Graph Summarization. In *ICDE*.
- [71] K. Shin, A. Ghoting, M. Kim, and H. Raghavan. 2019. SWeG: Lossless and Lossy Summarization of Web-Scale Graphs. In *WWW*.
- [72] M. Stella, C. S. Andreazzi, S. Selakovic, A. Goudarzi, and A. Antonioni. 2017. Parasite Spreading in Spatial Ecological Multiplex Networks. *J. Complex Networks* 5, 3 (2017), 486–511.

- [73] N. Tang, Q. Chen, and P. Mitra. 2016. Graph Stream Summarization: From Big Bang to Big Crunch. In *SIGMOD*.
- [74] R. L. Thorndike. 1953. Who Belongs in the Family? *Psychometrika* 18 (1953), 267–276.
- [75] Y. Tian, R. A. Hankins, and J. M. Patel. 2008. Efficient Aggregation for Graph Summarization. In *SIGMOD*.
- [76] Y. Tian and J. M. Patel. 2010. Interactive Graph Summarization. In *Link Mining: Models, Algorithms, and Applications*. 389–409.
- [77] H. Toivonen, F. Zhou, A. Hartikainen, and A. Hinkka. 2011. Compression of Weighted Graphs. In *KDD*.
- [78] I. Tsalouchidou, G. De Francisci Morales, F. Bonchi, and R. Baeza-Yates. 2016. Scalable dynamic graph summarization. In *BigData*.
- [79] S. White and P. Smyth. 2005. A Spectral Clustering Approach To Finding Communities in Graph. In *SDM*.
- [80] Y. Wu, S. Yang, M. Srivatsa, A. Iyengar, and X. Yan. 2013. Summarizing Answer Graphs Induced by Keyword Queries. *PVLDB* 6, 14 (2013), 1774–1785.
- [81] L. Xia, C. Huang, Y. Xu, P. Dai, X. Zhang, H. Yang, J. Pei, and L. Bo. 2021. Knowledge-enhanced Hierarchical Graph Transformer Network for Multi-behavior Recommendation. In *AAAI*.
- [82] X. Yan, X. J. Zhou, and J. Han. 2005. Mining Closed Relational Graphs with Connectivity Constraints. In *KDD*.
- [83] J. Yang, J. You, and X. Wan. 2021. Graph Embedding via Graph Summarization. *IEEE Access* 9 (2021), 45163–45174.
- [84] J. Zhang, S. S. Bhowmick, H. H. Nguyen, B. Choi, and F. Zhu. 2015. DaVinci: Data-driven Visual Interface Construction for Subgraph Search in Graph Databases. In *ICDE*.
- [85] X. Zhang and T. Özsu. 2019. Correlation Constraint Shortest Path over Large Multi-Relation Graphs. *PVLDB* 12, 5 (2019), 488–501.
- [86] P. Zhao, C. C. Aggarwal, and M. Wang. 2011. gSketch: On Query Estimation in Graph Streams. *PVLDB* 5, 3 (2011), 193–204.
- [87] F. Zhou, S. Mahler, and H. Toivonen. 2010. Network Simplification with Minimal Loss of Connectivity. In *ICDM*.