

Computing Graph Descriptors on Edge Streams

ZOHAIR RAZA HASSAN, Rochester Institute of Technology, USA

SARWAN ALI, Georgia State University, USA

IMDADULLAH KHAN, Lahore University of Management Sciences, Pakistan

MUDASSIR SHABBIR, Information Technology University, Pakistan

WASEEM ABBAS, University of Texas at Dallas, USA

Feature extraction is an essential task in graph analytics. These feature vectors, called graph descriptors, are used in downstream vector-space-based graph analysis models. This idea has proved fruitful in the past, with spectral-based graph descriptors providing state-of-the-art classification accuracy. However, known algorithms to compute meaningful descriptors do not scale to large graphs since: (1) they require storing the entire graph in memory, and (2) the end-user has no control over the algorithm's runtime. In this paper, we present streaming algorithms to approximately compute three different graph descriptors capturing the essential structure of graphs. Operating on edge streams allows us to avoid storing the entire graph in memory, and controlling the sample size enables us to keep the runtime of our algorithms within desired bounds. We demonstrate the efficacy of the proposed descriptors by analyzing the approximation error and classification accuracy. Our scalable algorithms compute descriptors of graphs with millions of edges within minutes. Moreover, these descriptors yield predictive accuracy comparable to the state-of-the-art methods but can be computed using only 25% as much memory.

CCS Concepts: • **Computing methodologies** → **Classification and regression trees; Knowledge representation and reasoning**; Distributed algorithms.

Additional Key Words and Phrases: Graph Descriptor, Edge Stream, Graph Classification

ACM Reference Format:

Zohair Raza Hassan, Sarwan Ali, Imdadullah Khan, Mudassir Shabbir, and Waseem Abbas. 2023. Computing Graph Descriptors on Edge Streams. *ACM Trans. Knowl. Discov. Data.* 1, 1, Article 1 (January 2023), 26 pages. <https://doi.org/00.0000/0000000>

1 INTRODUCTION

Graph analysis has a wide array of applications in various domains, from classifying chemicals based on their carcinogenicity [23] to determining the community structure in a friendship network [53] and even detecting discontinuities within instant messaging interactions [9]. The fundamental building block for analysis is a pairwise similarity (or distance) measure between graphs. However, efficient computation of such a measure is challenging: even the best-known solution for determining whether a pair of graphs are isomorphic has a quasi-polynomial runtime. Similarly, computing

Authors' addresses: Zohair Raza Hassan, zh5337@rit.edu, Rochester Institute of Technology, Rochester, USA; Sarwan Ali, sali85@student.gsu.edu, Georgia State University, Atlanta, USA; Imdadullah Khan, imdad.khan@lums.edu.pk, Lahore University of Management Sciences, Lahore, Pakistan; Mudassir Shabbir, mudassir.shabbir@itu.edu.pk, Information Technology University, Lahore, Pakistan; Waseem Abbas, waseem.abbas@utdallas.edu, University of Texas at Dallas, Dallas, USA.

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 ACM 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.

© 2023 Association for Computing Machinery.

1556-4681/2023/1-ART1 \$00.00

<https://doi.org/00.0000/0000000>

Graph Edit Distance [39], the minimum number of node/edge addition/deletions to interchange between two graphs is NP-HARD.

A relatively pragmatic approach is constructing fixed dimensional descriptors (vector embeddings) for graphs, allowing classical data mining algorithms that operate on vector spaces. Existing models using this approach can be categorized into (1) supervised models, which use deep learning methods to construct vector embeddings based on optimizing a given objective function [24, 32, 52] and (2) unsupervised models, which are based on graph-theoretic properties such as degree [18, 50], the Laplacian eigenspectrum [25], or the distribution of a fixed number of subgraphs [3, 16, 37, 41–43].

Unsupervised models construct general-purpose descriptors and do not require prior training on datasets. This approach has yielded great success; for example, descriptors based on spectral features (i.e., the graph’s Laplacian) provide excellent results on benchmark graph classification datasets [48, 50]. The order (number of vertices) and size (number of edges) of the graph and the number and nature of features computed directly determine the runtime and memory costs of the methods. By computing more statistics, one can construct more expressive descriptors. However, this approach does not scale well to real-world graphs due to their growing magnitudes [26].

Instead of storing and processing the entire graph, processing graphs as streams—one edge at a time—is a viable approach for limited memory settings [4]. The features are approximated from a representative sample of fixed size. This approach of trading-off accuracy for time and space complexity has yielded promising results on various graph analysis tasks such as graphlet counting [15], butterfly counting [38, 44], and triangle counting [45, 46]; despite storing a fraction of edges, these models have produced unbiased estimates with reasonably low error rates. Based on the success of these methods, our descriptors are designed to compute graph representations from edge streams, allowing us to compute features without storing the entire graph. In contrast, all existing descriptors and representation paradigms require storing the entire graph in memory.

This work is an extension of [22], wherein we proposed descriptors based on features obtained from graph streams. These descriptors are inspired by two existing works, the GRAPHLET KERNEL [43] and NETSIMILE [9], which compute local graph statistics as features. In this paper, we propose a new descriptor based on NETLSD [48] along with the proofs and experiments showcasing the said descriptor’s correctness and efficacy. We perform experiments on new benchmark datasets and provide data visualization of our proposed and NETLSD based embeddings using t -SNE.

Our contributions are summarized as follows:

- We propose simple graph descriptors that run on edge streams.
- We provide proofs to show how the features used in NETSIMILE [9] and NETLSD [48] can be computed using subgraph counts.
- We restrict our algorithms’ time and space complexity to scale linearly (for a fixed budget) in the order and size of the graph. We provide theoretical bounds on the time and space complexity of our algorithms.
- Empirical evaluation on benchmark graph classification datasets demonstrates that our descriptors are comparable to other state-of-the-art descriptors with respect to classification accuracy. Moreover, our descriptors can scale to graphs with millions of nodes and edges because we do not require storing the entire graph in memory.
- We perform data visualization to show the (global) distribution of data points in the proposed and state-of-the-art (SOTA) descriptors. The visualization results show that the SANTA preserves the data distribution better GABE and MAEVE and is comparable to the SOTA descriptor, NETLSD.

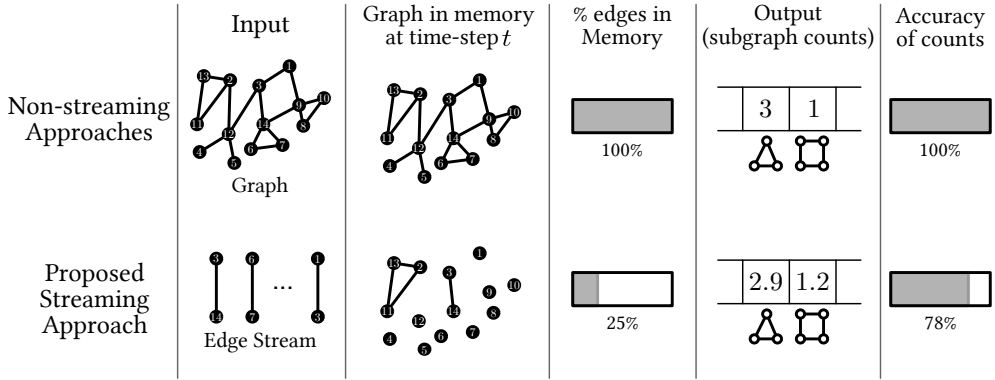


Fig. 1. This figure depicts the contrast between the typical approach for computing descriptors and our proposed approach. The descriptor in this example represents a graph by the counts of select subgraphs. Note how we tradeoff accuracy for memory consumption by keeping only a fraction of the graph in memory.

The remaining paper is organized as follows. We review some of the related work in Section 2 and give a formal problem description in Section 3. We provide detail of our descriptors in Section 4. Section 5 contains the experimental evaluation detail, including dataset statistics, preprocessing, hyperparameter values, and data visualization. In Section 6 we report the experimental results of our method. Finally, we conclude the paper in Section 7.

2 RELATED WORK

In this section, we review some closely related work on graph analysis. We discuss some distance/similar measures between graphs that are used in downstream machine learning algorithms. We also provide an overview of the basic paradigms for graph representation learning.

2.1 Pairwise Proximity Measure between Graphs

A fundamental building block for analyzing large graphs is evaluating pairwise similarity/distance between graphs. The *direct approach* to computing pairwise proximity considers the entire structure of both graphs. A simple and best-known distance measure between graphs is the *Graph Edit Distance* (GED) [39]. GED, like edit distance between sequences, counts the number of insertions, deletions, and substitutions of vertices and/or edges that are needed to transform one graph to the other. Runtimes of computing GED between two graphs are computationally prohibitive, restricting its applicability to graphs of very small orders and sizes. Another distance measure is based on permutations of vertices of one graph such that an error norm between the adjacency matrices of two graphs is minimum. Computing this distance and even relaxation of this distance is computationally expensive [7, 8]. When there is a valid bijection between vertices of the two graphs, then a similar measure, DELTACON [26], yields excellent results. However, requiring a valid bijection limits the applicability of DELTACON only to a collection of graphs on the same vertex set.

The representation learning approach for graph analysis maps graphs into a vector space. Vector space machine learning algorithms are employed using a pairwise distance measure between the vector representations of graphs. We discuss three broad approaches in this vein.

2.2 Kernel-Based Machine Learning Methods

The *kernel-based* machine learning methods represent each non-vector data item to a high dimensional vector. The feature vectors are based on counts (spectra) of all possible sub-structures of some fixed magnitude in the data item. A kernel function is then defined, usually as the dot-product of the pair of feature vectors. The pairwise kernel values between objects constitute a positive semi-definite matrix and serve as a similarity measure in the machine learning algorithm (e.g., SVM and kernel PCA). Explicit construction of feature vectors is computationally costly due to their large dimensionality. Therefore, in the so-called *kernel trick*, kernel values are directly evaluated based on objects. Kernel methods have yielded great successes for a variety of data such as images and sequences [10, 19, 27]. The most prominent graph kernels are the shortest-Path [11], Graphlet [43], the Weisfeller-Lehman [42], and the hierarchical [25] kernels. The computational and space complexity of the kernel matrix make kernel-based methods infeasible for large datasets of massive graphs.

2.3 Deep Learning Based Methods

The deep learning approach to representation learning is to train a *neural network* for embedding objects into Euclidean space. The goal here is to map ‘similar’ objects to ‘close-by’ points in \mathbb{R}^d . Deep learning-based methods and domain-specific techniques have been successfully used for embedding nodes in networks [13, 17, 21] and graphs [32, 52, 54, 55]. Vector-space-based machine learning methods are then employed on these embeddings for data analysis. However, these approaches are data-hungry and computationally prohibitive [40], hindering their scalability to graphs of large orders and sizes.

2.4 Descriptor Computation Methods

The *descriptor* learning paradigm differs from kernel methods in that the dimensionality of the feature vectors is much smaller than the kernel-based features. Unlike neural network-based models, the features are explainable and hand-picked using domain-specific knowledge [6, 9]. One such graph descriptor, NETSIMILE [9], represents a graph by a vector of aggregates of various vertex-level features. It considers seven features for each vertex, such as degree, clustering coefficient, and parameters of vertices’ neighbors and their “ego-networks,” and applies the aggregator functions, such as median, mean, standard deviation, skewness, and kurtosis, across each feature. Stochastic Graphlet Embedding [18] proposes a graph descriptor based on random walks over graphs to extract graphlets (sub-structures) of increasing order. Similar to this sub-structural approach is the Higher Order Structure Descriptor [3], which iteratively compresses graphlets within a graph to generate “higher-order” graphs and constructs histograms of the graphlet counts in each graph. More recently, FEATHER was introduced as a descriptor that computes node-level feature vectors using a complex characteristic function and aggregates these to construct graph embeddings [36]. There has been a trend towards using graph spectra [1, 2, 47] to learn descriptors [48, 50]. These descriptors are relatively computationally expensive but have excellent classification performance. An exact method, Von Neumann Graph Entropy (VNGE) is proposed in [12, 14] for graph comparison. Being an exact method, VNGE does not scale to large graphs. An approximate solution of NetLSD and VNGE, called SLAQ [49], computes spectral distances between graphs with multi-billion nodes and edges. Although computationally efficient, SLAQ keeps the entire graph in the memory during the processing, making it costly in terms of space efficiency.

Most of the above approaches require multiple passes over the entire input graph. The resulting space complexity renders them applicable only to graphs of small orders and sizes. On the other hand, real-world graphs are dynamic and enormous in their magnitude. Algorithms that perform a

single pass over the input stream and have low memory requirements [5] are best suited for modern-day graphs. An algorithm that computes the output with provable approximation guarantees is sufficient for the single-pass and sub-linear memory requirements. There have been few recent algorithms for counting specific substructures in a streamed graph owing to the inherent difficulty of the streaming model. These include approximately computing the number of triangles [46] in graphs, induced subgraphs of order three and four [15] in graphs, and cycles of length four in bipartite graphs [38].

3 PRELIMINARIES

In this section, we discuss the necessary prerequisites required to follow our work. We describe the graph nomenclature, followed by the description of graph descriptors, streams, and constraints imposed on our algorithm.

3.1 Graph Nomenclature

This section gives relevant notation and terminology for the rest of the paper, followed by a precise formulation of our main problem. A description of the notations for common terms used in this paper has been provided in Table 1. Notation tables specific to each descriptor have been provided in their sections.

Let $G = (V_G, E_G)$ be an undirected graph, where V_G is the set of vertices and E_G is the set of edges. We denote vertices of G by integers in the range $[0, |V_G| - 1]$. We refer to $|V_G|$ and $|E_G|$ as the order and size of G , respectively. In this paper, we only consider simple graphs (i.e., graphs with no self-loops and multi-edges) and unweighted graphs. For a vertex $v \in V_G$, we denote by $N_G(v)$, the set of neighbors of v , i.e., the set of vertices that are adjacent to v . More formally, $N_G(v) = \{u : (u, v) \in E_G\}$. The degree of a vertex v is denoted by d_G^v , i.e., $d_G^v := |N_G(v)|$. A pair of vertices $u, v \in V_G$ are said to be connected if there is a path between u and v , i.e., there exists a sequence of vertices $u = v_1, \dots, v_k = v$, where for $1 \leq i \leq k - 1$, $(v_i, v_{i+1}) \in E_G$. The length of a path is the number of vertices in it. A graph is called connected iff all pairs of vertices in V_G are connected.

A graph $G' = (V_{G'}, E_{G'})$ is called a subgraph of $G = (V_G, E_G)$ if $V_{G'} \subseteq V_G$ and $E_{G'} \subseteq E_G$ such that edges in $E_{G'}$ are incident only on the vertices present in $V_{G'}$, i.e., $E_{G'} \subseteq \{(u, v) : (u, v) \in E_G \wedge u, v \in V_{G'}\}$. If all edges incident on vertices in $V_{G'}$ are in $E_{G'}$ ($E_{G'} = \{(u, v) : (u, v) \in E_G \wedge u, v \in V_{G'}\}$), then G' is called an induced subgraph of G .

When vertices of a graph G_1 can be relabelled in such a way that we get another graph G_2 , then we say that G_1 and G_2 are isomorphic. In other words, G_1 and G_2 are isomorphic iff there exists a bijection $\pi : V_{G_1} \rightarrow V_{G_2}$ such that $E_{G_2} = \{(\pi(u), \pi(v)) : (u, v) \in E_{G_1}\}$. For a graph $F = (V_F, E_F)$, let H_G^F (resp., \widehat{H}_G^F) be the set of subgraphs (resp., induced subgraphs) of G that are isomorphic to F .

3.2 Graph Descriptors and Streams

A graph descriptor is a mapping from the family of all possible graphs (undirected, unweighted and simple, in our case) to a set of d -dimensional real vectors. More formally, let \mathcal{G} be the set of all possible graphs. A descriptor φ is a function, $\varphi : \mathcal{G} \rightarrow \mathbb{R}^d$. The primary motivation for using descriptors for graph analysis is to map graphs (possibly of varying sizes and orders) into a fixed-dimensional vector space, independent of the representation of graphs [9, 48]. A direct comparison of the number of certain subgraphs in two graphs of different orders and/or sizes is not very meaningful, as larger graphs will naturally have more subgraphs. Moreover, descriptors enable the application of vector-space-based machine learning algorithms for graph analysis tasks,

Notation	Description
G, F, G'	Common terms for graphs
V_G, E_G	Vertex and edge set for a graph G
$N_G(v), d_G^v$	Neighborhood and degree of a vertex $v \in V_G$
S, e_t	A stream of edges, and the edge arriving at time-step t
H_G^F	Set of subgraphs of G isomorphic to F
\widehat{H}_G^F	Set of induced subgraphs of G isomorphic to F
N_G^F	Estimate of $ H_G^F $
p_t^F	Probability of detecting F in \widetilde{E}_G , at the t^{th} time-step

Table 1. Notation table for common terms used throughout the paper.

often using the ℓ_2 -distance (Euclidean distance) as the proximity measure. Our descriptors are graph-theoretic and apply to graphs of varying magnitudes.

Let $S = e_1, e_2, \dots, e_{|E_G|}$ be a sequence of edges in a fixed order, i.e., $e_t = (u_t, v_t)$ is the t^{th} edge. We assume an online setting wherein the input graph is modeled as a stream of edges, i.e., we assume that elements of S are input to the algorithm one at a time. The following constraints are imposed on our algorithms:

- C1 **Constant Number of Passes:** The algorithm must do processing in a constant number of passes over the graph stream. Our algorithms require two passes at most.
- C2 **Limited Space:** The algorithm can store at most b edges during the execution. We refer to b as the budget and \widetilde{E}_G as the sample.
- C3 **Linear Complexity:** The time and space complexity of the algorithms must be linear in the order and size of the graph, with fixed b .

3.3 Estimating Connected Subgraph Counts on Edge Streams

In this section, we formally define the subgraph estimation problem within our constraints and describe the solution to this problem used throughout our proposed descriptors.

Problem 1 (Connected Subgraph Estimation on Edge Streams). *Let S be a stream of edges, $e_1, e_2, \dots, e_{|E_G|}$ for some graph $G = (V_G, E_G)$. Let $F = (V_F, E_F)$ be a connected graph such that $|V_F| \ll |V_G|$ (i.e. F is significantly smaller than G). Compute an estimate, N_G^F , of $|H_G^F|$ while storing at most b edges at any given instant.*

The basic strategy for solving Problem 1 involves two things: (1) an algorithm that counts the number of instances of a subgraph F that an edge belongs to, and (2) a sampling scheme that allows us to compute the probability of detecting an instance of F in our sample, denoted by p_t^F , at the arrival of the t^{th} edge [15, 45, 46]. The basic streaming algorithm maintains a representative sample of edges from the stream, and for each next edge e_t , it estimates the number of subgraphs in the sample containing the edge e_t . This estimate is scaled according to the sample size. At the arrival of e_t , the estimate of $|H_G^F|$ is incremented by $1/p_t^F$ for all instances of F that e_t belongs to in our sample $\widetilde{E}_G \cup \{e_t\}$. A pseudo-code is provided in Algorithm 1. This approach computes estimates equal to $|H_G^F|$ on expectation.

Theorem 1. *Algorithm 1 provides unbiased estimates: $\mathbb{E}[N_G^F] = |H_G^F|$.*

PROOF. Let h be a subgraph in H_G^F . We define X_h as a random variable such that $X_h = 1/p_t^F$ if h is detected at the arrival of e_t , and 0 otherwise. Clearly, $N_G^F = \sum_{h \in H_G^F} X_h$, and $\mathbb{E}[X_h] = (1/p_t^F) \times p_t^F = 1$. Thus,

$$\mathbb{E}[N_G^F] = \mathbb{E}\left[\sum_{h \in H_G^F} X_h\right] = \sum_{h \in H_G^F} \mathbb{E}[X_h] = \sum_{h \in H_G^F} 1 = |H_G^F|$$

□

When e_t arrives, the only subgraphs counted are the ones that e_t belongs to. This ensures that no subgraph is counted more than once. Due to its previous success in subgraph estimation [15, 45, 46], we utilize reservoir sampling [51]. With reservoir sampling, the probability of detecting a subgraph F at the arrival of e_t is equal to the probability that F 's other $|E_F| - 1$ edges are present in the sample after $t - 1$ time-steps. Thus, we can write:

$$p_t^F = \min\left\{1, \prod_{i=0}^{|E_F|-2} \frac{b-i}{t-1-i}\right\}$$

Algorithm 1 Compute-Estimate(S, F, b)

- 1: $\widetilde{E}_G \leftarrow \emptyset$
 - 2: $N_G^F \leftarrow 0$
 - 3: **for** $t = 1$ to $|E_G|$ **do**
 - 4: $G' \leftarrow (V_G, \widetilde{E}_G \cup \{e_t\})$
 - 5: $N \leftarrow$ number of instances of F in G' that e_t belongs to.
 - 6: $N_G^F \leftarrow N_G^F + N \times 1/p_t^F$
 - 7: Discard or store e_t in \widetilde{E}_G , based on the sampling method and b
 - 8: **return** N_G^F
-

To analyze the effect of the budget on our estimates, we derive an upper bound for the variance of N_G^F . Although loose, the bound shows that better estimates are obtained for any connected graph F with increasing b .

Theorem 2. Let N_G^F be the estimate of $|H_G^F|$ obtained using Algorithm 1 with reservoir sampling. Then,

$$\text{Var}[N_G^F] \leq |H_G^F|^2 \prod_{i=0}^{|E_F|-2} \frac{|E_G| - i}{b - i}$$

PROOF. The theorem is true when $b \geq |E_G| - 1$. Thus, we focus on the case when $b < |E_G| - 1$. As in Theorem 1, we define X_h as a random variable such that $X_h = 1/p_t^F$ if h is detected at the arrival of e_t , and 0 otherwise. It is clear from the definition of p_t^F that $p_t^F \geq p_{t+1}^F$ for all t , and thus $p_t^F \geq p_{|E_G|}^F$. Hence, $\text{Var}[X_h] = \mathbb{E}[X_h^2] - \mathbb{E}[X_h]^2 = 1/p_t^F - 1 \leq 1/p_{|E_G|}^F$. The Cauchy-Schwarz inequality can be used to bound the total variance like so:

$$\begin{aligned} \text{Var}[N_G^F] &= \sum_{h \in H_G^F} \sum_{h' \in H_G^F} \text{Cov}[X_h, X_{h'}] \leq \sum_{h \in H_G^F} \sum_{h' \in H_G^F} \sqrt{\text{Var}[X_h] \text{Var}[X_{h'}]} \\ &\leq \sum_{h \in H_G^F} \sum_{h' \in H_G^F} \frac{1}{p_{|E_G|}^F} = |H_G^F|^2 \prod_{i=0}^{|E_F|-2} \frac{|E_G| - 1 - i}{b - i} \end{aligned}$$

□

As observed in [45], we note that this methodology can be used to estimate vertex counts (the number of subgraphs that each vertex belongs to) as well. Moreover, Theorems 1 and 2 can also be extended to vertex counts.

3.4 Improving Estimation Quality with Multiple Workers

Shin et al. proposed a model for triangle estimation which takes advantage of a master machine and multiple worker machines that work in parallel. Each machine independently receives edge streams, estimates triangle counts, then sends them to the master machine, which aggregates each machine's estimate [45]. They show that using W worker machines decreases the estimates' variance by a factor of $1/W$. Thus, we use their approach to improve the quality of subgraph estimations used in our descriptors.

4 GRAPH DESCRIPTORS

In this section, we describe three graph descriptors: Graphlet Amounts via Budgeted Estimates (GABE), Moments of Attributes Estimated on Vertices Efficiently (MAEVE), and Spectral Attributes for Networks via Taylor Approximation (SANTA). These are based on the GRAPHLET KERNEL [43], NETSIMILE [9], and NETLSD [48], respectively. For each descriptor, we describe its features and how it can be computed using subgraph enumeration. We also analyze their algorithms to show that constraints $C1$, $C2$, and $C3$ (from Section 3.2) are met. Each descriptor's details have been summarized in Table 2.

Name	Summarized Description	# Passes	Time Complexity	Space Complexity
GABE	Normalized subgraph counts	1	$O(b \log b E_G)$	$O(b + V_G)$
MAEVE	Aggregated local features	1	$O(b E_G + V_G)$	$O(b + V_G)$
SANTA	Functions on eigenspectrum	2	$O(b \log b E_G)$	$O(b + V_G)$

Table 2. A summary of the proposed descriptors.

4.1 GABE: Graphlet Amounts via Budgeted Estimates

The first descriptor we propose is based on normalized subgraph counts. Subgraph counts have been popular in graph classification literature (e.g., [3, 18, 41, 43]) and have been shown to provide fruitful descriptors by capturing the prevalence of small local structures throughout a graph.

Let \mathcal{F}_k be the set of graphs with order k . In their work on the GRAPHLET KERNEL, Shervashidze et al. [43] propose measuring the similarity between two graphs G_1 and G_2 by counting the number of graphlets in \mathcal{F}_k and computing the inner product $\langle \phi_k(G_1), \phi_k(G_2) \rangle$, where for a given k and graphs $F_i \in \mathcal{F}_k$:

$$\phi_k(G) := \frac{1}{\binom{|V_G|}{k}} \left[\left| \widehat{H}_G^{F_1} \right| \quad \cdots \quad \left| \widehat{H}_G^{F_{|\mathcal{F}_k|-1}} \right| \quad \left| \widehat{H}_G^{F_{|\mathcal{F}_k|}} \right| \right]^\top$$

They compute the exact counts of all graphlets in \mathcal{F}_k for $k \in \{3, 4, 5\}$. Unfortunately, their algorithm uses adjacency matrices and adjacency lists, which take $O(|V_G|^2)$ and $O(|V_G| + |E_G|)$ space, respectively. Moreover, the time complexity is $O(|V_G| d^{k-1})$, where $d = \max_{v \in V_G} d_v^o$ is the maximum degree across all vertices in G . Thus, their algorithm does not scale well to large graphs. Although the authors introduce a sampling method to approximate $\phi_k(G)$, it requires storing the entire graph in memory and therefore does not meet our constraints.

Notation	Description
k	Maximum order of a subgraph enumerated in G by GABE
\mathcal{F}	Family of all graphs with at most four vertices
O	Overlap matrix
$\mathcal{H}_G^{\mathcal{F}}$	Vector of subgraph counts
$\widehat{\mathcal{H}}_G^{\mathcal{F}}$	Vector of induced subgraph counts

Table 3. Notation table for Section 4.1.

We construct our descriptors by estimating subgraph counts and using linear combinations of these counts to compute induced subgraph counts, similar to the methodology used by Chen et al. [15]. The linear combinations are based on the overlap of graphs of the same order. Using this approach, we estimate, for a given graph G , $\phi_k(G)$ for $k \in \{2, 3, 4\}$. Each $\phi_k(G)$ is concatenated to construct our final descriptor. There are 17 graphs with ≤ 4 vertices, each shown in Figure 2. Note that Chen et al. do not discuss the estimation of disconnected subgraphs. We discuss how we compute these in the section to follow.

4.1.1 Induced Subgraph Counts. Let $\mathcal{F} = \{F_1, F_2, \dots, F_{17}\}$ be the set of all graphs with at most four vertices. Let $\mathcal{H}_G^{\mathcal{F}}$ (resp., $\widehat{\mathcal{H}}_G^{\mathcal{F}}$) be a $|\mathcal{F}|$ -dimensional vector where the i^{th} entry corresponds to $|H_G^{F_i}|$ (resp., $|\widehat{H}_G^{F_i}|$). Let O be an ‘‘overlap matrix.’’ O is an $|\mathcal{F}| \times |\mathcal{F}|$ matrix such that the element $O(i, j)$ corresponds to the number of subgraphs of F_j isomorphic to F_i when F_i and F_j have the same number of vertices. The value is set to zero when the orders $|V_{F_i}|$ and $|V_{F_j}|$ are not equal.

Observe that $\mathcal{H}_G^{\mathcal{F}} = O\widehat{\mathcal{H}}_G^{\mathcal{F}}$. This is because for a single subgraph $F_i \in \mathcal{F}$, the overlap matrix counts the number of F_i 's induced in G , and the number of F_i 's that occur in induced F_j 's for each $F_j \in \mathcal{F}$ such that F_i is a subgraph of F_j . Note that O is invertible since it is an upper triangular matrix. Thus we can compute the vector of induced subgraphs using the formula $\widehat{\mathcal{H}}_G^{\mathcal{F}} = O^{-1}\mathcal{H}_G^{\mathcal{F}}$. Thus, our proposed approach is to compute $N_G^{F_i}$ using our estimation technique, and $\widehat{N}_G^{F_i}$ using the overlap matrix, where $N_G^{F_i}$ (resp., $\widehat{N}_G^{F_i}$) is the estimate of $|H_G^{F_i}|$ (resp., $|\widehat{H}_G^{F_i}|$).

The estimated counts of each subgraph are computed as follows:

- (1) *Connected Subgraphs.* The graphs $F_6, F_{13}, \dots, F_{17}$ are computed as described in Section 3.3; edge-centric algorithms were written to enumerate over all instances in the sample $\widetilde{E}_G \cup \{e_t\}$ and increment the estimates as described earlier. The counts of the star graphs, F_5 and F_{12} , are computed using the degrees of each vertex and the formulas written in Table 4. F_2 is simply equal to the number of edges in G .
- (2) *Disconnected Subgraphs.* Combinatorial formulas based on the estimates of connected subgraphs, $|E_G|$, and $|V_G|$. Note that the size of G can be computed by keeping track of the number of edges received, and the order can be computed by tracking the maximum vertex label, on account of each vertex being labeled in the range $[0, |V_G| - 1]$.

4.1.2 Time and Space Complexity. Let G' denote the graph represented by $\widetilde{E}_G \cup \{e_t\}$. We assume that G' is stored as an adjacency list, where the list of neighbors for each vertex is stored in a sorted, tree-like structure. Thus, determining if two vertices are adjacent takes $O(\log b)$ time.

The diameter of each connected graph in \mathcal{F} is 2. Thus, for an edge $e_t = (u_t, v_t)$, only vertices at most two hops away from u_t or v_t need to be visited. At most, three adjacency checks are needed




Fig. 2. All graphs on at most four vertices, and O , the overlap matrix. Zeros are omitted.

Graph	Formula	Graph	Formula	Graph	Formula
	$\binom{ V_G }{2}$		$\binom{ V_G }{3}$		$\binom{ V_G }{4}$
	$ E_G $		$ E_G (V_G - 2)$		$ E_G \binom{ V_G - 2}{2}$
	$\binom{ E_G }{2} - H_G^{F_5} $		$\sum_{v \in V_G} \binom{d_v^v}{2}$		$\sum_{v \in V_G} \binom{d_v^v}{3}$
	$ H_G^{F_5} (V_G - 3)$		$ H_G^{F_6} (V_G - 3)$	-	-

Table 4. Graphs and their corresponding subgraph count formulas.

to discover each connected graph. Hence, for a single edge, $2 \left(\sum_{w \in N_{G'}(u_t)} d_{G'}^w + \sum_{w \in N_{G'}(v_t)} d_{G'}^w \right) \times 3 \log b = O(b \log b)$ time is taken to process one edge. Thus, checking the entire graph takes $O(b \log b |E_G|)$ time. $O(|V_G|)$ integers are stored to keep track of the degrees of each $v \in V_G$. Since each value can be accessed in $O(1)$ time, the counts for F_5 and F_{12} can be updated each time an edge arrives in $O(1)$ time as well.

It takes $O(1)$ time to compute the remaining estimates. Thus, the total runtime is $O(b \log b |E|)$. Storing the adjacency list and degree array takes $O(b + |V_G|)$ space.

4.1.3 Effect of Increasing k . This method may be extended to further k to create richer descriptors. This would require implementing algorithms to find connected components on k vertices, deducing formulas to count the disconnected components, and constructing the overlap matrix to find the induced counts. However, obtaining counts for all subgraphs on k vertices requires finding k -cliques, which have $\binom{k}{2}$ edges. The probability of detecting larger cliques in the stream will decrease with increasing k . Thus, increasing too much larger k is likely to be unfeasible.

4.2 MAEVE: Moments of Attributes Estimated on Vertices Efficiently

NETSIMILE [9] proposed extracting local features for each vertex and aggregating them by taking various moments over their distribution. The features chosen by the authors are based on four social theories that allowed them to encompass the connectivity, transitivity, and reciprocity among the vertices and the control of information flow across graphs.

Similarly, we extract a subset of those features—chosen because they require at most one pass of the edge stream, listed in Table 6. As in NETSIMILE, the mean, standard deviation, skewness, and kurtosis for each feature are computed over the vertices. The only moment used in NETSIMILE ignored in our work is the median, left out to ensure that only one pass is needed over the vertices’ features.

4.2.1 Extracting Vertex Features. For a graph G , and a vertex $v \in V_G$, $I_G(v)$ (see Table 5 for description) is defined as the induced subgraph of G formed by $\{v\} \cup N_G(v)$, i.e., $V_{I_G(v)} = N_G(v) \cup \{v\}$ and $E_{I_G(v)} = \{(u, v) | u, v \in V_{I_G(v)} \wedge (u, v) \in E_G\}$. Note that $I_G(v)$ is also referred to as the “egonet” of v . We define $T_G(v)$ and $P_G(v)$ as the set of triangles that v belongs to and the set of three-paths (paths on three vertices) where v is an end-point, respectively. In Theorem 3 (described below), we show that each feature described in Table 6 can be calculated using values for d_G^v , $|T_G(v)|$, and $|P_G(v)|$. Thus, the vertex counts for triangles are estimated for each vertex, as described in Sections 3.3 and 4.1. Note that unlike in 4.1, the three-path estimates are not computed via the formula in Table 4 since this formula provides no information on the number of three-paths for each vertex. Moreover, the formula $\binom{d_G^v}{2}$ only provides us with the number of three-paths in which v is the middle vertex. Thus, an edge-centric algorithm is employed for each vertex to estimate the number of three-paths it ends at via the stored sample.

Notation	Description
$I_G(v)$	Subgraph induced on $\{v\} \cup N_G(v)$
$T_G(v)$	Number of triangles v belongs to
$P_G(v)$	Number of paths v belongs to, as an endpoint

Table 5. Notation table for common terms used throughout Section 4.2.






Degree	Clustering Coefficient	Avg. Degree of $N_G(v)$	Edges in $I_G(v)$	Edges leaving $I_G(v)$
d_G^v	$ T_G(v) / \binom{d_G^v}{2}$	$1 + P_G(v) / d_G^v$	$d_G^v + T_G(v) $	$ P_G(v) - 2 T_G(v) $
				

Table 6. Features extracted for each vertex, $v \in V_G$ for MAEVE, their formulae, and a figure highlighting the relevant edges. The filled-in vertex depicts v .

Theorem 3. *Each feature described in Table 6 can be expressed in terms of d_G^v , $|T_G(v)|$, and $|P_G(v)|$.*

PROOF. Observe that the degree and clustering coefficient of a vertex is already written in terms of d_G^v and $|T_G(v)|$. We will now show that the remaining three features can also be formulated in terms of d_G^v , $|T_G(v)|$, and $|P_G(v)|$.

Average Degree of Neighbors: Consider a vertex $u \in N_G(v)$. For each $w \in N_G(v) \setminus \{v\}$, w must end at a the three-path (v, u, w) . The only remaining edge for each $u \in N_G(v)$ is v itself. Note that when summing over the degrees of all vertices in $N_G(v)$, v appears once in each degree, and thereby d_G^v times in total. Hence $|P_G(v)| + d_G^v = \sum_{u \in N_G(v)} d_G^u$ and the average can be expressed as $\frac{1}{d_G^v} \sum_{u \in N_G(v)} d_G^u = 1 + \frac{|P_G(v)|}{d_G^v}$

Edges in $I_G(v)$: Let $X \subseteq E_{I_G(v)}$ be the set of all edges in $E_{I_G(v)}$ that are incident on v , and $\bar{X} = E_{I_G(v)} \setminus X$ be the complement of X . Clearly, $E_{I_G(v)} = X \cup \bar{X}$ and $X \cap \bar{X} = \emptyset$. Thus, $|E_{I_G(v)}| = |X| + |\bar{X}|$. By definition, there are exactly d_G^v edges incident on v , and each of them belongs to $E_{I_G(v)}$. Thus, $|X| = d_G^v$.

Now, consider any edge $(u, w) \in \bar{X}$. Recall that $u, w \neq v$, so, by the definition of $I_G(v)$, $u, w \in N_G(v)$. Thus, (u, w) must be part of the triangle $\{(u, v), (v, w), (u, w)\}$. Since each edge in \bar{X} forms a triangle incident on v , we have that $|\bar{X}| = |T_G(v)|$. Hence, $|E_{I_G(v)}| = |X| + |\bar{X}| = d_G^v + |T_G(v)|$.

Edges leaving $I_G(v)$: Consider a three-path $(v, u)(u, w)$. Clearly, if $w \notin N_G(v)$, (u, w) must be an edge leaving $I_G(v)$. Thereby, the number of edges leaving $I_G(v)$ must be all three-paths starting at v and ending at a vertex not in $N_G(v)$. Thus, we must account for all three-paths starting at v that lie in $I_G(v)$. Clearly, if $w \in E_{I_G(v)}$, then the following three-paths are formed: $(v, u)(u, w)$, and $(v, w), (v, u)$. Hence, each triangle in $T_G(v)$ contributes twice to the number of paths in $P_G(v)$, and we can formulate the feature as $|P_G(v)| - 2|T_G(v)|$. \square

Observe that each feature is a linear combination of our estimated variables, $|T_G(v)|$ and $|P_G(v)|$. Thus, we note that the features computed for each vertex are equal to the true value on expectation, as per Theorem 1 and the linearity of expectation.

4.2.2 Time and Space Complexity. We assume the same adjacency list structure described in Section 4.1. Let G' be the sampled graph. Three arrays of length $|V_G|$ are used to store the values of d_G^v , $P_G(v)$, and $|T_G(v)|$ for all $v \in V_G$. The degree of each vertex takes $O(1)$ time to update. Let $e_t = (u_t, v_t)$ be the edge arriving at time t . Due to the sorted nature of our adjacency list, triangles incident on e_t can be found by computing the intersection of $N_G(u_t)$ and $N_G(v_t)$ in $O(|N_G(u_t)| + |N_G(v_t)|)$ time. Counting three-paths also takes $O(|N_G(u_t)| + |N_G(v_t)|)$ time, as one pass over each neighborhood is required. Thus, the time taken to process each edge is $O(|N_G(u_t)| + |N_G(v_t)|) = O(b)$, and processing all edges takes $O(b|E_G|)$ time. After processing the entire edge stream, computing the moments over all arrays takes $O(|V_G|)$ time. Thus, the total runtime is $O(b|E_G| + |V_G|)$. The space complexity is $O(b + |V_G|)$, on account of storing b edges in the adjacency list and a constant number of arrays of length $|V_G|$.

4.3 SANTA: Spectral Attributes for Networks via Taylor Approximation

For a graph $G = (V_G, E_G)$, let A_G be its adjacency matrix. Let D_G be a diagonal matrix where $D_G(i, i)$ is the degree of vertex $v_i \in V_G$. Let $\mathcal{L}_G = I_G - D^{-\frac{1}{2}}AD^{-\frac{1}{2}}$ be the normalized Laplacian of G (see Table 7 for notation description), where I_G is the $|V_G| \times |V_G|$ identity matrix. Let λ_k be the k^{th} eigenvalue of \mathcal{L}_G , and $\Lambda_G = (\lambda_1, \lambda_2, \dots, \lambda_{|V_G|})$ refers to the eigenspectrum of \mathcal{L}_G . In [48], Tsitsulin et al. present NETLSD: a descriptor based on the spectral properties of a graph. NETLSD's descriptors are based on functions of the form $\psi_j : \Lambda_G \rightarrow \mathbb{R}$ which map \mathcal{L}_G 's eigenspectrum to a real number, based on a parameter j . For a set of parameters $\{j_1, j_2, \dots, j_m\}$, the vectors take the following form:

$$[\psi_{j_1}(\Lambda_G) \quad \psi_{j_2}(\Lambda_G) \quad \cdots \quad \psi_{j_{m-1}}(\Lambda_G) \quad \psi_{j_m}(\Lambda_G)]^T$$

The authors of NETLSD define six different functions based on two “kernels” and three normalization factors based on the eigenspectrums of complete graphs and their complements on $|V_G|$ vertices. Each function is of the form:

$$\psi_j(\Lambda_G) = \alpha \times \text{Re} \left(\sum_{\lambda_k \in \Lambda_G} e^{-j\beta\lambda_k} \right)$$

where α is a normalization factor dependent on $|V_G|$ and j , and $\beta \in \{1, i\}$. Each function has been mentioned in Table 8. Note that $\beta = 1$ for Heat kernel, and $\beta = i = \sqrt{-1}$ for the Wave kernel. For

Notation	Description
I_G	$ V_G \times V_G $ identity matrix
Λ_G	List of eigenvalues
\mathcal{L}_G	Normalized Laplacian of G
ψ_j	Function that maps Λ_G to a real number based on j
j	Parameter for ψ_j
n	Exponent of \mathcal{L}_G

Table 7. Notation table for common terms used throughout Section 4.3.

small values of j , Tsitsulin et al. suggest approximating the functions using the Taylor expansion:

$$\alpha \sum_{k=0}^{\infty} \frac{\text{tr}((-j\beta\mathcal{L}_G)^k)}{k!} = \alpha \text{tr}(I_G) - \alpha j\beta \text{tr}(\mathcal{L}_G) + \alpha \frac{(j\beta)^2}{2} \text{tr}(\mathcal{L}_G^2) + \dots$$

The authors of NETLSD discuss approximating $\psi_j(\Lambda_G)$ for small j using three Taylor terms. By enumerating over subgraphs, we propose using the first five terms of the Taylor expansion to construct a descriptor similar to NETLSD's for small values of j :

$$\psi_j(\Lambda_G) = \alpha \text{Re} \left(\text{tr}(I_G) - j\beta \text{tr}(\mathcal{L}_G) + \frac{(j\beta)^2}{2} \text{tr}(\mathcal{L}_G^2) - \frac{(j\beta)^3}{6} \text{tr}(\mathcal{L}_G^3) + \frac{(j\beta)^4}{24} \text{tr}(\mathcal{L}_G^4) \right)$$

In the remainder of this section, we discuss how subgraph enumeration can be used to compute $\text{tr}(\mathcal{L}_G^n)$ for $n \leq 4$ and a two-pass algorithm that can approximate NETLSD using the estimation scheme discussed previously.

4.3.1 Computing the Trace via Subgraph Enumeration. For an adjacency matrix A_G , $A_G^n(u, v)$ is the number of walks of length n from u to v . The n^{th} product of the Laplacian behaves similarly with the added facts that: (1) we must also consider the self-loops induced on each vertex due to the 1's in \mathcal{L}_G 's diagonal, and (2) the value added to $\mathcal{L}_G^n(u, v)$ by a walk will be a product of the "weights" of each of its edges, as each entry in the Laplacian corresponds to the following:

$$\mathcal{L}_G(u, v) = \begin{cases} 1, & \text{if } u = v \text{ and } d_G^u > 0 \\ -\frac{1}{\sqrt{d_G^u d_G^v}} & \text{if } (u, v) \in E_G \\ 0 & \text{otherwise} \end{cases}$$

Using these facts, one can assert the following:

Theorem 4. *The value of $\text{tr}(\mathcal{L}_G^n)$ can be computed for $n \in \{2, 3, 4\}$ by enumerating over all subgraphs on at most four vertices.*

PROOF. Clearly, $\mathcal{L}_G^n(u, u)$ is equal to the sum of the weights of all walks with $\leq n$ edges from a vertex u to itself. Thus, it is sufficient to enumerate all such walks and sum the weight of each walk. We do this by enumerating all relevant subgraphs, then adding a term that accounts for the weight of each walk in the subgraph and the number of walks within it. The largest subgraph induced by a walk of length n from a vertex to itself is a n -cycle, which has n vertices. The relevant subgraphs for each $n \in \{2, 3, 4\}$ are shown in Tables 9, 10, and 11. Observe that each coefficient of each term is determined by the number of walks of length n possible on the corresponding subgraph, and each

Kernel	Normalization		
	None	Empty	Complete
Heat	$\sum e^{-j\lambda}$	$\frac{1}{ V_G } \sum e^{-j\lambda}$	$\frac{\sum e^{-j\lambda}}{1 + (V_G - 1)e^{-j}}$
Wave	$\text{Re}(\sum e^{-ij\lambda})$	$\frac{1}{ V_G } \text{Re}(\sum e^{-ij\lambda})$	$\frac{\text{Re}(\sum e^{-ij\lambda})}{1 + (V_G - 1)\cos(j)}$

Table 8. Each cell represents a function of the form ψ_j proposed by the authors of NETLSD. All summations are taken over each eigenvalue $\lambda \in \Lambda_G$.

term is determined by the product of the weights of the edges as specified by the definition of the Laplacian. \square



Subgraph	Walks	Term	Subgraph	Walks	Term
	www	1		$wxw \quad xwx$	$\frac{2}{d_G^w d_G^x}$

Table 9. Subgraphs and terms relevant to computing $\text{tr}(\mathcal{L}_G^2)$.




Subgraph	Walks	Term	Subgraph	Walks	Term
	$wwww$	1		$wwxw \quad wxww \quad wxxw$ $xxwx \quad xwxw \quad xwxx$	$\frac{6}{d_G^w d_G^x}$
	$wxyw \quad xywx \quad ywxw$ $wyxw \quad xwyx \quad yxwy$	$-\frac{6}{d_G^w d_G^x d_G^y}$			

Table 10. Subgraphs and terms relevant to computing $\text{tr}(\mathcal{L}_G^3)$.

4.3.2 Computing the Descriptor on an Edge Stream. We propose a two-pass algorithm to compute the descriptor. In the algorithm's first pass, each vertex's degree is stored. In the second pass, the traces are computed; subgraphs are enumerated on the stream exactly as in the previous sections. When incrementing our count, the term to be added is multiplied by the probability of encountering it in the stream. We now show the validity of this method:

Theorem 5. *The approach proposed to approximate $\text{tr}(\mathcal{L}_G^n)$ provides unbiased estimates.*

PROOF. We present a proof similar to the one presented in Theorem 1. Let τ_n be the estimate of $\text{tr}(\mathcal{L}_G^n)$ provided by the algorithm described above. Let H_G^n be the set of subgraphs that are







Subgraph	Walks	Term	Subgraph	Walks	Term
	wwwww	1		wwwxw wxwxw wxwww wxxxw wxxww wxwww xxxwx xxwwx xwxxx xwwww xwwww xxwww	$\frac{12}{d_G^w d_G^x}$
	wxwxw xwxwx	$\frac{2}{(d_G^w d_G^x)^2}$		wyxyw xywyx yxywy ywyxy	$\frac{4}{d_G^w d_G^x d_G^y d_G^z}$
	wxyww wyxww xywxx yxxwy wwxyw wyxxw xxywx yxwwy wxxyw xwyxx xyxwx ywxxy wxyyw xxwyx xywxx yywxy wyxww xwwyx yxwyw ywwxw wwyxw xwyyx yyxwy ywxxw	$-\frac{24}{d_G^w d_G^x d_G^y}$		wxyzw yzwxw wzyxw yxwzy xyzwx zwxyz xwzyx zywxz	$\frac{8}{d_G^w d_G^x d_G^y d_G^z}$

Table 11. All subgraphs and terms relevant to computing $\text{tr}(\mathcal{L}_G^4)$.

observed to increment τ_n . For each $h \in H_G^n$, let δ_h be the term added to τ_n when h is discovered in the stream. Recall from our prior discussion that $\text{tr}(\mathcal{L}_G^n)$ can be defined as follows:

$$\text{tr}(\mathcal{L}_G^n) = \sum_{h \in H_G^n} \delta_h$$

Let X_h be a random variable such that $X_h = \delta_h \times \frac{1}{p_t}$ if h is discovered at the arrival of its last edge, and 0 otherwise, where p_t is the probability of detecting h . Clearly, $\mathbb{E}[X_h] = (\delta_h/p_t) \times p_t = \delta_h$. We now analyze the expectation of τ_n :

$$\mathbb{E}[\tau_n] = \mathbb{E}\left[\sum_{h \in H_G^n} X_h\right] = \sum_{h \in H_G^n} \mathbb{E}[X_h] = \sum_{h \in H_G^n} \delta_h = \text{tr}(\mathcal{L}_G^n)$$

□

4.3.3 Time and Space Complexity. The computation performed is similar to the one in Section 4.1, with the extra step of storing the degrees in the first pass, which takes $O(|E_G|)$ time. Computing the descriptors takes $O(1)$ time. Thus, the time complexity is $O(b \log b |E_G|)$. Likewise, the space complexity is $O(b + |V_G|)$.

5 EXPERIMENTAL SETUP

This section outlines the experimental setup, including the dataset statistics, hyperparameter values, and data preprocessing. We also introduce state-of-the-art methods for comparing results with our proposed model. We show the visual representation of the proposed and the existing descriptor by converting them into 2-dimensional representations. All experiments, except the ones on Malnet-TB, are performed on a single machine with 48 processors (2.50GHz Intel Xeon E5-2680v3) and 125 GB of memory. The experiments on Malnet-TB are run on a single machine with 16 processors (3.70GHz Intel Xeon W-2145) and 32 GB of memory. All algorithms were implemented¹ in C++ using an MPICHv3.2 backend. The code is built upon the Tri-Fly code, provided by Shin et al. [45]. For each experiment, 25 processors simulate 1 master machine and 24 worker machines, and each embedding is computed once.

¹<https://git.io/JEQmI>

5.1 Hyperparameters

Based on empirical observations, we use Canberra distance, $\left(d(\vec{x}, \vec{y}) := \sum_{i=1}^d \frac{|\bar{x}_i - \bar{y}_i|}{|\bar{x}_i| + |\bar{y}_i|}\right)$, as the distance metric to measure approximation error for GABE and MAEVE. While ℓ_2 -distance metric is used to evaluate SANTA. We note that these error metrics are inline with those used in the literature, (c.f. [9, 48]).

As observed later, one achieves reasonable estimates for SANTA with $j \leq 1$. Thus, as in [48], we use 60 evenly-spaced values on the logarithmic scale within the range $[0.001, 1]$ to construct the descriptors for SANTA. Note that when comparing SANTA to its actual values, the values produced by NETLSD are used. Thereby, the approximation error includes both the error introduced via subgraph estimation and the error via Taylor approximation.

5.2 Datasets Statistics

Our proposed model, along with the baselines and SOTA methods, are evaluated on various publicly available graph datasets, chosen primarily to showcase the efficacy of our model on large graphs. Eight graph classification datasets were selected from the TUDataset [31] repository: DD [42], CLB, RDT2, RDT5, and RDT12 [53], OHSU [34], GHUB [35], FMM² [33]. These datasets were selected due to the large size of the graphs within them relative to other datasets. The details for these datasets are provided in Table 12. Similarly, seven massive networks were selected from KONECT [28] (i.e., Florida, USA, CiteSeer, Patent, Flickr, Stanford, and UK) to showcase the scalability of our models. The details of these graphs are provided in Table 13.

- (1) REDDIT graphs³ were randomly sampled to construct a dataset to evaluate the approximation quality of our proposed methodology. Each graph represents a subreddit, wherein a vertex is a user within that subreddit, and an edge represents two users who have interacted within the subreddit. RDT2, RDT5, and RDT12 are datasets of REDDIT graphs, as described earlier.
- (2) DD is a bioinformatics dataset. Graphs in DD represent protein structures. A protein is represented as a graph, where the vertex represents amino acids, and there will be an edge between two vertices if they are connected less than 6 Angstroms apart.
- (3) OHSU is a bioinformatics dataset. OHSU graphs represent brain networks, wherein each vertex represents a region of the brain, and two regions are linked if they are correlated.
- (4) Graphs in CLB represent networks of researchers (each node is a single researcher) where edges represent that two researchers have collaborated.
- (5) GHUB graphs are social networks of developers (each developer is a node) who “starred” popular machine learning and web development repositories on Github to make the edges.
- (6) Graphs in FMM represent 3D point clouds of household objects, wherein each vertex represents an object, and two objects share an edge when there are nearby.
- (7) The Malnet dataset was used to test the efficacy of our models on large-scale classification tasks [20]. The dataset aims to distinguish malware based on “function call graphs.” Thus, each graph represents a program, wherein each vertex is a function within the program, and each edge represents an “inter-procedural” call. In lieu of using the entire dataset, we used two subsets: Trojan and Benign. We consider the binary classification task of distinguishing between these two sets.
- (8) Florida (FO) and USA (US) datasets in KONECT are the road network graphs, where each edge represents a road, and each vertex represents the intersection of two or more roads.

²<http://www.first-mm.eu/data.html>

³<https://dynamics.cs.washington.edu/data.html>

- (9) CiteSeer (CS) and Patent (PT) datasets in KONECT are the citation networks, where vertices represent documents and connected vertices represent documents that reference each other.
- (10) Flickr (FL) dataset in KONECT is the friendship network, where vertices represent users on social networks, and two vertices are connected if the users are “friends”.
- (11) Stanford (SF) and UK 2002 (U2) datasets in KONECT are the hyperlink networks, where vertices represent webpages, and connected vertices represent webpages that link to each other.

Dataset	Graphs	Classes	max $ V_G $	max $ E_G $	Avg. Deg.
FMM	41	11	5037	21774	4.50
OHSU	79	2	171	1646	4.33
DD	1178	2	5748	14267	4.98
RDT2	2000	2	3782	4071	2.34
RDT5	4999	5	3648	4783	2.25
CLB	5000	3	492	40120	37.39
RDT12	11929	11	3782	5171	2.28
GHUB	12725	2	957	9336	3.20
Malnet-TB	258373	2	551873	1639647	4.15

Table 12. Descriptions of each dataset used in our work for graph classification. For each dataset, we list the number of graphs, the number of classes, the largest order and size of a graph within the dataset, and the average degree across all graphs in the dataset.

Graph	$ V_G $	$ E_G $	Type	Description
Florida (FO) USA (US)	1070376 23947347	1343951 28854312	Road	Vertices are the intersections of two or more roads, edges represent roads
CiteSeer (CS) Patent (PT)	384054 3774768	1736145 16518937	Citation	Vertices are documents and are connected if one document references the other
Flickr (FL)	2302925	22838276	Friendship	Vertices represent users on social networks and are connected if the users are “friends”
Stanford (SF) UK 2002 (U2)	281903 18483186	1992636 261787258	Hyperlink	Vertices represent webpages and vertices are connected if one webpage links to the other

Table 13. Massive networks from KONECT listed alongside the number of vertices and edges, and descriptions.

For the preprocessing step, we convert each graph into an edge list. Duplicate edges and possible self-loops are removed from the list. If required, each vertex is relabelled to lie in the range $[0, |V_G| - 1]$. Finally, the list is randomly shuffled to ensure that the input stream is unbiased.

5.3 Existing State-of-the-Art Descriptors

We compare our models to the following state-of-the-art (SOTA) methods:

NETLSD [48]: NETLSD represents a graph based on the eigenspectrum of the graph’s Laplacian. Euclidean distance is used to compare embeddings, as suggested by the authors in their work. We report the best accuracy for each of the six variants of NETLSD.

FEATHER [36]: FEATHER’s descriptors are aggregated over characteristic function descriptors of each node in a graph. The default hyperparameters are used to construct the descriptors. We report the best accuracy for each of the three variants proposed.

SF [29]: The authors of SF proposed a “simple” baseline algorithm based on the eigenspectrum of a graph’s Laplacian. As suggested by the authors, the “embedding dimension” is set to the average number of nodes of each graph within a dataset.

Remark 6. Since no distance was suggested for FEATHER or SF, we compute results on Euclidean and Canberra distances and report the best accuracy.

Remark 7. Note that our models have no direct competitors, as no other graph classification paradigm is constructed to run under our proposed constraints. Despite this, we compare our model with the SOTA methods to show its effectiveness in terms of scalability.

5.4 Data Visualization

To visually compare different descriptors, we plot them using t -distributed Stochastic Neighbour Embedding (t -SNE) [30]. Given the graph descriptors, t -SNE computes a 2-dimensional representation of the feature vectors. Figure 3 shows the t -SNE based visualization on DD dataset for our methods (SANTA, GABE, and MAEVE on 25% and 50% budget) and NETLSD. Observe that as we increase the budget, the class-wise separation of the data becomes more prominent. Moreover, SANTA shows the most similar representation of data to NETLSD.

6 RESULTS AND DISCUSSION

In this section, we report the results of our experiments and show the changes in approximation performance with varying values of the budget b . We also report the accuracies of classifiers learned on these descriptors. Furthermore, we demonstrate the scalability of the corresponding streaming algorithms. The distance between the exact and the approximate descriptor (output of the algorithms) is referred to as the approximation error of the algorithms. Note that in the figures ahead, SANTA-XY corresponds to the SANTA descriptor variant with kernel x (H or W) and normalization Y (N, E, or C).

6.1 Approximation Quality

In this section, we test the approximation quality of our descriptors. We uniformly sampled 1,000 graphs of size 10,000 to 50,000 from REDDIT, representing interactions in various “sub-reddits”.

6.1.1 Effect of number of Taylor terms for SANTA. We first show in Figure 4 how increasing the number of Taylor terms affects the approximation quality of SANTA with respect to j . For 1000 linearly spaced values of $j \in [0.001, 1]$, the relative error (defined as $\|x - \hat{x}\|/x$, where x is the real value and \hat{x} is the approximation) across 1000 REDDIT graphs is averaged and plotted. Observe that increasing the number of Taylor terms allows us to better approximate values for larger j , enabling us to use a greater range of j .

Note that there is no need to check this for each normalization since the normalization is canceled out when computing relative error. Also, note that the values produced by four terms are ignored for the wave kernel since the values introduced in the fourth term are imaginary and are not used in the descriptor.

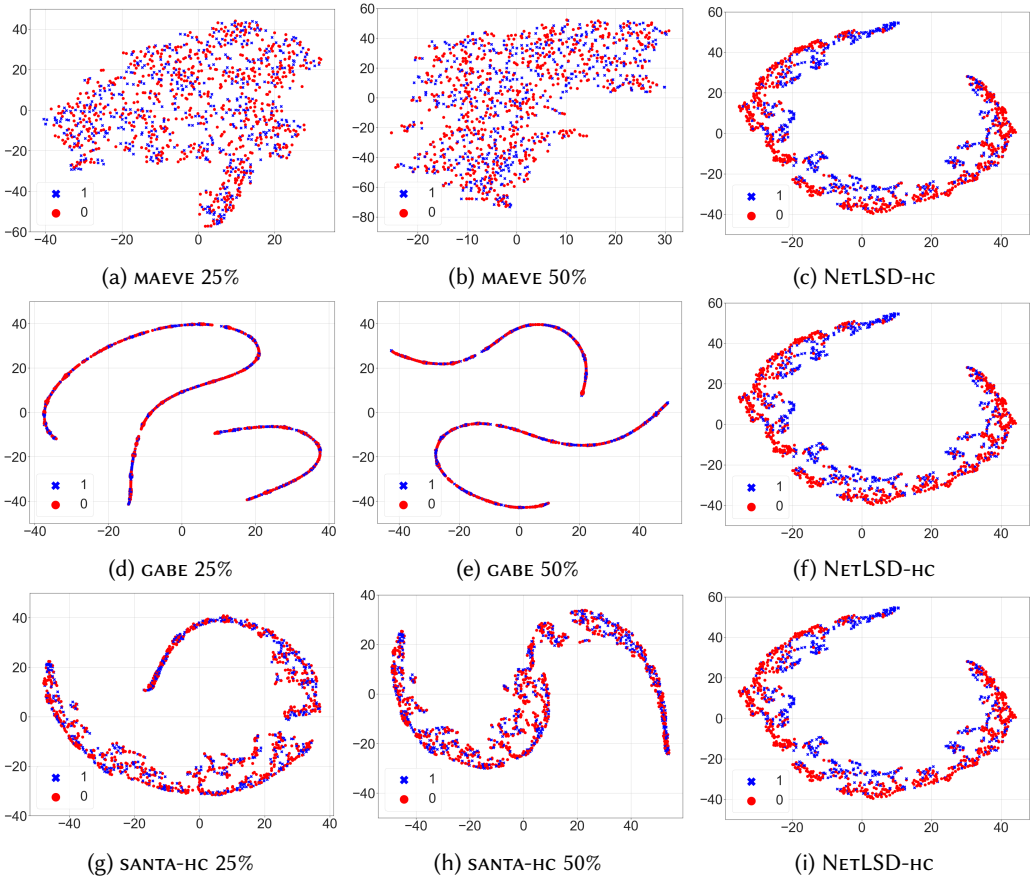


Fig. 3. t-SNE plots for different descriptors and budgets on DD dataset. Legends show true classes. The figure is best seen in color.

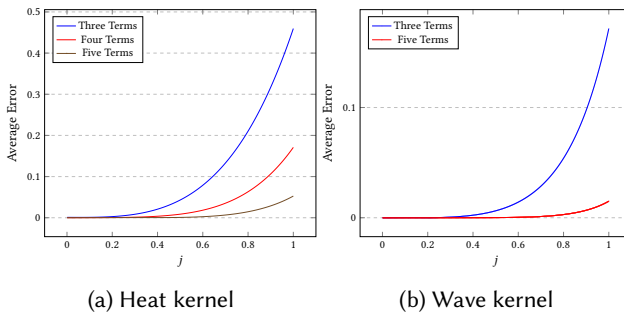


Fig. 4. Average relative error for $j \in [0.001, 1]$ of SANTA with varying number of Taylor terms.

6.1.2 *Effect of increasing the budget for each descriptor.* Figure 5 shows that the average approximation error across the sampled graphs decreases as the budget increases. Observe that normalized

versions of SANTA can achieve very low errors even with small values of b . Unfortunately, unnormalized variants of SANTA have very large errors and would likely be unfruitful in practical settings.

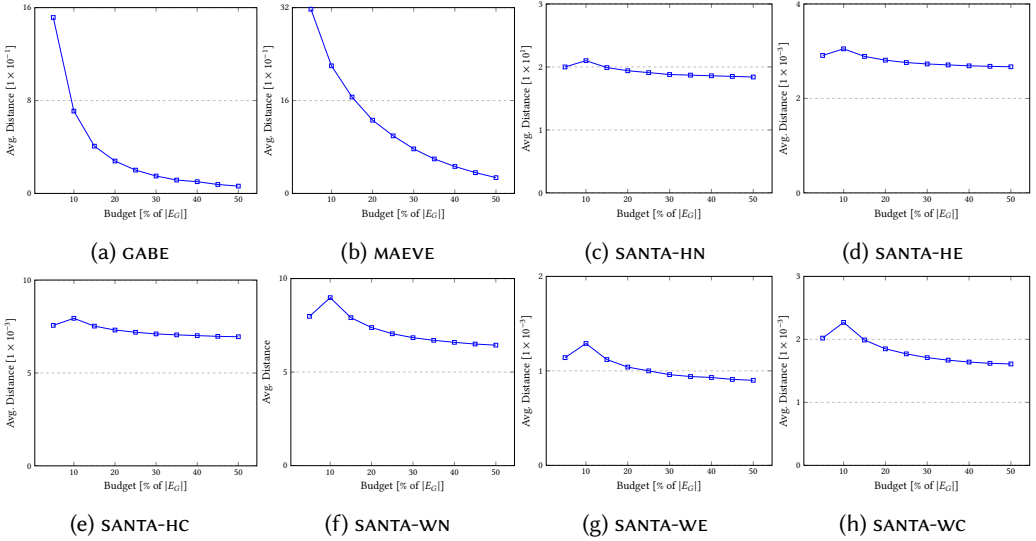


Fig. 5. Approximation errors with increasing budget (b) for GABE, MAEVE, and all variants of SANTA.

6.2 Graph Classification

We opted for a Nearest Neighbor classifier as in Tsitulin et al. [48] work on NETLSD. 10-fold cross-validation was performed for ten random splits of the dataset. The average accuracy for each fold is reported. Note that only two folds are used for FMM because each class has a small number of samples. The descriptors are computed for our models by using 25% and 50% of the number of edges of each graph.

6.2.1 Results on Different Variants of SANTA. In Table 15, we compare all variants of SANTA to find out which one works best. It is clear that SANTA-HC often provides the best results. For this reason, and because it has the lowest error across all variants in Figure 5, we recommend SANTA-HC for practical usage and compare it to other descriptors in the coming section.

In this same table, we show the results provided by NETLSD when using the same values for j . Despite the error added by the Taylor approximation and budgeted sampling, SANTA provides results comparable to NETLSD.

We observe better results for the datasets OHSU and FMM when the budget is smaller, sometimes more significant than those provided by NETLSD. We believe that due to the small size of these datasets, the noise added when approximating the embeddings is not eliminated by the classifier. Thus, we do not recommend using SANTA on smaller datasets without a larger budget; otherwise, the classifier may not be able to generalize.

Table 16 compares the classification accuracy of our proposed models and the benchmark descriptors. Despite using only a fraction of edges, our proposed descriptors provide results competitive to descriptors that have access to the entire graph in seven of the eight classification datasets. Unfortunately, SANTA is unable to compete with its competitors in most cases, despite giving results near to NETLSD when used on the same values of j (see Table 15).

6.2.2 *Comparing SANTA to SLaQ.* We report the comparison of the accuracy achieved by SANTA (across all variants) with that of SLaQ, a method introduced by Titsulin et al. [49] to approximate NetLSD, in Table 14. For DD and CLB datasets, we observe that SANTA outperforms SLaQ. For RDT5 and RDT12 datasets, although SLaQ outperforms SANTA, the predictive accuracy of SLaQ is not significantly higher compared to SANTA, despite SLaQ keeping the entire graph in memory.

Method	Budget	DD	CLB	RDT5	RDT12
SLaQ	$ E_G $	66.77	58.76	35.48	25.31
SANTA	$1/4 E_G $	68.16	63.80	35.32	24.68
	$1/2 E_G $	66.83	64.90	34.62	23.89

Table 14. Comparing reported classification accuracy of SANTA and SLaQ. Results within 1% of the best have been bold-faced.

6.2.3 *Performance on Large Classification Tasks.* To showcase the practical usage of our proposed methods, we performed graph classification on the Malnet-TB dataset on a computer with relatively weaker hardware. In this case, we used ten workers and $1/10|E_G|$ as our budget. The time taken and classification accuracy is provided in Table 17. Note that all of our proposed models can process 260K graphs with up to 550K vertices and 1.6M edges in $\approx 1 1/2$ days.

6.3 Scaling to Large Real-world Networks

In this section, we show the scalability of our proposed descriptors by running them on large real-world networks. For this purpose, we ran our algorithms on the networks listed in Table 13. For each graph, descriptors were estimated for $b \in \{100000, 500000\}$. In Table 18 and 19, we show the wall-clock time taken and the distance between the real and approximate vectors. Note that the lower values are better.

Note that to compute the real embeddings for SANTA, one would have to compute the eigenspectrum of each graph. Due to the intractability of this method, we approximate the true embeddings by approximating the eigenvalues using the largest and smallest eigenvalues of the Laplacian of each graph, as proposed in [48]. As per the authors' suggestion, we attempted to obtain 150 eigenvalues from each end of the spectrum. While this was not possible for all graphs, a minimum of 50 eigenvalues were used for each end, i.e., at least 100 eigenvalues were used to compute the NETLSD embeddings for each graph. Note that this was not possible for the UK 2002 graph due to its large size. Observe that we can process graphs with millions of edges with reasonably low approximation error. UK 2002, a graph with $\approx 260M$ edges, was processed under half an hour by all of our proposed models. We note that when $b = 500000$, GABE and SANTA take a significant amount of time to compute on the Stanford and Flickr graphs due to their dense nature. Thus, we posit that one must consider the graph's density when setting the value of b .

7 CONCLUSION

This paper proposes three graph descriptors and streaming algorithms with constant space complexity to construct them. Our descriptors extend the state-of-the-art graph descriptors and approximate their embeddings over graph streams. Experiments show that while using very less memory, our descriptors provide results comparable to SOTA descriptors, which store the entire graph in memory. We demonstrate the scalability of our algorithms to graphs with millions of edges (which is not possible for existing methods). We hope to introduce descriptors for attributed graphs that meet our

Variant	Method	Budget	DD	CLB	RDT2	RDT5	RDT12	OHSU	GHUB	FMM
HN	SANTA	$1/4 E_G $	66.22	61.90	76.02	35.12	22.38	54.50	54.88	26.80
		$1/2 E_G $	66.03	62.59	75.88	34.39	22.21	54.50	54.78	26.80
	NETLSD*	$ E_G $	66.44	63.29	75.82	33.50	21.74	56.98	55.75	27.14
HE	SANTA	$1/4 E_G $	63.98	63.80	63.77	34.90	21.42	69.96	54.56	39.70
		$1/2 E_G $	65.76	64.90	64.33	34.22	21.91	66.82	55.11	20.00
	NETLSD*	$ E_G $	60.75	64.08	61.98	29.44	19.47	52.66	57.19	21.49
HC	SANTA	$1/4 E_G $	68.16	63.44	79.14	35.32	24.68	67.98	55.99	38.76
		$1/2 E_G $	66.83	63.50	78.34	34.62	23.89	58.25	55.61	23.74
	NETLSD*	$ E_G $	65.99	64.77	75.96	37.02	25.12	55.95	55.03	35.39
WN	SANTA	$1/4 E_G $	66.70	62.49	75.68	35.08	22.76	55.30	55.32	26.80
		$1/2 E_G $	66.63	63.15	75.57	34.53	22.81	55.30	55.30	26.80
	NETLSD*	$ E_G $	66.19	63.01	75.64	33.40	22.23	54.14	58.08	28.60
WE	SANTA	$1/4 E_G $	61.55	62.52	65.10	34.09	21.66	67.59	55.04	24.37
		$1/2 E_G $	61.02	62.04	64.90	33.56	21.27	64.32	54.06	11.27
	NETLSD*	$ E_G $	59.35	64.46	62.14	26.99	19.05	60.61	58.20	15.08
WC	SANTA	$1/4 E_G $	64.15	61.25	74.26	31.45	21.43	58.48	55.12	24.38
		$1/2 E_G $	61.81	62.47	74.64	31.79	21.46	58.12	54.67	11.74
	NETLSD*	$ E_G $	64.81	62.97	75.10	29.39	21.56	47.93	56.60	19.01

Table 15. Classification accuracy (in %) using nearest neighbor classifier across all datasets for all variants of SANTA, as well as NETLSD modified to use the same values for j . Results within 1% of the best across all SANTA variants have been bold-faced.

Approach	Method	Budget	DD	CLB	RDT2	RDT5	RDT12	OHSU	GHUB	FMM
Benchmark	NETLSD	$ E_G $	70.36	74.27	82.85	41.23	30.90	73.79	55.73	27.14
	FEATHER	$ E_G $	63.57	73.14	83.22	43.09	34.33	62.77	60.95	26.81
	SF	$ E_G $	62.84	72.82	82.38	42.36	30.80	59.50	57.01	29.00
	MAEVE	$1/4 E_G $	59.44	68.42	85.04	41.15	32.57	49.07	61.99	12.90
		$1/2 E_G $	61.26	70.95	86.15	41.53	33.69	47.12	61.81	14.63
Proposed Descriptors	GABE	$1/4 E_G $	65.23	63.62	84.65	41.10	32.18	44.30	61.88	27.37
		$1/2 E_G $	69.08	65.23	85.35	40.63	32.96	41.02	62.72	25.35
	SANTA-HC	$1/4 E_G $	68.16	63.44	79.14	35.32	24.68	67.98	55.99	38.76
		$1/2 E_G $	66.83	63.50	78.34	34.62	23.89	58.25	55.61	23.74

Table 16. Accuracy of the nearest neighbor classifier on different datasets, descriptors, and benchmark methods. Results within 1% of the best have been bold-faced.

constraints in the future. Another interesting future direction is to explore neural networks that

	GABE	MAEVE	SANTA HN	SANTA HE	SANTA HC	SANTA WN	SANTA WE	SANTA WC
Accuracy	79.52	79.93	74.60	69.38	69.47	74.89	69.39	69.40
Avg. Time [s]	0.54	0.45	0.36	0.36	0.36	0.36	0.36	0.36
Max Time [min]	66.73	37.93	33.82	33.82	33.82	33.82	33.82	33.82
Total Time [hr]	38.59	32.41	25.63	25.63	25.63	25.63	25.63	25.63

Table 17. Results on Malnet-TB for GABE, MAEVE, and all variants of SANTA with $b = 1/10|E_G|$. We report the accuracy, the average and maximum amount of time taken for each graph in the dataset, and the total amount of time taken.

can process edge streams, combining the scalability of stream-based methods and the classification prowess of graph convolutional networks.

REFERENCES

- [1] M. Ahmad, S. Ali, J. Tariq, I. Khan, M. Shabbir, and A. Zaman. 2020. Combinatorial trace method for network immunization. *Information Sciences* 519 (2020), 215 – 228.
- [2] M. Ahmad, J. Tariq, M. Shabbir, and I. Khan. 2017. Spectral Methods for Immunization of Large Networks. *Australasian Journal of Information Systems* 21 (2017).
- [3] A. Ahmed, Z. Hassan, and M. Shabbir. 2020. Interpretable multi-scale graph descriptors via structural compression. *Information Sciences* 533 (2020), 169–180.
- [4] N. Ahmed, J. Neville, and R. Kompella. 2013. Network Sampling: From Static to Streaming Graphs. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 8, 2 (2013), 56.
- [5] S. Ali, M. Alvi, S. Faizullah, M. Khan, A. Alshantqi, and I. Khan. 2020. Detecting DDoS Attack on SDN Due to Vulnerabilities in OpenFlow. In *International Conference on Advances in Emerging Computing Technologies (AECT)*. 1–6.
- [6] S. Ali, M. Shakeel, I. Khan, S. Faizullah, and M. Khan. 2021. Predicting Attributes of Nodes Using Network Structure. *Transactions on Intelligent Systems and Technology (TIST)* 12, 2 (2021), 1–23.
- [7] L. Babai. 2016. Graph Isomorphism in Quasipolynomial Time. In *Symposium on Theory of Computing (STOC)*. 684–697.
- [8] J. Bento and S. Ioannidis. 2018. A Family of Tractable Graph Distances. In *International Conference on Data Mining (SDM)*. 333–341.
- [9] M. Berlingerio, D. Koutra, T. Eliassi-Rad, and C. Faloutsos. 2013. Network Similarity via Multiple Social Theories. In *International Conference Series on Advances in Social Network Analysis and Mining (ASONAM)*. 1439–1440.
- [10] L. Bo, X. Ren, and D. Fox. 2010. Kernel Descriptors for Visual Recognition. In *Neural Information Processing Systems (NeurIPS)*. 244–252.
- [11] K. Borgwardt and H. Kriegel. 2005. Shortest-Path Kernels on Graphs. In *International Conference on Data Mining (ICDM)*. 74–81.
- [12] Samuel L Braunstein, Sibasis Ghosh, and Simone Severini. 2006. The Laplacian of a graph as a density matrix: a basic combinatorial approach to separability of mixed states. *Annals of Combinatorics* 10, 3 (2006), 291–317.
- [13] S. Cao, W. Lu, and Q. Xu. 2015. Grarep: Learning graph representations with global structural information. In *International Conference on Information and Knowledge Management (CIKM)*. 891–900.
- [14] P. Chen, L. Wu, S. Liu, and I. Rajapakse. 2019. Fast Incremental Von Neumann Graph Entropy Computation: Theory, Algorithm, and Applications. In *International Conference on Machine Learning*. 1091–1101.
- [15] X. Chen and J. Lui. 2017. A Unified Framework to Estimate Global and Local Graphlet Counts for Streaming Graphs. In *International Conference Series on Advances in Social Network Analysis and Mining (ASONAM)*. 131–138.
- [16] Q. Duong, H. Ramampiaro, K. Nørnvåg, and T. Dam. 2021. Density Guarantee on Finding Multiple Subgraphs and Subtensors. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 15, 5 (2021), 32.
- [17] A. Duran and M. Niepert. 2017. Learning Graph Representations with Embedding Propagation. In *Neural Information Processing Systems (NeurIPS)*. 5119–5130.
- [18] A. Dutta and H. Sahbi. 2019. Stochastic Graphlet Embedding. *IEEE Transactions on Neural Networks and Learning Systems* 30, 8 (2019), 2369–2382.
- [19] M. Farhan, J. Tariq, A. Zaman, M. Shabbir, and I. Khan. 2017. Efficient Approximation Algorithms for Strings Kernel Based Sequence Classification. In *Neural Information Processing Systems (NeurIPS)*. 6935–6945.

		GABE	MAEVE	SANTA HN	SANTA HE	SANTA HC	SANTA WN	SANTA WE	SANTA WC
PT	Time [min]	0.52	0.77	1.11	1.11	1.11	1.11	1.11	1.11
	Distance	3.36	5.11	2.38	6.31	1.35	4.36	1.72	1.14
FL	Time [min]	5.48	3.48	4.96	4.96	4.96	4.96	4.96	4.96
	Distance	2.77	5.09	1.14	4.95	1.02	2.43	1.59	1.04
US	Time [min]	0.63	1.21	1.99	1.99	1.99	1.99	1.99	1.99
	Distance	5.24	11.39	18.30	7.64	1.92	6.32	0.37	0.28
U2	Time [min]	10.58	9.05	20.61	20.61	20.61	20.61	20.61	20.61
	Distance	6.48	9.61	-	-	-	-	-	-
FO	Time [min]	0.05	0.16	0.08	0.08	0.08	0.08	0.08	0.08
	Distance	2.07	6.67	0.75	6.96	1.76	0.21	0.27	0.21
CS	Time [min]	0.21	0.19	0.20	0.20	0.20	0.20	0.20	0.20
	Distance	1.07	3.09	0.19	4.97	1.03	0.39	1.52	1.01
SF	Time [min]	8.35	4.39	3.92	3.92	3.92	3.92	3.92	3.92
	Distance	1.10	3.55	0.20	6.98	1.50	0.35	1.86	1.23

Table 18. Approximation error and time taken for GABE, MAEVE, and all variants of SANTA with $b = 100000$. Accuracy results for U2 have been omitted since the graph was too large to obtain true values.

		GABE	MAEVE	SANTA HN	SANTA HE	SANTA HC	SANTA WN	SANTA WE	SANTA WC
PT	Time [min]	0.84	1.13	1.38	1.38	1.38	1.38	1.38	1.38
	Distance	2.65	3.14	2.38	6.32	1.35	4.36	1.72	1.14
FL	Time [min]	101.37	12.62	45.27	45.27	45.27	45.27	45.27	45.27
	Distance	2.66	3.95	1.17	5.08	1.05	2.40	1.56	1.03
US	Time [min]	0.90	1.34	1.94	1.94	1.94	1.94	1.94	1.94
	Distance	4.84	10.08	18.35	7.66	1.92	6.30	0.36	0.28
U2	Time [min]	17.23	19.18	26.88	26.88	26.88	26.88	26.88	26.88
	Distance	3.56	7.73	-	-	-	-	-	-
FO	Time [min]	0.12	0.17	0.14	0.14	0.14	0.14	0.14	0.14
	Distance	1.16	2.52	0.74	6.96	1.76	0.21	0.27	0.21
CS	Time [min]	1.88	0.71	0.83	0.83	0.83	0.83	0.83	0.83
	Distance	1.01	0.86	0.19	4.97	1.03	0.39	1.52	1.00
SF	Time [min]	174.54	29.10	54.37	54.37	54.37	54.37	54.37	54.37
	Distance	1.04	1.58	0.20	6.99	1.50	0.35	1.85	1.23

Table 19. Approximation error and time taken for GABE, MAEVE, and all variants of SANTA with $b = 500000$. Accuracy results for U2 have been omitted since the graph was too large to obtain true values.

[20] S. Freitas, Y. Dong, J. Neil, and D.H. Chau. 2021. A Large-scale Database for Graph Representation Learning. In *NeurIPS Datasets & Benchmarks*. 13.

[21] A. Grover and J. Leskovec. 2016. node2vec: Scalable feature learning for networks. In *International Conference on Knowledge Discovery and Data Mining (KDD)*. 855–864.

- [22] Z.R Hassan, M. Shabbir, I. Khan, and W. Abbas. 2020. Estimating Descriptors for Large Graphs. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*. 779–791.
- [23] C. Helma, R. King, S. Kramer, and A. Srinivasan. 2001. The Predictive Toxicology Challenge 2000-2001. *Bioinformatics* 17, 1 (2001), 107–108.
- [24] J. Jin, M. Heimann, D. Jin, and D. Koutra. 2021. Toward Understanding and Evaluating Structural Node Embeddings. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 16, 3, Article 58 (2021), 32 pages.
- [25] R. Kondor and H. Pan. 2016. The Multiscale Laplacian Graph Kernel. In *Neural Information Processing Systems (NeurIPS)*. 2982–2990.
- [26] D. Koutra, N. Shah, J. Vogelstein, B. Gallagher, and C. Faloutsos. 2016. DeltaCon: Principled Massive-Graph Similarity Function with Attribution. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 10, 3 (2016), 28:1–28:43.
- [27] P. Kuksa, I. Khan, and V. Pavlovic. 2012. Generalized Similarity Kernels for Efficient Sequence Classification. In *International Conference on Data Mining (SDM)*. 873–882.
- [28] J. Kunegis. 2013. KONECT: the Koblenz Network Collection. In *International Conference on World Wide Web (WWW)*. 1343–1350.
- [29] N. Lara and E. Pineau. 2018. A Simple Baseline Algorithm for Graph Classification. *Relational Representation Learning Workshop, NeurIPS 2018* abs/1810.09155 (2018), 7.
- [30] L. Maaten. 2014. Accelerating t-SNE using Tree-based Algorithms. *Journal of Machine Learning Research (JMLR)* 15, 1 (2014), 3221–3245.
- [31] C. Morris, N. Kriege, F. Bause, K. Kersting, P. Mutzel, and M. Neumann. 2020. TUDataset: A collection of benchmark datasets for learning with graphs. *ICML 2020 workshop on Graph Representation Learning and Beyond* abs/2007.08663 (2020), 11.
- [32] C. Morris, M. Ritzert, M. Fey, W. Hamilton, J. Lenssen, G. Rattan, and M. Grohe. 2019. Weisfeiler and Leman Go Neural: Higher-Order Graph Neural Networks. In *AAAI Conference on Artificial Intelligence (AAAI)*. 4602–4609.
- [33] M. Neumann, R. Garnett, C. Bauckhage, and K. Kersting. 2016. Propagation Kernels: Efficient Graph Kernels from Propagated Information. *Machine Learning* 102, 2 (2016), 209–245.
- [34] N. Qiang, Q. Dong, F. Ge, H. Liang, B. Ge, S. Zhang, Y. Sun, J. Gao, and T. Liu. 2021. Deep Variational Autoencoder for Mapping Functional Brain Networks. *IEEE Trans. Cogn. Dev. Syst.* 13, 4 (2021), 841–852.
- [35] B. Rozemberczki, O. Kiss, and R. Sarkar. 2020. Karate Club: An API Oriented Open-source Python Framework for Unsupervised Learning on Graphs. In *International Conference on Information and Knowledge Management (CIKM)*. 3125–3132.
- [36] B. Rozemberczki and R. Sarkar. 2020. Characteristic Functions on Graphs: Birds of a Feather, from Statistical Descriptors to Parametric Models. In *International Conference on Information and Knowledge Management (CIKM)*. 1325–1334.
- [37] S. Sanei-Mehri, A. Das, H. Hashemi, and S. Tirthapura. 2021. Mining Largest Maximal Quasi-Cliques. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 15, 5, Article 81 (2021), 21 pages.
- [38] S. Sanei-Mehri, Y. Zhang, A. Sariyüce, and S. Tirthapura. 2019. FLEET: Butterfly Estimation from a Bipartite Graph Stream. In *International Conference on Information and Knowledge Management (CIKM)*. 1201–1210.
- [39] A. Sanfeliu and K. Fu. 1983. A Distance Measure between Attributed Relational Graphs for Pattern Recognition. *IEEE Transactions on Systems, Man, and Cybernetics* 13, 3 (1983), 353–362.
- [40] M. Shakeel, A. Karim, and I. Khan. 2020. A Multi-Cascaded Model with Data Augmentation for Enhanced Paraphrase Detection in Short Texts. *Information Processing & Management* 57, 3 (2020), 1–19.
- [41] P. Shao, Y. Yang, S. Xu, and C. Wang. 2021. Network Embedding via Motifs. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 16, 3 (2021), 20.
- [42] N. Shervashidze, P. Schweitzer, E. Leeuwen, K. Mehlhorn, and K. Borgwardt. 2011. Weisfeiler-Lehman Graph Kernels. *Journal of Machine Learning Research (JMLR)* 12 (2011), 2539–2561.
- [43] N. Shervashidze, S. Vishwanathan, T. Petri, K. Mehlhorn, and K. Borgwardt. 2009. Efficient graphlet kernels for large graph comparison. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*. 488–495.
- [44] A. Sheshbolouki and M. Özsu. 2022. SGrapp: Butterfly Approximation in Streaming Graphs. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 16, 4 (2022), 43.
- [45] K. Shin, M. Hammoud, E. Lee, J. Oh, and C. Faloutsos. 2018. Tri-Fly: Distributed Estimation of Global and Local Triangle Counts in Graph Streams. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*. 651–663.
- [46] L. Stefani, A. Epasto, M. Riondato, and E. Ufpl. 2016. TRIÈST: Counting Local and Global Triangles in Fully-Dynamic Streams with Fixed Memory Size. In *International Conference on Knowledge Discovery and Data Mining (KDD)*. 825–834.
- [47] J. Tariq, M. Ahmad, I. Khan, and M. Shabbir. 2017. Scalable Approximation Algorithm for Network Immunization. In *Pacific Asia Conference on Information Systems (PACIS)*. 200.
- [48] A. Tsitsulin, D. Mottin, P. Karras, A. Bronstein, and E. Müller. 2018. NetLSD: Hearing the Shape of a Graph. In *International Conference on Knowledge Discovery and Data Mining (KDD)*. 2347–2356.

- [49] A. Tsitsulin, M. Munkhoeva, and B. Perozzi. 2020. Just SLAQ when you Approximate: Accurate Spectral Distances for Web-scale Graphs. In *Proceedings of the Web Conference 2020*. 2697–2703.
- [50] S. Verma and Z. Zhang. 2017. Hunt For The Unique, Stable, Sparse And Fast Feature Learning On Graphs. In *Neural Information Processing Systems (NeurIPS)*. 88–98.
- [51] J. Vitter. 1985. Random Sampling with a Reservoir. *ACM Trans. Math. Software* 11, 1 (1985), 37–57.
- [52] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. 2019. How Powerful are Graph Neural Networks?. In *International Conference on Learning Representations (ICLR)*. 17.
- [53] P. Yanardag and S. Vishwanathan. 2015. Deep Graph Kernels. In *International Conference on Knowledge Discovery and Data Mining (KDD)*. 1365–1374.
- [54] L. Yang, Y. Guo, D. Jin, H. Fu, and X. Cao. 2018. 3-in-1 correlated embedding via adaptive exploration of the structure and semantic subspaces. In *International Joint Conference on Artificial Intelligence (IJCAI)*. 3613–3619.
- [55] L. Yang, Y. Wang, J. Gu, C. Wang, X. Cao, and Y. Guo. 2020. JANE: Jointly Adversarial Network Embedding. In *International Joint Conference on Artificial Intelligence (IJCAI)*. 1381–1387.