

Summarizing Labeled Multi-Graphs

Dimitris Berberidis¹ Pierre J. Liang² Leman Akoglu¹ ✉

¹ Carnegie Mellon University, Heinz College of Information Systems and Public Policy
{dbermper, lakoglu}@andrew.cmu.edu

² Carnegie Mellon University, Tepper School of Business
liangj@tepper.cmu.edu

Abstract. Real-world graphs can be difficult to interpret and visualize beyond a certain size. To address this issue, graph summarization aims to simplify and shrink a graph, while maintaining its high-level structure and characteristics. Most summarization methods are designed for homogeneous, undirected, simple graphs; however, many real-world graphs are *ornate*; with characteristics including node labels, directed edges, edge multiplicities, and self-loops. In this paper we propose TG-SUM, a *versatile* yet rigorous graph summarization model that (to the best of our knowledge, for the first time) can handle graphs with *all* the aforementioned characteristics (and *any* combination thereof). Moreover, our proposed model captures basic sub-structures that are prevalent in real-world graphs, such as cliques, stars, etc. TG-SUM compactly quantifies the information content of a complex graph using a novel encoding scheme, where it seeks to minimize the total number of bits required to encode (i) the summary graph, as well as (ii) the corrections required for reconstructing the input graph *losslessly*. To accelerate the summary construction, it creates super-nodes efficiently by merging nodes *in groups*. Experiments demonstrate that TG-SUM facilitates the visualization of real-world complex graphs, revealing interpretable structures and high-level relationships. Furthermore, TG-SUM achieves better trade-off between compression rate and running time, relative to existing methods (only) on comparable settings.

Keywords: graph summarization · super graph · labeled multi-graph · MDL

1 Introduction

Given a directed labeled multi-graph G , how can we construct a small summary graph g that reflects the high-level structures and relationships in G ? How can we find a succinct g that is yet an accurate representation, which requires a small amount of corrections to recover the original G ? With the advent of technology, not only the size but also the complexity of real-world graphs have grown immensely. Today graph data often contains node labels, multi-edges, etc. Graph summarization aims to find high-level structural patterns and most salient information in large complex graphs to enable efficient storage, processing, visualization and interpretation.

A large body of existing graph summarization techniques is for *plain* graphs with homogeneous unlabeled nodes [15,25,17,16,24,27,21]. However there exist numerous real-world graphs with multiple node labels; including transaction networks containing nodes (i.e. accounts) of various types (cash, revenue, expense, etc.) or heterogeneous graphs such as publication records among entities of various types (paper, author, venue, etc.). We refer to both kinds as node-labeled, or simply *labeled graphs*. Moreover, a vast majority of prior work are for summarizing *simple* [15,25,17,16,21,29,20,26,9] *undirected* [15,25,17,16,24,27,21,29,26] graphs, whereas the edges in real-world graphs may repeat (e.g., multiple transactions between two accounts, multiple exchanges

between two email addresses, etc.) which are called *multi-graphs*. As is the case for transaction and email graphs, among others, the edges can also be *directed*.

In this work we propose (to the best of our knowledge; see, Table 1) the *first* method called TG-SUM, for *multi-Type* (i.e. node-labeled) *multi-Graph SUM*marization, with directed edges and possible self-loops. (See Sec. 2, and a recent comprehensive survey [19].) Besides, TG-SUM is *versatile* in that it can also handle graphs with any combination of those properties (i.e., (un)directed, plain/labeled, simple/multi- or weighted graphs).

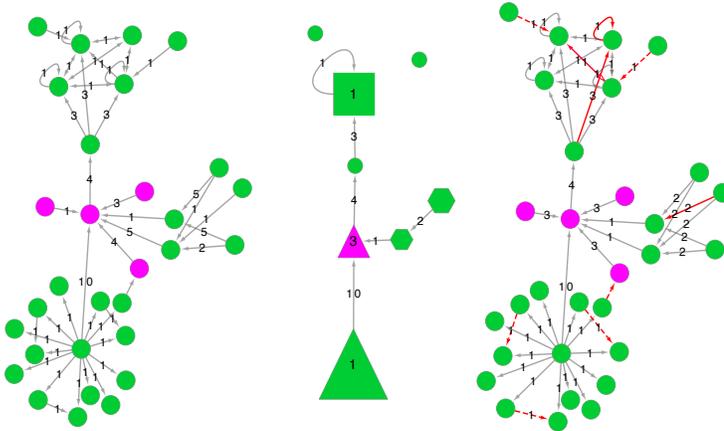


Fig. 1: (best in color) Ex. input graph (left), its summary/super-graph (middle), and the decompressed graph (right) w/ edge corrections in red, where dashed $--$ (solid $---$) are edges that need to be added (removed) for lossless reconstruction. See text for description of the scalars, node color, size, and shape.

A super-node is also marked with a scalar (i.e., weight), representative of the edge multiplicities among its nodes. A super-edge is placed between two super-nodes whose constituent nodes are sufficiently well-connected, and is also marked with a scalar (i.e., weight) that best represents the edge multiplicities inbetween.

We aim to construct a small summary graph, which accurately reflects the input graph. Here, succinctness and accuracy are in trade-off; the coarser the summary graph, the more information about the original graph is lost. We design a novel two-part *lossless* encoding scheme, describing (i) the summary graph and (ii) the corrections required to reconstruct the input graph losslessly. Treating the total number of encoding bits as a cost function, we design algorithms to find a summary with a small total cost. In summary, our main contributions are:

- **Summarizing LMDS-Graphs:** TG-SUM is the first graph summarization method that can simultaneously handle node-Labels, edge Multiplicities, edge Directions, and Self-loops if any, as well as any combination thereof (Sec. 2). LMDS-Graphs form the most general type of graphs, and frequently offer the most complete representation of domain specific data (e.g., accounting, social media, etc.).
- **Novel Super-graph Model:** We build an interpretable summary graph consisting of super-nodes and super-edges. Super-nodes capture nodes with similar connectivity, e.g., within stars or cliques. A representative multiplicity is found for each set of summarized edges by a super-edge and within a super-node. (Sec. 3.1)

Our goal is to output a small yet representative summary that facilitates the visualization, by which, improves the understanding of the overall structure of an input graph. To this end, we model a *summary graph* (or *super-graph*) as a collection of labeled super-nodes and weighted super-edges. As illustrated in Fig. 1, we merge structurally similar nodes of the same type/label (depicted by color) into super-nodes capture prevalent structural constructs found in real-world graphs, such as stars and cliques [16] (depicted by shape).

- **Novel Two-part Lossless Encoding Scheme:** We introduce a novel scheme for encoding (i) the summary graph as well as (ii) the corrections for reconstructing the input graph losslessly. (Sec. 3.2)
- **Efficient Search Algorithms:** We propose a fast algorithm for searching groups of structurally-similar nodes to merge into super-nodes. Our candidate set generation avoids pairwise computations and allows multi-resolution summaries. (Sec. 4)

Extensive experiments on 12 real-world graphs with varying properties show the versatility of TG-SUM, which reveals insights that are otherwise hard to draw. TG-SUM also achieves better compression with increasing allotted running time relative to the baselines on comparable (i.e., non-LMDS) settings, and scales near-linearly with input size. (Sec. 5)

Reproducibility: Source code for TG-SUM and all public-domain datasets are shared at <https://anonymous.4open.science/r/4243006d-fede-49bd-b549-02387edf7ccd/>.

2 Related Work

Table 1: **Our method matches *all* specs**, while prior work miss one or more of the graph or summary properties. * depicts the cases that the method can be modified to handle, although details are not given in the corresponding paper.

Method	Input graph				Summary		
	Labeled nodes	Directed edges	edge Multiplicity	Self-loops	Super-graph	Lossless	Multiresolution
Khan <i>et al.</i> [15], SWeG [25]					✓	✓	
GraSS [17]					✓	✓	✓
VoG [16]						✓	
Riondato <i>et al.</i> [24]			✓	✓	✓		✓
Toivonen <i>et al.</i> [27]			✓		✓		✓
CoSum [29]	✓				✓		
Liu <i>et al.</i> [20]	✓	✓			✓		✓
SNAP [26]	✓	✓*			✓		✓
Subdue [9]	✓	✓		✓	✓		✓
Navlakha <i>et al.</i> [21]	✓*	✓*			✓	✓	✓
TG-SUM [this paper]	✓	✓	✓	✓	✓	✓	✓

Riondato *et al.*[24] and Toivonen *et al.*[27] are some of the few summarization methods that can accommodate weighted edges, but not labeled nodes or directed edges. Among the methods that can handle graphs with multiple node labels, CoSum [29], Liu *et al.*[20], and SNAP [26] build a coarser graph by only merging the nodes of the same label into super-nodes. Differently, Subdue

Graph *summarization* and graph *compression* techniques, while related, exhibit a key distinction. The former typically aims to simplify an input graph into a coarser one, while reflecting its prominent structure. On the other hand, the latter aims at reducing the storage requirements of a graph, often enabling speedy querying, while maintaining a certain level of query accuracy [2,6,7,14,18]. (See [5] for a recent survey.) In this work we focus on graph summarization, with a goal to extract a simplified overview of key structural patterns within an input graph. Table 1 provides a structured comparison of related work, with respect to the properties of the input graph as well as properties of the summary.

Most graph summarization techniques are designed for unlabeled, undirected, and simple graphs without edge multiplicities, weights, or self-loops [25,15,17,16]. Closely-related are graph-pooling methods used within graph neural networks to gradually reduce the dimension of the layers; see. e.g. [28].

[9] replaces frequent sub-graphs that potentially contain different labels with a super-node, which makes the interpretation of the summary graph harder.

Closest to our work is the approach by Navlakha *et al.*[21], which iteratively merges nodes into super-nodes as long as the description length of the input graph decreases. Thanks to its simple model and algorithm, it can be modified to handle labeled graphs, specifically by restricting the node merges to same-label nodes. However, its model is unable, nor is it trivial to modify, to accommodate edge weights/multiplicities. All in all, there is *no existing work* that can summarize labeled multi-graphs – with labeled nodes, directed and multi-edges and self-loops. (See [19] for an extended survey and Table 1 therein.) While our TG-SUM is the first of its kind, it is *versatile* in that it can also accommodate graphs with any combination of those properties.

Besides input graph properties, prior work can also be classified w.r.t. the properties of the summary. Here, we focus on summaries where the output is itself a (coarser) graph, called the summary or super graph. VoG [16] is an exception in Table 1; it identifies key sub-structures (stars, (near)cliques, etc.) however does not provide any super-edges, i.e., its summary graph is disconnected. Second, the summary may be lossy; including only the coarse summary graph [24,27,29,20,26,9], or lossless; consisting of both the summary and the corrections necessary to fully reconstruct the input graph [25,15,17,16,21]. Finally, a desired characteristic of a summary is multi-granularity; where the coarseness or resolution of the summary graph can be adjusted on demand [17,24,27,20,26,9,21], via appropriately altering some of the model parameters. Notably, TG-SUM exhibits all of these three properties: super graph output, lossless and multi-resolution summary.

3 Graph Summary Design and Encoding

3.1 Summary and Decompression

Given a directed graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, \mathcal{T}\}$ with edge multiplicities $m(e) \in \mathbb{N}, \forall e \in \mathcal{E}$, node labels/types $\ell(u) \in \mathcal{T}, \forall u \in \mathcal{V}$, and self-loops, we define the *summary* and *decompressed* graphs as follows.

3.1.1 Summary graph (or super-graph) model Let $\mathcal{G}_s = \{\mathcal{V}_s, \mathcal{E}_s\}$ be the sets of super-nodes and directed super-edges that define the summary graph topology. Each super-node $v \in \mathcal{V}_s$ is annotated by four components: (i) its label $\ell(v)$ (depicted by color), (ii) the number $|\mathcal{S}_v|$, $\mathcal{S}_v \subset \mathcal{V}$, of nodes it contains (depicted by size), (iii) the *glyph* $\mu(v) \in \mathcal{M}$ it represents (depicted by shape), and (iv) the *representative* multiplicity $m(v)$ of the edges it summarizes (depicted as a scalar inside the glyph). For each super-edge $e \in \mathcal{E}_s$, we let $m(e)$ be the *representative* multiplicity of the edges it captures, depicted as a scalar on the super-edge. (We describe how to find the “representative” multiplicity of a set of edges in Sec. 4.2.)

Fig. 1 (left) and (middle) respectively depict an example input graph and its corresponding summary graph. Apart from unmerged simple nodes that are depicted as plain circles, the set \mathcal{M} of possible glyphs that TG-SUM supports contains: 1) **Clique** (square), 2) **In-star** (triangle), 3) **Out-star** (triangle), and 4) **Disconnected** set (hexagon). Such structures are commonly found in real-word graphs [16]. For instance, a clique can represent a tightly-knit group of friends in a social network, while an out-star can capture spam-like activity in an email or call network. Moreover, using glyphs has been shown to yield easily interpretable visualizations [10].

3.1.2 Decompression The summary graph \mathcal{G}_s decompresses *uniquely* and *unambiguously* into $\mathcal{G}' = \text{dec}(\mathcal{G}_s) = \{\mathcal{V}, \mathcal{E}'\}$ according to simple and intuitive rules (e.g., Fig. 1 (right)). First, every

super-node expands to the set of nodes it contains, all of which also inherit the super-node’s label. The nodes are then connected according to the super-node’s glyph: for out(in)-stars a node defined as the hub points to (is pointed by) all other nodes, for cliques all possible directed edges are added between the nodes, and for disconnected sets no edges are added. Moreover, a super-node self-loop expands to self-loops on every node it contains. On the other hand, super-edges expand to sets of edges that have the same direction. If the involved source/target glyphs are not stars, all the nodes contained in the source glyph point to all the nodes contained in the target glyph. For stars, the expanded incoming and out-going super-edges are only connected to the star’s hub. Finally, all expanded edges inherit their corresponding “parent” super-node’s or super-edge’s representative multiplicity.

Apart from enabling a clear interpretation of a given summary, the decompression rules help quantify how well the summary represents the original graph. For example, the pink triangle with representative multiplicity 3 in Fig. 1 (middle) expands to an in-star with all edges having multiplicity 3 as shown in Fig. 1 (right). While the topology is perfectly captured (pink nodes form a perfect in-star), the expanded multiplicities are not always equal to the original ones. On the other hand, expanding the green triangle perfectly captures the edge multiplicities (all are 1), but only approximates the topology, as the original green subgraph also contains some edges between the spokes of the hub node.

3.2 Model Encoding

Following the two-part Minimum Description Length paradigm [12], we aim to identify a summary graph \mathcal{G}_s that minimizes the total description cost of the full graph, that is,

$$\mathcal{G}_s^* := \arg \min_{\mathcal{G}_s} L(\mathcal{G}_s) + L(\mathcal{G}|\mathcal{G}'), \quad \text{s.t.} \quad \mathcal{G}' = \text{dec}(\mathcal{G}_s) \quad (1)$$

where $L(\mathcal{G}_s)$ measures the number of bits required to encode the summary graph, and $L(\mathcal{G}|\mathcal{G}')$ the bits needed to encode the corrections (or extra-information) for reconstructing the original graph \mathcal{G} from the (uniquely and unambiguously) decompressed \mathcal{G}' . These costs can be quantified as follows.

3.2.1 Encoding the summary graph We first encode the size of the summary graph $L_{\mathbb{N}}(|\mathcal{V}_s|)$, and the number of labels $L_{\mathbb{N}}(|\mathcal{T}|)$.³ For each super-node, $\log_2 |\mathcal{T}|$ bits are used to record its label, $\log_2 |\mathcal{M}|$ for its glyph, $L_{\mathbb{N}}(|\mathcal{S}_v|)$ for its size, $\mathcal{L}_{\mathbb{N}}(m(v))$ for the within-glyph representative multiplicity, $\log_2(|\mathcal{V}_s|)$ for the number of super-nodes in \mathcal{G}_s that it points to, and $\log_2 \binom{|\mathcal{V}_s|}{|\mathcal{N}(v)|}$ to identify the specific set of super-nodes it points to, where $\mathcal{N}(v)$ denotes the set of direct (out)neighbors. For each super-edge, $L_{\mathbb{N}}(m(e))$ bits are used for the representative multiplicity. In total, the number of bits required to encode a summary graph is given as

$$L(\mathcal{G}_s) = L_{\mathbb{N}}(|\mathcal{V}_s|) + L_{\mathbb{N}}(|\mathcal{T}|) + \sum_{v \in \mathcal{V}_s} L_{\text{SNODE}}(v), \quad \text{where} \quad (2)$$

$$L_{\text{SNODE}}(v) = \log_2 |\mathcal{T}| + \log_2 |\mathcal{M}| + L_{\mathbb{N}}(|\mathcal{S}_v|) + \mathcal{L}_{\mathbb{N}}(m(v)) + \log_2(|\mathcal{V}_s|) + \log_2 \binom{|\mathcal{V}_s|}{|\mathcal{N}(v)|} + \sum_{z \in \mathcal{N}(v)} L_{\mathbb{N}}(m(v, z)) \quad (3)$$

³ $L_{\mathbb{N}}(k) = 2 \log k + 1$ bits are required to encode an arbitrarily large natural number k , using the variable-length prefix-free encoding; see, Ex. 2.4 in [12].

3.2.2 Encoding the corrections For the overall cost for corrections, we first compute the number of bits used to correct the *topology* of the expanded (i.e., decompressed) graph, followed by the number of bits needed to represent the true *multiplicities*. Regarding the topology, we first map the expanded nodes back to the original node-set \mathcal{V} . This costs $L_{\text{MAP}}(v) = \log_2 \binom{|\mathcal{V}|}{|\mathcal{S}_v|} + \mathbb{1}_{\{\mu(v)=\text{STAR}\}} \log_2 |\mathcal{S}_v|$ bits per super-node v (the latter term identifying the hub of a star). Subsequently, we have two types of edge corrections: Either adding edges that exist in the original graph but not in the expanded graph (i.e., positive corrections) or removing edges from the expanded graph because they do not exist in the original graph (i.e., negative corrections).

These costs are compactly encoded for every expanded super-edge and every expanded super-node (glyph), using the binomial encoding $L(\mathcal{E}_{\text{COR}}) = L_{\mathbb{N}}(|\mathcal{E}_{\text{COR}}|) + \log_2 \binom{|\mathcal{E}_{\text{COR}}^{\text{max}}|}{|\mathcal{E}_{\text{COR}}|}$, where \mathcal{E}_{COR} denotes the possible set of corrections (positive or negative), and $\mathcal{E}_{\text{COR}}^{\text{max}}$ the largest set that \mathcal{E}_{COR} can possibly be. For example, for positive edge corrections in a disconnected set, we have $\mathcal{E}_{\text{COR}}^{\text{max}} = \mathcal{S}_v \times \mathcal{S}_v$, and similarly for negative edge corrections in a clique. For super-edges, corrections are computed according to the decompression rules (see Sec. 3.1). For the (few) edges in the original graph between super-nodes v and z that are not represented by a super-edge, the corrections are always positive, and $\mathcal{E}_{\text{COR}}^{\text{max}} = \mathcal{S}_v \times \mathcal{S}_z$.

The binomial encoding arises from using the uniform code over all the lexicographically ordered *subsets* of possible corrections. An alternative to this, as suggested in [17] and [24], would be to encode each correction *individually* using an optimal prefix code. Then, interpreting $p = |\mathcal{E}_{\text{COR}}|/|\mathcal{E}_{\text{COR}}^{\text{max}}|$ as the “probability” of each correction, we would need $L_{\text{ENTR}} = H(\text{Ber}[p]) \cdot |\mathcal{E}_{\text{COR}}^{\text{max}}|$ bits, where $H(\cdot)$ is the Shannon entropy, and $\text{Ber}[p]$ is a Bernoulli with parameter p . Denoting $|\mathcal{E}_{\text{COR}}| = n'$ and $|\mathcal{E}_{\text{COR}}^{\text{max}}| = n$, we can show that our binomial encoding is more efficient.

Theorem 1. *It holds that $L_{\text{ENTR}} \geq \log_2 \binom{n}{n'}$, for all $n > n' > 0$.*

Proof. See Appendix 7.

Theorem 1 establishes that the binomial encoding always gives a more compact measure of information required for corrections. Having corrected the edge topology, we compute the cost of correcting the edge multiplicities. Since any edge e not included in a glyph or super-edge does not have a representative multiplicity, its multiplicity correction is encoded by $L_{\mathbb{N}}(m(e))$, encoding its true value. The reason for using $L_{\mathbb{N}}(\cdot)$ to encode multiplicities is the fact that, for most real graphs, multiplicities follow a power-law distribution. Since the vast majority has small values, $L_{\mathbb{N}}(\cdot)$ will generally be a more “compact” encoding compared to a uniform code based on the maximum multiplicities. For expanded super-nodes and super-edges with representative multiplicity m , we obtain the cost of correcting the multiplicities as

$$L_{\text{DIFF}}(\mathcal{E}_{\text{sup}}, m) = \sum_{e \in \mathcal{E}_{\text{sup}}} \ell_{\text{diff}}(m(e), m), \quad (4)$$

where \mathcal{E}_{sup} in this context is the set of all edges contained in said super-node or super-edge, and

$$\ell_{\text{diff}}(m', m) = \begin{cases} 1, & m = m' \\ 2 \log_2(|m - m'|) + 3, & m \neq m' \end{cases}, \quad (5)$$

bits are needed to encode the *difference* between a true multiplicity m' and its representative m . Note that, since $L_{\mathbb{N}}(\cdot)$ only holds for natural numbers (see footnote³), one extra bit is required to indicate whether the difference is 0, and one more for the sign of the difference.

4 Graph Summary Search

The discrete optimization problem in (1) has a very large set of feasible solutions, and needs to be approximated efficiently. Towards this goal, we follow a two-step process, where we first generate a list of (possibly overlapping) groups of nodes, which we term *candidate* node-sets (see Sec. 4.1), and then decide which ones to merge into super-nodes. These candidates have varying size and quality (i.e., structural-similarity). Larger candidates with low quality compress the graph more (reduced $L(\mathcal{G}_s)$), but also typically require more corrections (increased $L(\mathcal{G}|\mathcal{G}')$). Clearly, the best candidates have both high quality and large size. For this reason, we first sort the candidate sets in descending order with respect to the product of their size and quality. We then process the sorted list from top to bottom, and merge the candidate sets into super-nodes, updating the summary graph accordingly (see Sec. 4.2). To ensure the quality of summarization, we only monitor the overall total cost, and only commit to a given candidate if $\Delta_{\text{cost}} = \text{Cost_After} - \text{Cost_Before} < 0$. This offers two benefits: (1) We avoid the cumbersome process of merging nodes in pairs (i.e. two at a time) and instead merge *in groups*, and (2) We achieve ability to summarize at *multiple resolutions*. The overview is given in Algorithm 1.

Algorithm 1: TG-SUM: Summarizing Labeled Multi-Graphs

Input: directed labeled multi-graph \mathcal{G}

- 1 Construct candidate node-sets (Sec. 4.1);
- 2 Sort candidates w.r.t. (size \times quality);
- 3 **for** every candidate set in list **do**
- 4 Merge unmarked nodes in set and decide glyph (Sec. 4.2.1);
- 5 Decide super-edges (Sec. 4.2.2) ;
- 6 Compute representative multiplicities (Sec. 4.2.3);
- 7 Mark candidate node-set as merged;
- 8 **if** $\Delta_{\text{cost}} < 0$ **then**
- 9 Commit to merged super-node and its super-edges;

10 **Return** summary graph \mathcal{G}_s ;

4.1 Candidate Set Generation

4.1.1 Measuring candidate quality To quantify a candidate set’s quality, we first need to define a proper metric of structural node similarity. For undirected graphs, the Jaccard similarity between two nodes v and v' is given as $J^U(v, v') = \frac{|\mathcal{N}^U(v) \cap \mathcal{N}^U(v')|}{|\mathcal{N}^U(v) \cup \mathcal{N}^U(v')|}$, and simply measures the proportion of common neighbors that they share. Naïvely using $J^U(\cdot, \cdot)$ on directed graphs is straightforward by ignoring the directions of the edges, however, it may yield misleading results by often over-estimating the true node similarity.

Consider the simple example depicted in Fig. 2, where any pair in the set $\{B, C, D, E, F\}$ has undirected similarity 1, while the fact that $\{B, C, D\}$ are incoming and $\{E, F\}$ are outgoing means that the set is far from an ideal candidate. To mitigate such inconsistencies, we introduce the following extension of Jaccard that may also accommodate directed graphs, by taking into account the similarity of both *Incoming* and *Outgoing* edges.

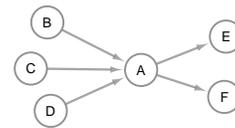


Fig. 2: Directed Jaccard similarity treats incoming (B, C, D) and outgoing (E, F) neighbors separately.

Definition 1. *The directed Jaccard similarity between any two pair of nodes v, v' of a directed graph is given as*

$$J^D(v, v') = \frac{|\mathcal{N}^I(v) \cap \mathcal{N}^I(v')| + |\mathcal{N}^O(v) \cap \mathcal{N}^O(v')|}{|\mathcal{N}^I(v) \cup \mathcal{N}^I(v')| + |\mathcal{N}^O(v) \cup \mathcal{N}^O(v')|} \quad (6)$$

First, it can easily be observed that for undirected graphs, $J^D(v, v') = J^U(v, v')$, since $\mathcal{N}^I(v) = \mathcal{N}^O(v) = \mathcal{N}^U(v)$. Note however, that in our example, $J^D(v, v')$ becomes 0 for all cross-pairs between $\{B, C, D\}$ and $\{E, F\}$, effectively creating two separate groups. In general for directed graphs, $J^D(v, v')$ will be more “informed” than $J^U(v, v')$, typically yielding lower similarity scores.

Having defined our similarity, we measure a candidate’s quality according to the minimum directed Jaccard among all pairs of nodes it contains. Thus, given the following definition,

Definition 2. *Any set $\mathcal{C} \subseteq \mathcal{V}$, is t -bounded if $J^D(v, v') \geq t \quad \forall (v, v') \in \mathcal{C} \times \mathcal{C}$.*

We use the t -bounded-ness of a candidate to serve as a pessimistic valuation of its quality. In addition, given that we are interested in a *collection* of candidate sets, we would like the sets to be *non-redundant* defined as follows.

Definition 3. *Let \mathcal{C}_S be a collection of candidate sets, each one accompanied by a bound t . We call \mathcal{C}_S non-redundant, if for any $\mathcal{C} \in \mathcal{C}_S$ that is t -bounded, there exists no t' -bounded $\mathcal{C}' \in \mathcal{C}_S$, such that $t' \geq t$ and $\mathcal{C} \subset \mathcal{C}'$.*

Simply put, non-redundancy ensures that none of the candidate sets is a strict subset of another set of higher or equal quality.

4.1.2 Incremental LSH To group nodes according to their similarity, we first utilize Locality Sensitive Hashing (LSH) [4]. Specifically for every node v , we generate a set of r *minhash* signatures

$$h_j(v) := \min_{z \in \mathcal{N}^D(v)} f_j(z) \quad \forall j = 1, \dots, r \quad (7)$$

where f_j ’s are independent and uniform hash functions (see, e.g., [4] for implementation details), and $\mathcal{N}^D(v) := \mathcal{N}^I(v) \parallel \mathcal{N}^O(v)$ is the concatenated adjacency list of node v that includes all incoming and outgoing neighbors separately. It can then be shown that $\Pr \{h_j(v) = h_j(v')\} = J^D(v, v')$; that is, two nodes share a minhash signature with probability proportional to their directed Jaccard similarity.

Since the r hash functions are independent, it follows that $\Pr \{\mathbf{h}(v) = \mathbf{h}(v')\} = (J^D(v, v'))^r$, where $\mathbf{h}(v) := [h_1(v), \dots, h_r(v)]^T$ is the r -length minhash signature vector of node v . If the nodes are hashed into buckets according to their r minhash signatures, the equality gives the probability that two nodes hash into the same bucket. By collecting b hash-tables corresponding to b bands of r minhash signatures, the probability that v and v' hash to the same bucket at least once is

$$\Pr \{\mathbf{h}_i(v) = \mathbf{h}_i(v') \exists i = 1, \dots, b\} = 1 - \left(1 - (J^D(v, v'))^r\right)^b \quad (8)$$

Interestingly, for sufficiently large r and b , the RHS expression in Eq. (8) when viewed as a function of $(J^D(v, v'))$ approximates a step function around the threshold $t = \left(\frac{1}{b}\right)^{\frac{1}{r}} \in (0, 1]$, meaning that with high probability v and v' will belong in a t -bounded set. To avoid repeating the

entire process for different values of b , we incrementally generate and add more bands of minhash node signatures, that in turn hash nodes into new buckets. The new buckets are then merged with any overlapping existing buckets, gradually coalescing into larger clusters that are *approximately* t -bounded, with $t = (\frac{1}{b})^{\frac{1}{r}}$ decreasing as b increases. This is exactly how we obtain larger candidate sets, albeit of lower quality, incrementally by the addition of new bands.

4.1.3 Filtered LSH While the incremental LSH described in the previous section efficiently guides the process of forming candidate sets, merged buckets are not guaranteed to be t -bounded due to the false alarm probability. For this purpose, we maintain an undirected similarity graph \mathcal{G}_{sim} , where an edge (v, v') is *guaranteed* to appear if and only if $J^D(v, v') \geq t$. Intuitively, \mathcal{G}_{sim} serves as a data structure where large t -bounded candidates appear as maximal cliques. As new LSH buckets appear and clusters are updated, we compute $J^D(v, v')$ for newly coalesced pairs of nodes (v, v') , and add the latter as an edge to \mathcal{G}_{sim} if $J^D(v, v') \geq t$. If the threshold is not satisfied, the computed value $J^D(v, v')$ is not discarded, but cached into a max-heap since it may satisfy a lower t in one of the subsequent iterations as b is increased. To further reduce the complexity of verifying pair similarities, we rely on the Jaccard upper bound $J^D(v, v') \leq \min(d, d') / \max(d, d')$, where $d = |\mathcal{N}^D(v)|$. If the upper bound of a pair (v, v') is smaller than the lowest threshold $t_{\min} = (\frac{1}{b_{\max}})^{\frac{1}{r}}$ that will ever be required, we can safely ignore (v, v') since it will never form a valid edge. Leveraging this property, we maintain nodes sorted in every cluster according to their in-plus-out-degrees d , and for a given v' in the array, we only consider similarities with v 's in the window $[d_v t_{\min}] \leq d_v \leq [d_v / t_{\min}]$.

As mentioned earlier, candidate sets are collected as maximal cliques in \mathcal{G}_{sim} . To ensure that the set of candidates is *non-redundant* (cf. Def.n 3), we maintain for every node the size of the maximum clique that it has been found to belong in. Every time a new clique is discovered, we update the maximum-sizes for all the nodes it contains using the clique's size. As new edges are added to \mathcal{G}_{sim} , we examine every node for newly emerged cliques, and we rely on the heuristic in [23] to prune the search by avoiding the evaluation of cliques that cannot exceed the size of the previously-found maximum clique. This way, we achieve two goals: 1) We ensure that the new cliques discovered at every iteration are not subsets of (or equal to) any previous clique, and 2) We reduce computation by early-stopping of the clique discovery process for most nodes.

4.2 Merging Candidates: Glyphs, Super-edges, Multiplicities

Every time a candidate set \mathcal{C} is tested, we deploy subroutines that efficiently update the summary graph, by making decisions regarding (1) the glyph that will be assigned to the merged set of nodes, (2) super-edges that emerge (or disappear) due to the merging, as well as (3) representative multiplicities for the set of edges summarized by the glyph and its super-edges.

4.2.1 Glyph decision rules To preserve super-node label homogeneity, a candidate set that contains nodes of different labels is first split into same-label subsets. Each subset is merged into a separate super-node using the procedure described below. Hereafter, the term candidate set refers to such a label-homogeneous subset.

For the glyph decision, we first identify the number of directed edges $E_{\mathcal{C}}$ that are included in the subgraph induced on nodes that corresponds to \mathcal{C} in the candidate set. Consequently, if $E_{\mathcal{C}} \geq |\mathcal{C}|(|\mathcal{C}| - 1)/2$, i.e., at least half of all possible directed edges are present, then we decide **Clique** since it most likely is the best glyph in terms of number of edge corrections. For sparsely-edged candidate sets that do not pass the clique threshold, we proceed to test for the presence of

stars. If there is a suitable out-/in-star present in \mathcal{C} , then its hub will be the highest out/in-degree node in \mathcal{C} . We therefore compute d_{\max}^I and d_{\max}^O , that is, the largest in and out degrees in the induced subgraph, and compute the following proxy correction cost for encoding an in-star

$$\text{Cost}_{\text{IN}} = (|\mathcal{C}| - 1 - d_{\max}^I) + (E_{\mathcal{C}} - d_{\max}^I), \quad (9)$$

and similarly Cost_{OUT} for out-star using d_{\max}^O . Intuitively, the first term of Eq. (9) is the number of edges that will have to be removed from the full decompressed star, while the second part is the number of edges that cannot be “explained” by the star and will have to be added. We then compare Cost_{IN} and Cost_{OUT} with $E_{\mathcal{C}}$, i.e., the number of edges that will have to be added if we decide that \mathcal{C} is a disconnected set. If only Cost_{IN} (or only Cost_{OUT}) is smaller than $E_{\mathcal{C}}$, then we decide **In-star** (or **Out-star**). If both Cost_{IN} and Cost_{OUT} are smaller than $E_{\mathcal{C}}$, then we choose the smallest of the two. Finally, if neither Cost_{IN} nor Cost_{OUT} are smaller than $E_{\mathcal{C}}$, we decide that \mathcal{C} is a **Disconnected** set.

4.2.2 Super-edge decision rule Having decided the glyph of \mathcal{C} , we merge any outgoing and incoming edges and/or super-edges into “bundles” of edges and their corresponding multiplicities. We then obtain the topology-based correction costs of merging or not merging each bundle into a super-edge (recall Sec. 3.2). Furthermore, we need to extract a representative multiplicity (see below) and compute the cost of encoding the corrections as the difference between the true and the representative multiplicity (recall Eq. (4)). If the total cost (topology and multiplicities) of representing each bundle of edges with a super-edge is lower than the cost of *not* representing it, then the corresponding super-edge (and its representative multiplicity) is added to the summary.

4.2.3 Finding representative multiplicities For every newly-formed super-node as well as each potential super-edge, we find the representative multiplicity m^* as

$$m^* := \arg \min_m L_{\text{DIFF}}(\mathcal{E}_{\text{sup}}, m), \quad (10)$$

where $L_{\text{DIFF}}(\mathcal{E}_{\text{sup}}, m)$ is defined in Eq. (4), and \mathcal{E}_{sup} is the set of all edges contained in a given super-node or bundled by a super-edge. Although this 1D-optimization problem is not convex, we find that the dichotomous search algorithm [8] finds the optimal solution in most cases, and runs in $\mathcal{O}(|\mathcal{E}_{\text{sup}}| \log_2 R)$ time, where $R = \max_{e \in \mathcal{E}_{\text{sup}}} m(e) - \min_{e \in \mathcal{E}_{\text{sup}}} m(e)$, i.e., the dynamic range of multiplicities.

4.3 Complexity Analysis

Time: First we analyze the complexity of the candidate generation algorithm. For a given LSH band of r rows, generating the minhash signatures is $\mathcal{O}(|\mathcal{E}|r)$, where \mathcal{E} is the set of all directed edges in \mathcal{G} . Further, filtering the merged buckets takes $\mathcal{O}(\bar{d}|C_{\max}|^2)$, where \bar{d} is the sum of average in- and out-degrees in \mathcal{G} , which is the average time required for computing the directed Jaccard for a given pair of nodes, and C_{\max} is the largest cluster that is created from merging buckets. For the discovery of emerging cliques $\mathcal{O}(\gamma|C_{\max}|^2)$ is required in the worst case, where γ is the fraction of nodes in the largest cluster that will not be pruned. Overall, using b bands, $\mathcal{O}(b[|\mathcal{E}|r + (\bar{d} + \gamma)|C_{\max}|^2])$ computations are required for the generation of candidate sets. During the merging phase, the complexity of processing a given candidate set is linear w.r.t. the edges adjacent to all the nodes in the set. Therefore, processing say, top K candidate sets is $\mathcal{O}(\bar{d}\bar{C}K)$, where \bar{C} is the average candidate set size. Note that the total complexity is linear in the number of edges in \mathcal{G} .

Space: The overall space complexity is dominated by the candidate set generation, specifically, by maintaining the similarity graph \mathcal{G}_{sim} together with its cached edges. Given C_{max} , $\mathcal{O}(|C_{\text{max}}|^2)$ memory would be required in the worst case.

5 Experiments

5.1 Setup

Datasets. We experimented with real-world graphs of a wide variety of sizes and characteristics, including the Moreno sheep⁴ interaction network; a senator-to-senator network extracted from the 2009-2010 US Congress dataset [3] and the Political Blogs network [1], both with political affiliation labels; DBLP 4-area⁵ co-authorship network with labels denoting author areas; the Cora and Citeseer citation networks [11] that are labeled by publication venue; and finally, transaction networks from 3 (anonymous) corporations that we collaborated with. We also used unlabeled protein-protein interaction network for Homo Sapiens [22]; Stanford CS web network⁶; and the Enron⁷ email network. See Table 2 for a summary of network characteristics.

Table 2: Real-world graphs used in experiments. * depicts naturally directed graphs that are typically treated as undirected. For SH, HW, KD #labels is given for EB/FS labeling.

Name	#nodes	#m-edges	#labels	Lbl.	Dir.	Mult.	S-loop
moreno	0.03K	0.25K	2	✓	✓	✓	
senate	0.1K	2.4K	2	✓			
polblogs	1.5K	19K	2	✓	*		
4-area	27.2K	67K	5	✓			
cora	2.7K	10.6K	6	✓	*		
citeseer	3.3K	9.2K	7	✓	*		
SH trans	0.25K	301K	11/27	✓	✓	✓	✓
HW trans	0.32K	268K	11/60	✓	✓	✓	✓
KD trans	2.3K	648K	10/29	✓	✓	✓	✓
PPI	3.9K	77K	-				
stanford	10K	74K	-		*		
enron	70K	277K	-		*		

Baselines and Configurations. There is *no existing method* for LMDS-graph summarization (cf. Table 1), thus we compare only under simplified settings, w.r.t. running time and compression rate. Moreover, TG-SUM is only comparable to lossless methods. On unlabeled simple graphs, we use VoG [16] as a baseline (others in Table 1 do not have public source code available), with default settings⁸. Further, we modify the *Randomized* algorithm of Navlakha *et al.* [21] to accommodate node labels and edge directions, and compare on all graphs, ignoring the edge multiplicities. All experiments were run on a PC with 9th generation i-7 processor and 128GB RAM.

5.2 Qualitative Evaluation: TG-SUM at Work

Moreno sheep. This is a typical complex interaction network between animals. Here, an edge (u, u') is recorded every time sheep u is observed to show dominating behavior over sheep u' .

⁴<http://moreno.ss.uci.edu/data.html>

⁵<https://dblp.uni-trier.de/>

⁶<https://sparse.tamu.edu/Gleich/wb-cs-stanford>

⁷<https://sparse.tamu.edu/LAW/enron>

⁸https://github.com/GemsLab/VoG_Graph_Summarization

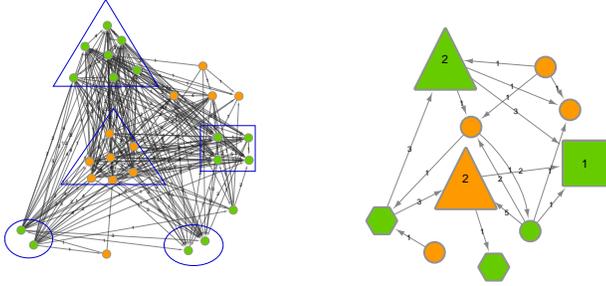


Fig. 3: Moreno sheep graph (left), summary graph (right).

the two out-stars reveal local hierarchies for old and young sheep correspondingly, where a “hub” sheep dominates the spokes. The summary also includes a clique of 4 young sheep that is dominated by most other groups, and two disconnected sets, each containing a couple of sheep that exhibit similar behavior (one couple being dominant, and the other dominated).

Co-authorship. In the 4-area coauthorship network, each author belongs to one of 4 areas (Databases, Info. Retrieval, Data Min., and Machine Learn.). We introduce a fifth category for authors that publish in more than one area. Analysis of the resulting glyphs (see Table 3) reveals different connectivity patterns among the areas. For instance, ML authors belong in cliques with higher frequency than those of other areas.

On the other hand, DB authors have a higher tendency to form disconnected groups that collaborate with one author. Of high interest is the case of the fifth category of interdisciplinary (mixed-area) authors, whom remain single (unmerged) at a significantly higher ratio than single-area authors. Visual inspection of different parts of the summary (see e.g. Fig. 4) reveals that the reason for mixed-area authors not being merged is because they frequently act as “bridges” that connect authors from two or more different areas. This showcases the important role of interdisciplinary researchers as the “glue” of the academic community.

Senate. The dataset contains the (positive or negative) votes of 108 senators for 696 congressional bills. The senators are labeled as Republicans (red), Democrats (blue), or independent (green). We construct an undirected graph where two senators are connected by an edge if the cosine similarity of their votes is larger than 0.3. The graph is plotted in Fig. 5, along with two summaries at different resolutions, leading to the following observations. Interestingly, while most democratic senators eventually form a clique, there is a smaller group of East coast senators, including prominent Democratic figures such as Joe Biden, Hillary Clinton, and Ted Kennedy that do not merge with the main body and form their own separate clique. Furthermore, this clique of Democrats is directly linked to certain Republicans, such as the Florida-based Mel Martinez, who has most recently opposed Trump openly and explicitly expressed his preference for Joe Biden⁹. The second observation is that Republican senators overall exhibit a more fragmented voting behavior, splitting into multiple cliques of comparable size. This

Ages of the sheep are available, and we use them to label the individuals as young (age ≤ 7 , green) or old (age > 7 , orange). Due to the relatively large number of recorded interactions, the full network in Fig. 3 (left) is difficult to interpret visually. The proposed TG-SUM can mitigate this problem by identifying the main structures, and building the summary graph in Fig. 3 (right). For this type of graph,

Table 3: Breakdown of 4-area summary graph

	DB	DM	IR	ML	mixed
Cliques ratio	43%	42%	38%	47%	36%
Disconnected ratio	40%	33%	38%	29%	18%
Single ratio	17%	25%	24%	24%	46%

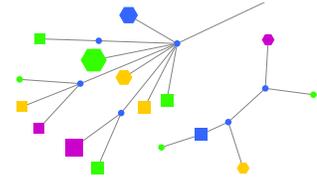


Fig. 4: 4-area coauthorship graph summary visualized partly. Mixed-area authors (blue) tend to “bridge” authors from a single area. (best in color)

⁹<https://theweek.com/speedreads/609406/former-rnc-chairman-says-wont-vote-trump-wishes-joe-biden-run>

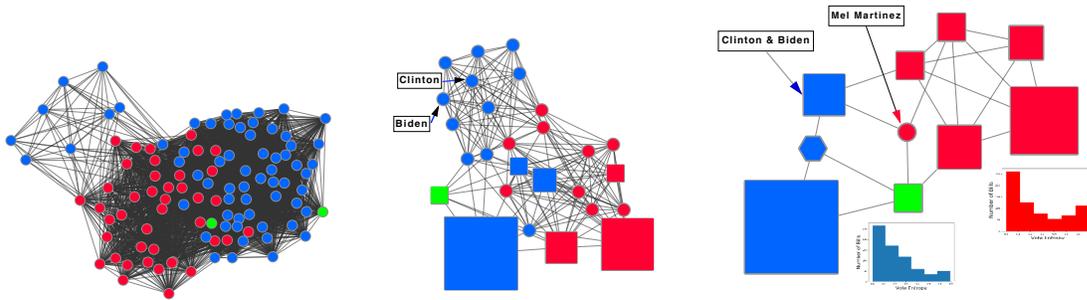


Fig. 5: (left) original US Senate graph, (middle) high resolution ($b = 2$) summary, (right) low resolution ($b = 5$) summary.

is corroborated by computing the entropies of the votes for all the bills, for Democrats and Republicans separately. Intuitively, bills with high entropy indicate a low degree of agreement on the subject. By plotting the histograms of the voting entropies (see Fig. 5 (right)) for the two groups, it becomes apparent that Republican votes exhibit higher entropy (median = 0.21) than Democrats (median = 0.16). Moreover, Republicans have more than 100 highly-contested bills (entropy ≈ 0.7), while the same for Democrats is less than 50.

5.3 Quantitative Evaluation: Evaluating Financial Accounts Labeling

In this section, we show how we employed TG-SUM to quantitatively address a domain-specific problem, specifically, evaluating a pre-existing *labeling*, i.e., the set of types pre-assigned to the nodes in an accounting network that connects business accounts via credit/debit transaction relations.

A business entity’s Chart of Accounts (COA) lists, and also pre-assigns a label to, each distinct account used in its ledgers. Such labeling helps companies prepare their aggregate financial statements (FS). For example, the FS caption “Cash and Cash Equivalents” is used to describe the total sum of all liquid assets tracked in a number of accounts; e.g., currencies, checking accounts, etc. In the US, FS captions are not uniform across corporations. In fact, the data we have from 3 different companies (anonymized as SH, HW, and KD in Table 2) each contains different FS captions.

How suitable is a given FS labeling?¹⁰ Can a different labeling be shown to be quantitatively better than another?

To this end, our collaborator (an accounting expert) designed a new labeling (referred as EB for economic bookkeeping), relabeling the accounts based on their primary economic nature. Specifically, EB organizes them into operating versus financing and long- versus short-term accounts. Expert knowledge suggests that EB improves over FS captions in a couple of ways. First, EB splits the accounts under a generic FS caption like “inter-company accounts” to reflect the financial or operating nature of within-firm activities. EB also consolidates excessive “over-labeling” in FS, combining a few economically similar FS captions into a single label; e.g., combining “accounts payable” and “accrued expenses” into “short-term operating liabilities”.

Transaction networks. Ideally, accounts of the same label should “behave” similarly in the system. This behavior can be discerned from the real-world usage data, in particular the transactions graph,

¹⁰Evaluating COA is beneficial, as it can become stale when corporations change due to market conditions or acquisitions and mergers. It can also assess any proposed new COA labeling on past data before implementation of the new system.

where accounts are connected through credit/debit relations. Under a more suitable labeling, the accounts with the same label should have more structural similarity and yield better compression.

Table 4: Evaluating account labelings in financial networks

Dataset	Labeling	Shuffled	Actual	norm. gain (%)
SH	EB	0.28	0.32	5.6 %
	FS	0.25	0.27	2.7 %
HW	EB	0.36	0.47	17.0 %
	FS	0.16	0.27	13.0 %
KD	EB	0.33	0.42	13.7 %
	FS	0.31	0.39	12.0 %

To compare EB vs. FS, we employ TG-SUM on each graph using one or the other labeling separately, and record the compression rate. Next, we shuffle the labels (within each setting) randomly, and employ TG-SUM again. “Shuffled” and “Actual” compression rates are reported in Table 4 (the former averaged over 20 random shuffles). EB rates are higher—this is not surprising as EB has fewer labels as compared to FS (See $|\mathcal{T}|$ in Table 2), and hence TG-SUM has higher degree of freedom to merge nodes on EB-labeled graphs. As such, Actual values are not directly comparable. What is comparable is the difference from Shuffled, that is, how much the labeling can improve on top of the random assignment of the *same* set of labels. Here, the absolute difference is always equal or larger for EB.

However, even the absolute difference of compression rates is not fair to compare—it is harder to compress a graph that has been compressed quite a bit even further. For EB, Shuffled rates are already high. Improving over Shuffled even by the same amount proves EB superior to FS. Therefore, we report the normalized gain; defined as $(\text{Actual} - \text{Shuffled}) / (1 - \text{Shuffled})$, which shows that our expert-designed EB labeling is better, for the aforementioned reasons. All in all, TG-SUM can be employed as a data-driven tool to quantitatively assess the connection between a labeling and the actual role of nodes in practice as reflected in the data.

5.4 Quantitative Evaluation: Compression rate, Running time, Scalability

Quantitatively, we measure summarization performance in terms of both (1) running time, as well as (2) the size reduction achieved in terms of bits (including bits required for correction). Specifically, upon obtaining the total number of bits (as given by the encoding scheme of each method), we measure the compression ratio as $\text{Compress Ratio} = \frac{\text{Bits_Before} - \text{Bits_After}}{\text{Bits_Before}} \in [0, 1)$, that is the fraction of the encoding cost that has been reduced by summarization. We compare with the Navlakha algorithm [21], which we modified to handle edge directions and node labels. We also compare against VoG only on unlabeled simple (i.e., unweighted) graphs (see Table 1), where cost is also measured in bits [16]. We run TG-SUM by gradually increasing b , to increase the number of candidate sets and obtain multi-resolution summaries. A larger number of candidates is expected to yield higher compression ratio, albeit at the cost of increased running time—hence enabling the user to choose a suitable trade-off in practice.

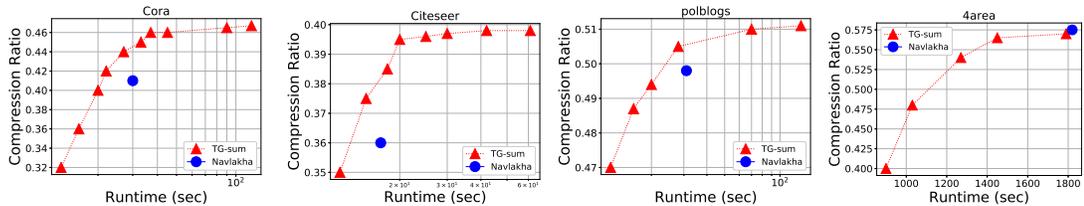


Fig. 6: Compression ratio vs. runtime on undirected graphs

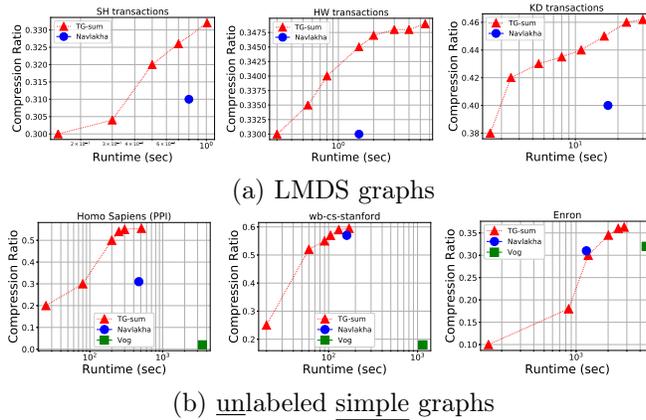


Fig. 7: Compression ratio vs. runtime

creasing size synthetic directed k -out graph, where nodes are incrementally added and connected to $k = 10$ of the existing nodes, simulating a preferential attachment process [13].

Results in Fig. 8 (left) show that, unlike the modified Navlakha, the runtime of TG-SUM grows in a near-linear fashion. In addition, we sample increasing percentage of the edges uniformly at random from the 4-area graph to quantify scalability as a function of number of edges where the node set is fixed. As Fig. 8 (right) shows, TG-SUM is scalable to large graphs, with runtime linear in the number of edges (w/ slope ≈ 1), empirically confirming the complexity analysis in Sec. 4.3.

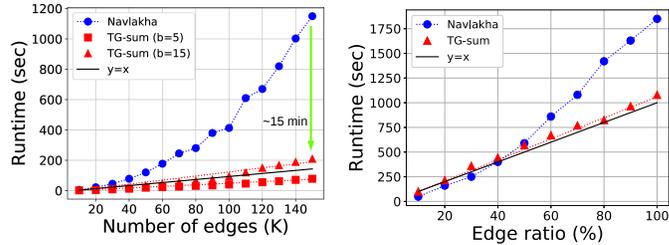


Fig. 8: TG-SUM scales linearly; (left) simulated graphs, (right) 4-area with increasingly sampled edges.

6 Conclusion

We introduced TG-SUM, a versatile graph summarization algorithm that (for the first time) can handle *directed*, *node-labeled*, *multi*-graphs with possible self-loops (or *any* combination). Built on a novel encoding scheme, TG-SUM seeks to minimize the total encoding cost of (i) a summary graph, and (ii) the corrections to reconstruct the input graph losslessly. It efficiently finds structurally-similar nodes to create super-nodes of larger sizes incrementally, producing multi-resolution summaries. Extensive experiments show that TG-SUM (1) provides insights into the high-level structure of real-world graphs, (2) achieves better trade-off between compression and runtime relative to baselines (only) on comparable settings, and (3) scales linearly in the number of edges.

Acknowledgements This work has been sponsored by the U.S. National Science Foundation CAREER 1452425 and the PwC Risk and Regulatory Services Innovation Center at Carnegie Mellon University. Any conclusions expressed in this material are those of the author and do not necessarily reflect the views, expressed or implied, of the funding parties.

Results are given in Fig. 6 (for labeled *undirected* graphs), in Fig. 7(a) (for LMDS graphs), and in Fig. 7(b) (for *unlabeled simple* graphs). TG-SUM remarkably outperforms the alternatives in terms of compression ratio in almost all cases. In other cases, it achieves comparable compression within similar or even lower running time. In absolute terms, it achieves roughly 30–60% compression across these various real-world graphs with up to hundreds of thousands of edges and tens of distinct labels.

Finally, we measure the scalability of TG-SUM by first generating an

References

1. Adamic, L.A., Glance, N.: The political blogosphere and the 2004 us election: divided they blog. In: Proceedings of the 3rd Intl. Workshop on Link Discovery. pp. 36–43 (2005)
2. Adler, M., Mitzenmacher, M.: Towards compressing web graphs. In: DCC. IEEE Comp. Soc. (2001)
3. Akoglu, L.: Quantifying political polarity based on bipartite opinion networks. In: ICWSM (2014)
4. Andoni, A., Indyk, P.: Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Commun. ACM* pp. 117–122 (2008)
5. Besta, M., Hoefler, T.: Survey and taxonomy of lossless graph compression and space-efficient graph representations (2018), arxiv:1806.01799
6. Boldi, P., Vigna, S.: The webgraph framework i: compression techniques. In: WWW. pp. 595–602 (2004)
7. Buehrer, G., Chellapilla, K.: A scalable pattern mining approach to web graph compression with communities. In: WSDM. pp. 95–106. ACM (2008)
8. Chong, E.K., Zak, S.H.: An introduction to optimization. John Wiley & Sons (2004)
9. Cook, D.J., Holder, L.B.: Substructure discovery using minimum description length and background knowledge. *Journal of Artificial Intelligence Research* **1**, 231–255 (1993)
10. Dunne, C., Shneiderman, B.: Motif simplification: improving network visualization readability with fan, connector, and clique glyphs. In: SIGCHI. pp. 3247–3256 (2013)
11. Giles, C.L., Bollacker, K.D., Lawrence, S.: Citeseer: An automatic citation indexing system. In: Proc. of the Conf. on Digital Libraries. pp. 89–98 (1998)
12. Grünwald, P.D.: The Minimum Description Length Principle. The MIT Press, Cambridge, MA (2007)
13. Jeong, H., Néda, Z., Barabási, A.L.: Measuring preferential attachment in evolving networks. *EPL* **61**(4), 567 (2003)
14. Kang, U., Faloutsos, C.: Beyond ‘caveman communities’: Hubs and spokes for graph compression and mining. In: ICDM. pp. 300–309 (2011)
15. Khan, K.U., Nawaz, W., Lee, Y.K.: Set-based approximate approach for lossless graph summarization. *Computing* **97**(12) (2015)
16. Koutra, D., Kang, U., Vreeken, J., Faloutsos, C.: Summarizing and understanding large graphs. *Statistical Analysis and Data Mining* **8**(3), 183–202 (2015)
17. LeFevre, K., Terzi, E.: Grass: Graph structure summarization. In: SDM. pp. 454–465. SIAM (2010)
18. Liakos, P., Papakonstantinou, K., Sioutis, M.: Pushing the envelope in graph compression. In: CIKM. pp. 1549–1558 (2014)
19. Liu, Y., Safavi, T., Dighe, A., Koutra, D.: Graph summarization methods and applications: A survey. *ACM Comput. Surv.* **51**(3) (2018)
20. Liu, Z., Yu, J.X., Cheng, H.: Approximate homogeneous graph summarization. *Info. and Media Tech.* **7**(1), 32–43 (2012)
21. Navlakha, S., Rastogi, R., Shrivastava, N.: Graph summarization with bounded error. In: SIGMOD. pp. 419–432. ACM (2008)
22. Oughtred, R., Stark, C., Breikreutz, B.J., Rust, J., Boucher, L., Chang, C., Kolas, N., O’Donnell, L., Leung, G., McAdam, R., et al.: The biogrid interaction database **47**(D1), D529–D541 (2019)
23. Pattabiraman, B., Patwary, M.M.A., Gebremedhin, A.H., Liao, W.k., Choudhary, A.: Fast algorithms for the maximum clique problem on massive sparse graphs. In: Intl Wksp. AMWG (2013)
24. Riondato, M., García-Soriano, D., Bonchi, F.: Graph summarization with quality guarantees. *Data Mining and Knowledge Discovery* **31**(2), 314–349 (2017)
25. Shin, K., Ghoting, A., Kim, M., Raghavan, H.: Sweg: Lossless and lossy summarization of web-scale graphs. In: WWW. pp. 1679–1690. ACM (2019)
26. Tian, Y., Hankins, R., Patel, J.: Efficient aggregation for graph summarization. In: SIGMOD (2008)
27. Toivonen, H., Zhou, F., Hartikainen, A., Hinkka, A.: Compression of weighted graphs. In: KDD (2011)
28. Ying, Z., You, J., Morris, C., Ren, X., Hamilton, W., Leskovec, J.: Hierarchical graph representation learning with differentiable pooling. In: NeurIPS. pp. 4800–4810 (2018)
29. Zhu, L., Ghasemi-Gol, M., Szekely, P., Galstyan, A., Knoblock, C.A.: Unsupervised entity resolution on multi-type graphs. In: Intl Semantic Web Conf. pp. 649–667 (2016)

7 Appendix

Proof of Theorem 1

First we expand

$$\begin{aligned}
L_{\text{ENTR}} &= n(-p \log_2 p - (1-p) \log_2(1-p)) \\
&= n' \log_2(n/n') + (n-n') \log_2(n/(n-n')) \\
&= \log_2 \frac{n^{n'}}{n'^{n'}} + \log_2 \frac{n^{(n-n')}}{(n-n')^{(n-n')}} \\
&= \log_2 \frac{n^n}{n'^{n'}(n-n')^{(n-n')}}.
\end{aligned}$$

Thus, to prove the inequality in Theorem 1, it suffices to show that

$$\binom{n}{n'} \leq \frac{n^n}{n'^{n'}(n-n')^{(n-n')}}.$$

Using Stirling's upper and lower bound on the factorial

$$\sqrt{2\pi n} n^{n+1/2} e^{-n} \geq n! \geq e n^{n+1/2} e^{-n},$$

we can bound

$$\begin{aligned}
\binom{n}{n'} &= \frac{n!}{n'!(n-n')!} \\
&\leq \frac{\sqrt{2\pi n} n^n \sqrt{n} e^{-n}}{e n'^{n'} \sqrt{n'} e^{-n'} e (n-n')^{n-n'} \sqrt{(n-n')} e^{-(n-n')}} \\
&= \frac{n^n}{n'^{n'}(n-n')^{(n-n')}} \cdot \frac{\sqrt{2\pi} \sqrt{n}}{e^2 \sqrt{n'}(n-n')} \\
&\leq \frac{n^n}{n'^{n'}(n-n')^{(n-n')}}
\end{aligned}$$

since

$$\frac{\sqrt{2\pi} \sqrt{n}}{e^2 \sqrt{n'}(n-n')} < \frac{\sqrt{n}}{e \sqrt{n'}(n-n')} \leq \frac{\sqrt{n}}{e \sqrt{n-1}} \leq 1,$$

for all $n > n' > 0$, which concludes the proof. \square