

Fast and Flexible Top-k Similarity Search on Large Networks

JING ZHANG, Tsinghua University, Renmin University of China

JIE TANG, Tsinghua University

CONG MA, Tsinghua University

HANGHANG TONG, Arizona State University

YU JING, Tsinghua University

JUANZI LI, Tsinghua University

WALTER LUYTEN, KU Leuven

MARIE-FRANCINE MOENS, KU Leuven

Similarity search is a fundamental problem in network analysis and can be applied in many applications, such as collaborator recommendation in co-author networks, friend recommendation in social networks, and relation prediction in medical information networks. In this paper, we propose a sampling-based method using *random paths* to estimate the similarities based on both common neighbors and structural contexts efficiently in very large homogeneous or heterogeneous information networks. We give a theoretical guarantee that the sampling size depends on the error-bound ϵ , the confidence level $(1 - \delta)$, and the path length T of each random walk. We perform an extensive empirical study on a Tencent microblogging network of 1,000,000,000 edges. We show that our algorithm can return top- k similar vertices for any vertex in a network $300\times$ faster than the state-of-the-art methods. We develop a prototype system of recommending similar authors to demonstrate the effectiveness of our method.

CCS Concepts: • **Information systems** → **Data mining**;

General Terms: Design, Algorithms, Performance

Additional Key Words and Phrases: Vertex similarity, Similarity search, Social network, Random path, Heterogeneous information network

ACM Reference format:

Jing Zhang, Jie Tang, Cong Ma, Hanghang Tong, Yu Jing, Juanzi Li, Walter Luyten, and Marie-Francine Moens. 2010. Fast and Flexible Top-k Similarity Search on Large Networks. *Proc. ACM Meas. Anal. Comput. Syst.* 9, 4, Article 39 (March 2010), 30 pages.

DOI: 0000001.0000001

Author's addresses: Jing Zhang is with Department of Computer Science and Technology, Tsinghua University and Information School, Renmin University of China, Email: zhang-jing@ruc.edu.cn; Jie Tang*, Cong Ma, Yu Jing and Juanzi Li are with the Department of Computer Science and Technology, Tsinghua University, Beijing, China, 100084. E-mail: ma-c11@mails.tsinghua.edu.cn, {jietang, yujing5b5d, li-juanzi}@tsinghua.edu.cn; *Jie Tang is the corresponding author. Hanghang Tong is with School of Computing, Informatics, and Decision Systems Engineering, ASU. E-mail: hanghang.tong@asu.edu; Walter Luyten and Marie-Francine Moens with Katholieke Universiteit Leuven, Leuven, Belgium. E-mail: walter.luyten@med.kuleuven.be, sien.moens@cs.kuleuven.be.

ACM acknowledges that this contribution was authored or co-authored by an employee, or contractor of the national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only. Permission to make digital or hard copies for personal or classroom use is granted. Copies must 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. To copy otherwise, distribute, republish, or post, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2010 ACM. 2476-1249/2010/3-ART39 \$15.00

DOI: 0000001.0000001


1 INTRODUCTION

Similarity search is a fundamental problem in network analysis and can be applied in many applications, such as collaborator recommendation in co-author networks, friend recommendation in social networks, and drug-protein relation prediction in biological information networks.

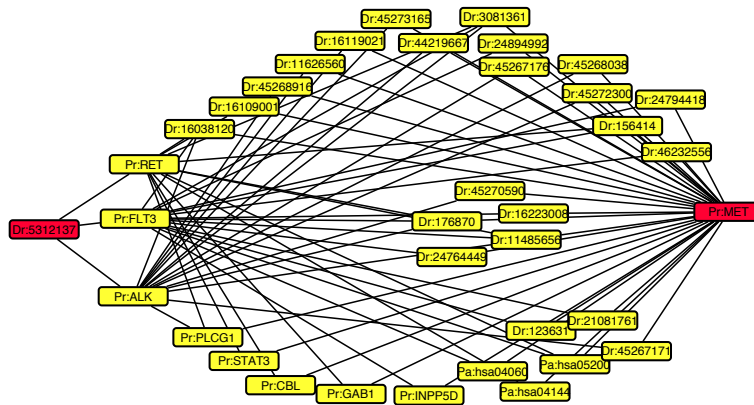
The problem has been extensively studied. In general, existing research work follows two basic principles. The first one is that two vertices are considered similar if they have many direct or indirect common neighbors in a network. For example, in a scientific co-author network shown in figure 1(a), we can say that the authors in the red circle are similar as Barabási, because they have more collaborations with Barabási than others. We category the methods following the first principle into neighborhood similarity. Jaccard index [23] and Cosine similarity [4] are two basic metrics of neighborhood similarity. However, they estimate the similarity in a local fashion. Though some work, such as SimRank [24] and VertexSim [34], use the entire network to compute similarity, they are essentially based on the transitivity of similarity in the network. Meanwhile, some researchers extend the neighborhood similarity to multi-typed neighborhood similarity based on different relation types in a heterogeneous information network [50]. The typed similarities provide extra semantic explanations for similarities. For example, in a medical information network shown in figure 1(b), we can say that the drug named 5312137 is closely related with the protein named MET, because they are connected by many paths of different semantics. Here, we use “related” instead of “similar”, because the type of the source vertex and the target vertex may not be the same in a heterogeneous information network. The second principle is that two vertices are considered equivalent if they play the same structural role—this can be further quantified by degree, closeness centrality, betweenness, and other network centrality metrics [15]. For example, in figure 1(a), we can say the authors denoted in green are similar as barabási, because they are all in the center places of certain sub-networks. Similarly, the authors denoted in red are similar as Robert, because they are all in some tight-knit groups, and those denoted in blue are similar as Rinzel, because they are in periphery places. We category the methods following the second principle into structure similarity. For example, RoleSim calculates the similarity between disconnected vertices by changing the initialization of SimRank [27], and ReFex [20] is a feature-based method to calculate vertex similarity by defining a vector of features for each vertex.

When networks get larger, the efficiency issue has become one of the biggest challenges in similarity search, because most existing similarity methods in networks are iterative and have a high computation cost. For example, SimRank results in a complexity of $O(I|V|^2\bar{d}^2)$, where $|V|$ is the number of vertices in a network; \bar{d} is the average degree of all vertices; I is the number of iterations to perform the SimRank algorithm. It is clearly infeasible to apply SimRank to large-scale networks. For example, in our experiments, when dealing with a network with 500,000 edges, even the fast (top- k) version of SimRank [33] requires more than five days to complete the computation for all vertices (as shown in Table 2). Although much effort has been made to improve the computational efficiency, such as fast random walk with restart [16, 52] and fast SimRank [32, 33], they usually only focus on improving the efficiency of neighborhood similarity search in homogeneous information networks, while, ignoring other kinds of similarity search. Thus in this paper, we aim at designing a similarity method that is flexible enough to measure both neighborhood similarity and structure similarity quickly in large homogeneous or heterogeneous information networks.

To achieve the goal of neighborhood similarity search, we define a new similarity metric, referred to as path similarity. The basic idea behind this is that two vertices have a high similarity if they frequently appear on the same paths. Then we propose a sampling-based method, referred to as Panther, based on a novel idea of *random path* to estimate path similarity. Specifically, given a network, we perform R random walks, each starting from a randomly picked vertex and walking T steps. The path similarities are calculated efficiently based on the generated paths and the inverted index from vertex to paths. We provide theoretical proof that the sample size, $R = \frac{c}{\epsilon^2} (\log_2 \binom{T+1}{2} + 1 + \ln \frac{1}{\delta})$, only depends on the path length T of each random walk, for a given error-bound



(a) Homogeneous information network.



(b) Heterogeneous information network.

Fig. 1. Case studies of similar nodes in (a) a homogeneous information network and (b) a heterogeneous information network. We use a scientific co-author network [38] as the case of a homogeneous information network. In the network, the authors having close collaborations with a given author can be treated as similar authors to the given author. In another viewpoint, similar authors can be also treated as those in similar positions to that of the given author. For example, the authors in similar position with Barabási are denoted in green, similar to that of Robert are in red, and similar to that of Rinzel are in blue. We use a medical network as the case of a heterogeneous information network. In the network, the proteins with many paths connected with a drug can be treated as related proteins to the given drug. The prefix "Dr" in the vertex label denotes Drug, "Pr" denotes Protein, and "Pa" denotes Pathway.

ε and confidence level $1 - \delta$. Then we extend Panther to Panther_m and Panther_v . Specially, Panther_m is used to estimate meta-path based path similarity in a heterogeneous information network, and Panther_v is used to achieve the goal of structure similarity search.

We evaluate the efficiency of the proposed methods on a large microblogging “following” network from Tencent.¹ and show the results in table 2. Clearly, our methods are much faster than the comparison methods. Panther_v achieves a $300\times$ speed-up over the fastest comparison method on a Tencent subnetwork of 443,070 vertices and 5,000,000 edges. Our methods are also scalable. Panther is able to return top- k similar vertices for all vertices in a network with 51,640,620 vertices and 1,000,000,000 edges. On average, it only need 0.0001 second to perform a top- k search for each vertex.

We build a prototype system to demonstrate the effectiveness of the method Panther. Specifically, in the system Aminer.org, for each searched expert, we recommend similar authors. Figure 2 shows that similar authors of Professor Jiawei Han are Professor Xifeng Yan, Professor Philip Yu and so on, all of whom share many direct or indirect co-authors with Professor Jiawei Han in the co-author network. All codes and datasets used in this paper are publicly available.²

This article is an extension of prior work [59]. Compared to the prior work, we have the following new contributions: (1) definition of a new similarity metric in a heterogeneous information network, named as meta-path similarity (Section 3); (2) proposal of a sampling algorithm for calculating the new metric efficiently (Section 5); (3) empirical evaluation of effectiveness of the newly proposed algorithm for top- k similarity search in a large medical information network (Section 8.4.3); (4) development of a prototype system of similar expert recommendation in Aminer.org based on the proposed sampling method (Figure 2); (5) adding new motivation of similarity search in heterogeneous information networks in Section 1, unified problem definition of three similarity metrics in Section 3, including path similarity, meta-path similarity and vector similarity, and the classified related work in Section 2.

Organization Section 2 reviews the related work. Section 3 formulates the problem and defines three similarity metrics. In Section 4, we detail the proposed method, for top- k path similarity search, and provide a theoretical analysis. Section 5 and Section 6 introduce the extended methods for top- k meta-path similarity search and vector similarity search respectively. Section 7 compares our proposed methods with existing methods. Section 8 presents experimental results to validate the efficiency and effectiveness of our methods. Finally, Section 9 concludes the paper.

2 RELATED WORK

In this section, we review the neighborhood similarity and structural similarity in homogeneous information networks and heterogeneous information networks respectively.

Neighborhood similarity in homogeneous information networks. Early neighborhood similarity measures, including bibliographical coupling [30] and co-citation [49], are based on the assumption that two vertices are similar if they have many common neighbors. However, they cannot estimate similarity between vertices without common neighbors. A direct consequence is that two nodes with no common neighbors will be treated as not similar at all. Thus, several measures have been proposed to address this problem. For example, Katz [29] counts two vertices as similar if there are more and shorter paths between them. Random walk with restart (RWR) [39] measures the similarity between v_i and v_j as the steady-state probability that v_i will finally walk at v_j . Tsourakakis learns a low-dimension vector for each vertex from the adjacent matrix and calculate similarities between the vectors [53]. Jeh and Widom propose SimRank, which follows a basic recursive intuition that two

¹<http://t.qq.com>

²<https://aminer.org/billboard/Panther>

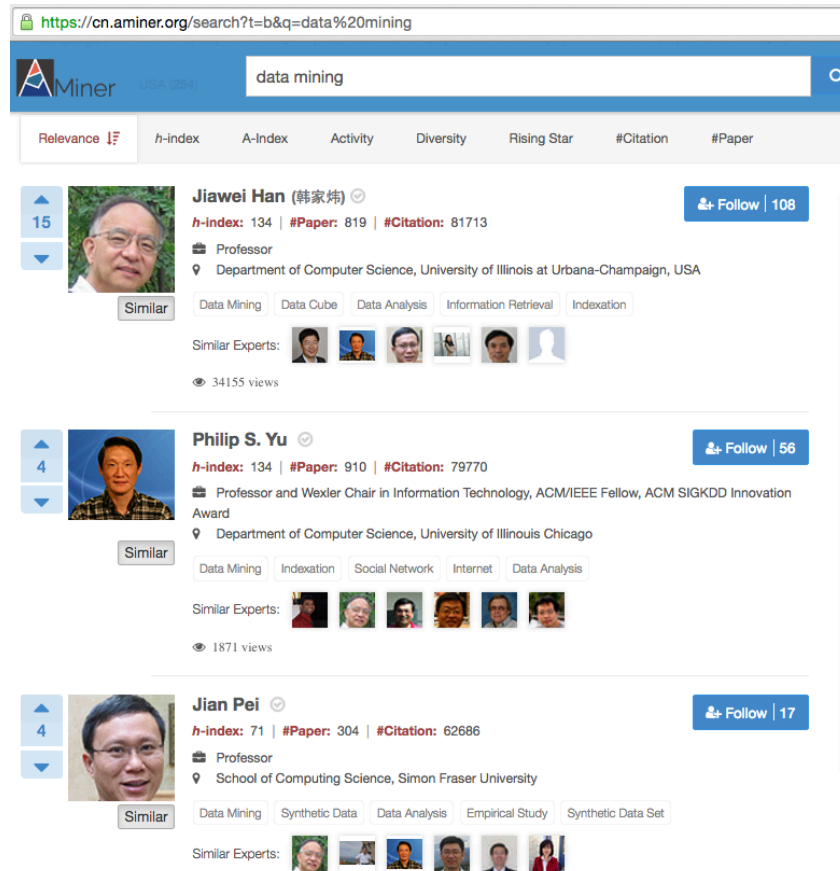


Fig. 2. A prototype system of similar expert recommendation in Aminer.org.

nodes are similar if they are referenced by similar nodes [24]. Leicht et al. develop an asymmetrical version of SimRank named vertex similarity. It is based on the assumption that two vertices are similar if any pair of their neighbors are similar [38]. However, all the SimRank-based methods share a common drawback: their computational complexities are too high. For example, SimRank requires $O(IN^2d^2)$ time and $O(N^2)$ space, where I is the number of iterations, N is the number of vertices and d is the average degree over all vertices. Further studies have been done to reduce the computational complexity of SimRank [32, 33]. Fast-random-walk-based graph similarities, such as in [16, 46, 52], have also been studied. However, the efficiency can still be improved when the networks get larger.

Neighborhood similarity in heterogeneous information networks. Several researchers extend the neighborhood similarity to multi-typed neighborhood similarity in heterogeneous information networks. The state-of-the-art work is proposed by Sun et al., who measure the similarities between vertices by enumerating all the paths following a given meta-path, where the meta-path is used to represent the semantics of a path [50]. Yun et al. solve the problem of similarity join in a heterogeneous information network [58], while this paper targets at solving the problem of top- k similarity search. Shi et al. propose HeteSim to conduct top- k similarity search

in a heterogeneous information network [48]. Their method also suffers from the efficiency issue, and thus the experimental dataset is not big enough. We extend our sampling-based method to solve meta-path based path similarity in a heterogeneous information network efficiently.

Structure similarity in homogeneous information networks. For structure similarity, Blondel et al. provide a HITS-based recursive method to measure similarity between vertices across two different graphs [5]. RoleSim [27] extends SimRank allowing to calculate similarities between disconnected vertices. Similar to SimRank, the computational complexity of the two methods is very high. Feature-based methods can match vertices with similar structures. The basic idea is to define several features for each vertex and then calculate the Euclidean distance between feature vectors of two vertices as their structure similarity. For example, Burt counts 36 kinds of triangles in one's ego network to represent a vertex's structural characteristics [7]. In the same way, vertex centrality, closeness centrality, and betweenness centrality [14] of two different vertices can be compared, to produce a structure similarity measure. However, any of the above metrics is too limited to explain the structural characteristic of a vertex. Aoyama et al. present a fast method to estimate similarity search between objects, instead of vertices in networks [3]. ReFex [19, 20] defines basic features such as degree, the number of within/out-egonet edges, and recursive features as the aggregated values of these features over neighbors. The computational complexity of ReFex depends on the recursive times. Although they use pruning technique to reduce the complexity, no theoretical proof is given to show how many recursive times is enough. More references about feature-based similarity search in networks can be found in the survey [44]. Structure similarity in heterogeneous information networks cannot be directly or easily extended from the sampling-based method, thus will be studied in the future.

Graph Sampling. Our problem of estimating top- K similar vertices is related to estimating frequency of subgraph patterns in a network. Traditional research extensively study how to enumerate the number of subgraph structures in a given graph, such as the number of triangles [25, 35, 41], the count of 4-node subgraphs [26], and the macro frequency and micro frequency of 2,3, and 4-node connected and disconnected subgraphs [45]. A few studies can be generalized to any type of subgraphs. For example, Kashtan et al. [28] and Rahman et al. [42] propose random edge enumeration algorithms to sample different kinds of subgraphs. Wernicke et al. propose a random node enumeration algorithm to uniformly sample different subgraphs [56]. Ahmed et al. propose a general edge sampling framework to estimate the number of triangles, connected paths of length 2, clustering coefficient and so on [2]. The methods that are used to estimate frequency of any type of subgraphs can be used to solve our problem. However, the difference lies in that the weight of a path is also considered when counting paths in our problem. Although Duffield et al. propose a weighted reservoir sampling method [12], however, they solve the problem of aggregation from data streams, and the weight is defined for the probability that discards the keys of aggregated data. The problem is totally different from the problem in this paper, i.e., similarity search in a graph, and thus it is not clear how to directly use the method proposed in that paper to solve our problem. Other work has been conducted on graph stream sampling. For example, Sarma et al. uniformly sample nodes from graph streams to estimate pagerank scores [47]. Buriol et al. [6] and Pavan et al. [40] estimate the number of triangles in graph streams. Cormode use a min-wise hash function to sample edges nearly uniformly to maintain the cascaded summaries of the graph stream [10]. Aggarwal et al. propose a structural reservoir sampling method for structural summarization [1]. Graph stream sampling methods assume the edges arrive as a stream, and the target is to improve space and time complexity for fundamental problems. The graph is assumed to be static in this paper. We will study how to calculate the proposed path similarity in graph streams in the future.




Fig. 3. Illustration of similarity metrics. The vertices in red color are source vertices. Others are in white color. (a) Path similarity: top-2 similar vertices of v_1 are v_3 and v_2 ; (b) Meta-path similarity: following meta-path DDP, the similarity between v_1 and v_5 is 0.5, and is 1.0 following meta-path DPDP, where D represents *Drug* and P represents *Protein*; (c) Vector similarity: v_3 is more similar with v_1 than v_2 .

3 PROBLEM FORMULATION

In this section, we first define the problem of top- k similarity search, and then define three similarity metrics that will be used in solving the problem, i.e., path similarity, and its two extended similarity metrics, meta-path similarity and vector similarity. Path similarity and meta-path similarity are both neighborhood similarity metrics, and vector similarity belongs to the category of structure similarity.

Definition 3.1. Information Network. Let $G = (V, E, W)$ denote a directed network, where V is a set of $|V|$ vertices and $E \subset V \times V$ is a set of $|E|$ edges between vertices. We use $v_i \in V$ to represent a vertex and $e_{ij} \in E$ to represent an edge from vertex v_i to vertex v_j . Let W be a weight matrix, with each element $w_{ij} \in W$ representing the weight associated with edge e_{ij} .

If there exists an edge from vertex v_i to v_j , we also build an inverse edge from v_j to v_i . We use $\mathcal{N}(v_i)$ to indicate the set of neighboring vertices of vertex v_i . Our purpose is to find the top- k similar vertices for any vertex in the network. Precisely, the problem can be defined as, given a network $G = (V, E, W)$ and a query vertex $v \in V$, how to find a set $X_{v,k}$ of k vertices that have the highest similarities to vertex v , where k is a positive integer and the similarity metric is defined as any of the following three ones.

Path similarity. We define a new similarity metric, referred to as *path similarity*. The basic idea of path similarity is that two vertices are similar if they frequently appear on the same paths. The principle is similar to that in Katz [29]. To begin with, we first define T -path as a sequence of vertices $p = (v_1, \dots, v_{T+1})$, which consists of $T + 1$ vertices and T edges.³ Let Π denotes all the T -paths in G and $w(p)$ be the weight of a path p . The weight can be defined in different ways. Given this, the path similarity between v_i and v_j is defined as:

$$S(v_i, v_j) = \frac{\sum_{p \in P(v_i, v_j)} w(p)}{\sum_{p \in \Pi} w(p)}, \quad (1)$$

where $P(v_i, v_j)$ is a subset of Π that contain both v_i and v_j . Figure 3(a) shows an example of path similarity. Considering vertex v_1 , the ranking of the similarities with other vertices is $S(v_1, v_3) > S(v_1, v_2) > S(v_1, v_4) > S(v_1, v_5) > S(v_1, v_6)$. If calculating by SimRank [24], the ranking is $S(v_1, v_5) > S(v_1, v_6) > S(v_1, v_3) > S(v_1, v_2) > S(v_1, v_4)$. We can see among the first-order neighbors, i.e., v_2, v_3 and v_4 , v_1 is most similar to v_3 by both path

³Vertices in the same path are not necessary to be distinct.

similarity and SimRank, because according to path similarity, v_3 shares most paths with v_1 , and according to SimRank, v_3 shares most neighbors with v_1 . Among the second-order neighbors, i.e., v_5 and v_6 , v_1 is more similar to v_5 by both the similarity metrics, and the reason is the same as that of the first-order neighbors. The difference is that SimRank ranks the second-order neighbors before the first-order neighbors. Because according to SimRank, two directly connected vertices always treat each other as two different neighbors, and thus the similarity is weakened. However, path similarity counts paths of all lengths and ranks the first-order neighbors before the second-order neighbors, because first-order neighbors share more paths than the second-order neighbors.

Meta-path similarity. The proposed path similarity is limited in homogeneous information networks, which contains a single type of vertices and edges. Thus path similarity cannot distinguish the semantics among paths that connect two vertices. However, in real world, many networks are inherently heterogeneous, involving multiple types of vertices and edges, such as medical information networks, bibliographic networks, social networks [31] and so on. In those networks, the paths connecting two vertices present different semantics. In many cases, it is interesting to know the semantics of the paths, which may help understand the reasons why two vertices are closely related to each other. To distinguish the semantics among paths, we extend path similarity in homogeneous information networks to meta-path based path similarity in heterogeneous information networks. We first introduce the definitions of heterogeneous information network and meta-path.

Definition 3.2. Heterogeneous information network. A heterogeneous information network can be defined as a multi-typed directed network $G = (V, E, W; \phi, \mathcal{A}, \mathcal{R})$, where V , E and W are the same notations as those in the information directed network. There is a vertex type mapping function $\phi : V \rightarrow \mathcal{A}$ with \mathcal{A} as the set of vertex types, i.e., each vertex $v \in V$ belongs to a particular vertex type in \mathcal{A} . Similarly, there is also an edge type mapping function $\psi : E \rightarrow \mathcal{R}$ with \mathcal{R} as the set of edge types, i.e., each edge $e \in E$ belongs to a particular edge type in \mathcal{R} .

Note that when there is only one vertex type and one edge type, i.e., $|\mathcal{A}| = 1$ and $|\mathcal{R}| = 1$, the network reduces to a homogeneous information network. If there exists an edge type R from vertex type \mathcal{A}_i to \mathcal{A}_j , we also build an inverse relation of R from \mathcal{A}_j to \mathcal{A}_i .

A typical example is a medical information network, with multiple types of vertices such as compounds/drugs, disease, proteins, side effects, pathways and so on. Accordingly, multiple types of directed edges can be defined between different types of vertices. For example, an express relationship from a drug to a protein can be represented by edge $(v_1, v_2) \in E$, where $v_1, v_2 \in V$, $\phi(v_1) = Drug$, $\phi(v_2) = Protein$, and $\psi(v_1, v_2) = Drug \xrightarrow{express} Protein$; symmetrically, the edge (v_2, v_1) represents that the protein v_2 is expressed by the drug v_1 .

Because of the multiple types for vertices and edges, the paths from one vertex to another can also be associated with multiple types. We use the concept meta-path [50] to represent the type/semantics of a path.

Definition 3.3. Meta-path. In a heterogeneous network G , a T -length meta-path is an ordered sequence of T edge types connecting two vertices with type \mathcal{A}_1 and type \mathcal{A}_{T+1} , denoted by $\mathcal{M} = (\mathcal{A}_1 \xrightarrow{R_1} \mathcal{A}_2 \xrightarrow{R_2} \dots \xrightarrow{R_T} \mathcal{A}_{T+1})$, where $\mathcal{A}_i \in \mathcal{A}$ and $\mathcal{R}_i \in \mathcal{R}$. An instantiation of \mathcal{M} is a path in G , denoted by $p = (v_1 v_2 \dots v_{T+1})$, satisfying $\phi(v_i) = \mathcal{A}_i, \forall i = 1, 2, \dots, T + 1$ and $\psi(v_i, v_{i+1}) = \mathcal{R}_i, \forall i = 1, 2, \dots, T$. In addition, we represent a set of path instances following a meta-path as $P_{\mathcal{M}}(v_i, v_j)$.

For example, in a medical information network, $\mathcal{M} = (Drug \xrightarrow{bind} Protein \xrightarrow{bind} Drug \xrightarrow{bind} Protein)$ represents a meta-path, where $Drug, Protein \in \mathcal{A}$ and $Drug \xrightarrow{bind} Protein, Protein \xrightarrow{bind} Drug \in \mathcal{R}$. An instantiation of this meta-path connects a drug and a protein by their directly connected protein and drug.

Given a meta-path $\mathcal{M} = (\mathcal{A}_1 \xrightarrow{R_1} \mathcal{A}_2 \xrightarrow{R_2} \dots \xrightarrow{R_T} \mathcal{A}_{T+1})$, the path space is changed to the T -paths instantiated by meta-path \mathcal{M} and can be denoted as $\Pi_{\mathcal{M}}$. Then meta-path based path similarity, abbreviated as meta-path similarity, from v_i to v_j is defined as

$$S_{\mathcal{M}}(v_i, v_j) = \frac{\sum_{p \in P_{\mathcal{M}}(v_i, v_j)} w(p)}{\sum_{p \in \Pi_{\mathcal{M}}} w(p)}, \quad (2)$$

where $P_{\mathcal{M}}(v_i, v_j)$ is a subset of paths instantiated by meta-path \mathcal{M} , between v_i and v_j .⁴ Figure 3(b) shows an example of meta-path similarity. When considering $Drug \xrightarrow{bind} Drug \xrightarrow{bind} Protein$ (DDP), the top similar vertices of v_1 are v_5 and v_6 , while only v_5 is similar with v_1 when following the meta-path $Drug \xrightarrow{bind} Protein \xrightarrow{bind} Drug \xrightarrow{bind} Protein$ (DPDP). In the figure, the paths with meta-path DDP may be explained as that drug v_1 can bind to the proteins v_5 or v_6 of v_1 's similar drugs v_3 or v_4 , and the paths with meta-path DPDP may be explained as that drug v_1 can bind to protein v_2 which shares another drug v_3 with protein v_5 . Note that actually the above problem is to find the top- k related vertices instead of similar vertices, because the type of the source vertex and the target vertex may not be the same in a heterogeneous information network.

Vector similarity. One limitation of path similarity is that the estimated top- k similar vertices have a bias to close neighbors, though in principle it considers the structural information. We therefore present an extension of path similarity, referred to as vector similarity. The idea is to augment each vertex with a feature vector, which is expected to represent the structural characteristic of the vertex (e.g., star network v.s chain network) as opposed to the identity of neighbors of the vertex (e.g., connect to Bill Gates) [20]. To construct the feature vector, we follow the intuition that the topology structures of two vertices are similar to each other if the probabilities of the two vertices linking to all other vertices are similar to each other [21]. Given a vertex, the probability distribution can be represented by the top- D ranked similarities between it and other vertices, where D is an integer number.

$$\theta(v_i) = (S(v_i, v_{(1)}), S(v_i, v_{(2)}), \dots, S(v_i, v_{(D)})), \quad (3)$$

where $S(v_i, v_{(d)})$ denotes the d -th largest similarity between v_i and another vertex $v_{(d)}$ and the similarity metric can be chosen as any kind of neighborhood similarities.

Naturally, the vector similarity between v_i and v_j can be calculated as the reciprocal Euclidean distance between their feature vectors.

$$S_v(v_i, v_j) = \frac{1}{\|\theta(v_i) - \theta(v_j)\|}. \quad (4)$$

Figure 3(c) shows an example of vector similarity. Based on the definition of vector similarity, we can see that v_3 is more similar with v_1 than v_2 because the structural characteristics of v_3 is more similar with these of v_1 , even though the nodes are in two disconnected networks.

4 PANTHER

We propose Panther to quickly conduct top- k similarity search based on the defined path similarity. A straightforward method to address problem is to first calculate the similarity $S(v_i, v_j)$ between vertex v_i and any other vertex v_j in the network and then select a set $X_{v_i, k}$ of k vertices that have the highest similarities to vertex v_i . However it is in general difficult to scale up to large networks. One important idea is to obtain an approximate

⁴Vertices in the same path are distinct to better present the semantics of a meta-path.

set $X_{v,k}^*$ for each vertex. From the accuracy perspective, we aim to guarantee that the difference between the two sets $X_{v,k}^*$ and $X_{v,k}$ is less than a threshold $\varepsilon \in (0, 1)$, i.e.,

$$\text{Diff}(X_{v,k}^*, X_{v,k}) \leq \varepsilon,$$

with a probability of at least $1 - \delta$.

The difference between $X_{v,k}^*$ and $X_{v,k}$ can be also viewed as the error-bound of the approximation. We propose a sampling-based method to approximate the top- k vertex similarity. In statistics, sampling is a widely used technique to estimate a target distribution [54]. Unlike traditional sampling methods, we propose a random path sampling method, named Panther, to estimate the predefined path similarity. We will explain in details how the method can guarantee the error-bound and how it is able to efficiently achieve the goal.

Random Sampling. To calculate Eq (1), we need to enumerate all T -paths in G . However, the time complexity is exponentially proportional to the path length T , and thus is inefficient when T increases. So we propose a sampling-based method to estimate the path similarity. Since path similarity between two vertices can be cast as estimating the probability that two vertices appear on a same path, our goal is to estimate the probability based on the sampled paths from the whole path space to guarantee a small error-bound with a high probability. Specifically, we randomly sample R paths from the network and recalculate Eq (1) as

$$S(v_i, v_j) = \frac{\sum_{p \in P(v_i, v_j)} w(p)}{\sum_{p \in P} w(p)}, \quad (5)$$

where P is the set of sampled paths.

To generate a path, we randomly select a vertex v_i in G as the starting point, and then conduct random walks of T steps from v_i using t_{ij} as the transition probability from vertex v_i to v_j .

$$t_{ij} = \frac{w_{ij}}{\sum_{v_k \in \mathcal{N}(v_i)} w_{ik}}, \quad (6)$$

where w_{ij} is the weight between v_i and v_j . In an unweighted network, the transition probability can be simplified as $1/|\mathcal{N}(v_i)|$.

We define $w(p)$ based on the random walk theory [13].

$$w(p) = \prod_{i=1, j=i+1}^T t_{ij}.$$

The path weight also represents the probability that a path p is sampled from Π ; thus, $w(p)$ in Eq. (5) is absorbed in the random walk process. So we can rewrite the equation as follows:

$$S(v_i, v_j) = \frac{|P(v_i, v_j)|}{R}, \quad (7)$$

where R is the number of the sampled paths.

Algorithm 2 summarizes the process for generating R random paths. To calculate Eq. (7), the time complexity is $O(RT)$, because it has to enumerate all R paths. To improve the efficiency, we build an inverted index of vertex-to-path [4]. Using the index, we can retrieve all paths that contain a specific vertex v with a complexity of $O(1)$. Then Eq. (7) can be calculated with a complexity of $O(\bar{R}T)$, where \bar{R} is the average number of paths that contain a vertex and is proportional to the average degree \bar{d} . Details of the algorithm are presented in Algorithm 1, where lines 1-5 regard the preprocessing steps, and line 6 refers to the top- k similarity searching for a vertex.

Theoretical Analysis. We give theoretical analysis for the random path sampling algorithm. In general, the path similarity can be viewed as a probability measure defined over all paths Π . Thus we can adopt the results from Vapnik-Chernovenkis (VC) learning theory [54] to analyze the proposed sampling-based algorithm. To begin with, we will introduce some basic definitions and fundamental results from Vapnik-Chernovenkis theory, and then demonstrate how to utilize these concepts and results to analyze our method.

Let $(\mathcal{D}, \mathcal{R})$ be a range space, where \mathcal{D} denotes a domain, and \mathcal{R} is a range set on \mathcal{D} . For any set $B \subseteq \mathcal{D}$, $P_{\mathcal{R}}(B) = \{B \cap A : A \in \mathcal{R}\}$ is the projection of \mathcal{R} on B . If $P_{\mathcal{R}}(B) = 2^B$, where 2^B is the powerset of B , we say that the set B is shattered by \mathcal{R} . The following definitions and theorem derive from [43].

Definition 4.1. The Vapnik-Chervonenkis (VC) dimension of \mathcal{R} , denoted as $VC(\mathcal{R})$, is the maximum cardinality of a subset of \mathcal{D} that can be shattered by \mathcal{R} .

We give an example to explain the concept of VC dimension. For example, if a range set \mathcal{R} is collections of intervals on a line, where each interval classifies the points inside the interval as 1 and those outside the interval as 0, \mathcal{R} can shatter 2 points in a line, but not 3, because no interval can classify the three points on a line as "1 0 1". Thus, $VC(\mathcal{R}) = 2$.

Let $S = \{x_1, \dots, x_n\}$ be a set of i.i.d. random variables sampled according to a distribution ϕ over the domain \mathcal{D} . For a set $A \subseteq \mathcal{D}$, let $\phi(A)$ be the probability that an element sampled from ϕ belongs to A , and let the empirical estimation of $\phi(A)$ on S be

$$\phi_S(A) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}_A(x_i),$$

where $\mathbf{1}_A$ is the indicator function with the value of $\mathbf{1}_A(x)$ equals 1 if $x \in A$, and 0 otherwise.

The question of interest is that how well we can estimate $\phi(A)$ using its unbiased estimator, the empirical estimation $\phi_S(A)$. We first give the goodness of approximation in the following definition.

Definition 4.2. Let \mathcal{R} be a range set on \mathcal{D} , and ϕ be a probability distribution on \mathcal{D} . For $\varepsilon \in (0, 1)$, an ε -approximation to (\mathcal{R}, ϕ) is a set S in \mathcal{D} such that

$$\sup_{A \in \mathcal{R}} |\phi(A) - \phi_S(A)| \leq \varepsilon.$$

One important result of VC theory is that if we can bound the VC-dimension of \mathcal{R} , it is possible to build an ε -approximation by randomly sampling points from the domain according to the distribution ϕ . This is summarized in the following theorem.

THEOREM 4.3. *Let \mathcal{R} be a range set on a domain \mathcal{D} , with $VC(\mathcal{R}) \leq d$, and let ϕ be a distribution on \mathcal{D} . Given $\varepsilon, \delta \in (0, 1)$, let S be a set of $|S|$ points sampled from \mathcal{D} according to ϕ , with*

$$|S| = \frac{c}{\varepsilon^2} \left(d + \ln \frac{1}{\delta} \right),$$

where c is a universal positive constant. Then S is a ε -approximation to (\mathcal{R}, ϕ) with probability of at least $1 - \delta$.

According to the above theory, we set the domain in our problem to be Π —the set of all paths with length T in the graph G . Accordingly, we define the range set \mathcal{R}_G on Π to be

$$\mathcal{R}_G = \{P(v_i, v_j) : v_i, v_j \in V\}.$$

It is a valid range set, since it is the collection of subsets P_{v_i, v_j} of domain Π . We first show an upper bound of the VC dimension of \mathcal{R}_G in Lemma 4.4. The proof is inspired by Riondato and Kornaropoulos [43].

LEMMA 4.4. $VC(\mathcal{R}_G) \leq \log_2 \binom{T+1}{2} + 1$

PROOF. We prove the lemma by contradiction. Assume $VC(\mathcal{R}_G) = l$ and $l > \log_2 \binom{T+1}{2} + 1$. By the definition of VC-dimension, there is a set $Q \subseteq \Pi$ of size l that can be shattered by \mathcal{R}_G . That is, we have the following statement:

$$\forall S_i \subseteq Q, \exists P_i \in \mathcal{R}_G, \text{ s.t. } P_i \cap Q = S_i,$$

where P_i is the i -th range. Since each subset $S_i \subseteq Q$ is different from the other subsets, the corresponding range P_i that making $P_i \cap Q = S_i$ is also different from the other ranges. Moreover, the set Q is shattered by \mathcal{R}_G if and only if $\{P_i \cap Q : P_i \in \mathcal{R}\} = 2^Q$. Thus $\forall p \in Q$, there are 2^{l-1} non-empty distinct subsets $S_1, \dots, S_{2^{l-1}}$ of Q containing the path p . So there are also 2^{l-1} distinct ranges in \mathcal{R}_G that contain the path p , i.e.

$$|\{P_i | p \in P_i \text{ and } P_i \in \mathcal{R}_G\}| = 2^{l-1}.$$

In addition, according to the definition of range set, $\mathcal{R}_G = \{P(v_i, v_j) : v_i, v_j \in V\}$, a path belongs to the ranges corresponding to any pair of vertices in path p , i.e., to the pairwise combinations of the vertices in p . According to the definition, a T -path contains $T + 1$ vertices and some paths may contain a same vertex more than once. Thus the number of ranges in \mathcal{R}_G that p belongs to is no more than the combinatorial number $\binom{T+1}{2}$, i.e.,

$$|\{P_i | p \in P_i \text{ and } P_i \in \mathcal{R}_G\}| \leq \binom{T+1}{2}.$$

On the other hand, from our preliminary assumption, we have $l > \log_2 \binom{T+1}{2} + 1$, which equals to $\binom{T+1}{2} < 2^{l-1}$. Thus,

$$|\{P_i | p \in P_i \text{ and } P_i \in \mathcal{R}_G\}| \leq \binom{T+1}{2} < 2^{l-1}.$$

Hence, we reach a contradiction: it is impossible to have 2^{l-1} distinct ranges $P_i \in \mathcal{R}_G$ containing p . Since there is a one-to-one correspondence between S_i and P_i , we get that it is also impossible to have 2^{l-1} distinct subset $S_i \subseteq Q$ containing p . Therefore, we prove that Q cannot be shattered by \mathcal{R}_G and $VC(\mathcal{R}_G) \leq \log_2 \binom{T+1}{2} + 1$. \square

We now provide theoretical guarantee for the number of sampled paths. How many random paths do we need to achieve an error-bound ε with probability $1 - \delta$? We define a probability distribution on the domain Π . $\forall p \in \Pi$, we define

$$\phi(p) = \text{prob}(p) = \frac{w(p)}{\sum_{p \in \Pi} w(p)}.$$

We can see that the definition of $S(v_i, v_j)$ in Eq.(1) is equivalent to $\phi(P(v_i, v_j))$. This observation enables us to use a sampling-based method (empirical average) to estimate the original path similarity (true probability measure).

Plugging Lemma (4.4) into Theorem (4.3), we obtain:

$$R = \frac{c}{\varepsilon^2} \left(\log_2 \binom{T+1}{2} + 1 + \ln \frac{1}{\delta} \right).$$

That is, with at least R random paths, we can estimate the path similarity between any two vertices with the desired error-bound and confidence level. The above equation also implies that R only depends on the path length T , given an error-bound ε , and a confidence level $1 - \delta$.

ALGORITHM 1: Panther

Input: A network $G = (V, E, W)$, path length T , parameters ϵ, c, δ, k and a vertex v .

Output: top- k similar vertices with regard to v .

Calculate sample size $R = \frac{c}{\epsilon^2} (\log_2 \binom{T+1}{2} + 1 + \ln \frac{1}{\delta})$;

GenerateRandomPath(G, R);

foreach $p_n \in P(v)$ **do**

foreach Unique $v_j \in p_n$ **do**

$S(v, v_j)_+ = \frac{1}{R}$;

end

end

Retrieve top- k similar vertices according to $S(v, v_j)$.

ALGORITHM 2: GenerateRandomPath

Input: A network $G = (V, E, W)$ and sample size R .

Output: Paths $\{p_r\}_{r=1}^R$ and vertex-to-path index $\{P(v_i)\}_{i=1}^N$.

Calculate transition probabilities between every pair of vertices according to Eq. (6) ;

Initialize $r = 0$;

repeat

 Sample a vertex $v = v_i$ uniformly at random ;

 Add v into p_r ;

 Add p_r into the path set of v , i.e., $P(v)$;

repeat

 Randomly sample a neighbor v_j according to transition probabilities from v to its neighbors;

 Set current vertex $v = v_j$;

 Add v into p_r and add p_r into $P(v)$;

until $|p_r| = T + 1$;

$r+ = 1$;

until $r = R$;

5 PANTHER_M

Top- k meta-path similarity search can be explained as: given a heterogeneous information network $G = (V, E, W; \phi, \mathcal{A}, \mathcal{R})$, a meta-path $\mathcal{M} = (\mathcal{A}_1 \xrightarrow{R_1} \mathcal{A}_2 \xrightarrow{R_2} \dots \xrightarrow{R_T} \mathcal{A}_{T+1})$ and a positive integer k , any vertex $v \in V$ with type $\phi(v) = \mathcal{A}_1$, how to retrieve the top- k related vertices of v based on meta-path similarity, i.e., the type of the retrieved vertices is \mathcal{A}_{T+1} and the path instances from v to the retrieved vertices are instantiated by \mathcal{M} .

Similarly, in large scale networks, we need to obtain an approximate set of the related vertices that can guarantee an error bound with a confidence interval. We will introduce how the sampling-based method can efficiently estimate meta-path similarity and guarantee the error-bound.

Random sampling for meta-path similarity. To calculate Eq. (2), we need to enumerate all paths instantiated by meta-path \mathcal{M} from v_i to v_j , which is also inefficient. So we estimate meta-path similarity approximately by sampling meta-path based paths:

$$S_{\mathcal{M}}(v_i, v_j) = \frac{\sum_{p \in P_{\mathcal{M}}(v_i, v_j)} w(p)}{\sum_{p \in P_{\mathcal{M}}} w(p)}, \quad (8)$$

ALGORITHM 3: Panther_m

Input: A network $G = (V, E, W; \phi, \mathcal{A}, \mathcal{R})$, a vertex v , parameters $\varepsilon, c, \delta, k$, meta-path $\mathcal{M} = (\mathcal{A}_1 \xrightarrow{R_1} \mathcal{A}_2 \xrightarrow{R_2} \dots \xrightarrow{R_T} \mathcal{A}_{T+1})$.

Output: top- k related vertices with regard to v .

Calculate transition probabilities between every pair of vertices according to Eq. (9);

Initialize $r = 0$;

repeat

 Sample a vertex $v_c = v_i$ uniformly at random from the nodes with source type \mathcal{A}_1 ;

for $t = 1; t \leq T; t++$ **do**

 Randomly sample a neighbor v_t according to transition probabilities from v_c to its neighbors with type \mathcal{A}_{t+1} ,
 satisfying $\psi(v_c, v_t) = \mathcal{R}_t$;

 Set current vertex $v_c = v_t$;

end

if $t = T$ **then**

$S_{\mathcal{M}}(v_i, v_c) += \frac{1}{R}$;

end

until $r = R$;

Retrieve top- k similar vertices according to $S_{\mathcal{M}}(v, v_j)$;

where $P_{\mathcal{M}}$ is the sampled paths instantiated by \mathcal{M} .

To generate a path instantiated by \mathcal{M} , we randomly select a vertex v_i in G with the type \mathcal{A}_1 as the starting point, and then conduct random walks of T steps from v_i using t_{ij} as the transition probability from node v_i to v_j .

$$t_{i,i+1}^{(\mathcal{R}_i)} = \frac{w_{i,i+1}}{\sum_{v_k \in \mathcal{N}(v_i) \cap \psi(v_i, v_k) = \mathcal{R}_i} w_{ik}}, \quad (9)$$

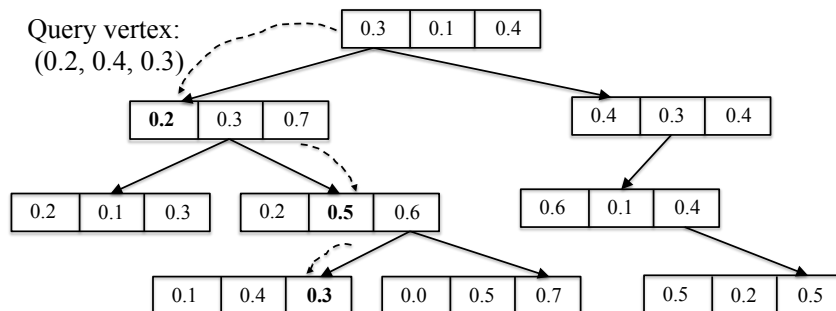
where we restrict the neighbors of v_i into those associated with edge type \mathcal{R}_i in meta-path \mathcal{M} . Accordingly, the weight of a path $w(p_{\mathcal{M}}) = \prod_{i=1}^T t_{i,i+1}^{\mathcal{R}_i}$. Because $w(p_{\mathcal{M}})$ is absorbed in the random walk process, Eq. (8) can be rewritten as follows:

$$S_{\mathcal{M}}(v_i, v_j) = \frac{|P_{\mathcal{M}}(v_i, v_j)|}{R}. \quad (10)$$

Panther_m only considers the paths between two vertices initialized by a given meta-path \mathcal{M} . Thus, we only count the source and target vertices in a path, rather than each distinct vertex in a path. Therefore, we do not need to record the paths and the vertex-to-path index. Instead, when generating random paths, we directly update the similarities between source and target vertices in a same random path. Thus, the space complexity is reduced to $O(|V|\bar{d})$. The updated Algorithm is presented in Algorithm 3.

Theoretical analysis. Theoretical proof is similar to that of Panther. For Panther_m, the domain \mathcal{D} is changed to the paths instantiated by a meta-path \mathcal{M} , i.e., $\Pi_{\mathcal{M}}$. Range set \mathcal{R} is changed to $\mathcal{R}_G = \{P_{\mathcal{M}}(v_i, v_j) : v_i, v_j \in V\}$.

Accordingly, we need to prove the lemma $VC(\mathcal{R}_G) \leq 1$. We also prove the lemma by contradiction. Assume $VC(\mathcal{R}_G) = l$ and $l > 1$. Similar to the proof of Panther, we also know that there are 2^{l-1} distinct ranges contain a path p . In addition, in Panther_m, each path only belongs to the ranges corresponding to the source and target vertices in the path, i.e., the number of ranges in \mathcal{R}_G that p belongs to is equal to 1. However, from the preliminary $l > 1$, we know that $2^{l-1} \neq 1$. Hence, we reach a contradiction and prove the given lemma.

Fig. 4. Data structure of the index built in Panther_v.

The definition of $S_M(v_i, v_j)$ in Eq. (2) is exactly a probability distribution on the domain Π_M . Thus, we can plug $VC(\mathcal{R}_G) \leq 1$ into Theorem (4.3) and obtain: $R = \frac{c}{\epsilon^2} (1 + \ln \frac{1}{\delta})$.

6 PANTHER_V

According to the definition of vector similarity, for each vertex, we construct a vector by top- D ranked similarities between it and other vertices according to Eq. (3). The similarity metric can be chosen as any kind of neighborhood similarity. We determine to use our proposed path similarity, which can be quickly estimated by Panther.

Specifically, for vertex v_i in the network, we first calculate the similarity between v_i and all the other vertices using Panther. Then we construct a feature vector for v_i by taking the largest D similarity scores as feature values.

Finally, the similarity between v_i and v_j is calculated as the reciprocal Euclidean distance between their feature vectors according to Eq. (3).

Index of feature vectors. Again, we use the indexing techniques to improve the algorithm's efficiency. We build a memory based kd-tree [55] index for feature vectors of all vertices. Then given a vertex, we can retrieve top- k vertices in the kd-tree with the least Euclidean distance to the query vertex efficiently. At a high level, a kd-tree is a generalization of a binary search tree that stores points in D -dimensional space. In level h of a kd-tree, given a node v , the $h\%D$ -th element in the vector of each node in its left subtree is less than the $h\%D$ -th element in the vector of v , while the $h\%D$ -th element of every node in the right subtree is no less than the $h\%D$ -th element of v . Figure 4 shows the data structure of the index built in Panther_v. Based on the index, we can query whether a given point is stored in the index very fast. Specifically, given a vertex v , if the root node is v , return the root node. If the first element of v is strictly less than the first element of the root node, look for v in the left subtree, then compare it to the second element of v . Otherwise, check the right subtree. It is worth noting that we can easily replace the kd-tree with any other index methods, such as a r-tree. The algorithms for calculating feature vectors of all vertices and the similarity between vertices are shown in Algorithm 4, where lines 1-8 regard the preprocessing steps, and line 9 refers to the top- k similarity searching for a vertex.

7 COMPARISON WITH EXISTING METHODS

In general, most of the existing methods result in high complexities. For example, the time complexities of SimRank [24], RWR [39] and RoleSim [27] are all proportional to $|V|^2$. TopSim, the top- k version of SimRank, is more efficient, however, is still exponentially proportional to the number of random walk steps. As our method Panther_v, ReFex [19, 20] is also to construct a feature vector for each vertex and then calculate the Euclidean

ALGORITHM 4: Panther_v

Input: A network $G = (V, E, W)$, path length T , parameters $\varepsilon, c, \delta, D, k$ and a vertex v .

Output: top- k similar vertices with regard to v .

Calculate sample size $R = \frac{c}{\varepsilon^2} (\log_2 \binom{T}{2} + 1 + \ln \frac{1}{\delta})$;

GenerateRandomPath(G, R);

foreach $v_i \in V$ **do**

foreach $p_n \in P(v_i)$ **do**

foreach Unique $v_j \in p_n$ **do**

$S(v_i, v_j) += \frac{1}{R}$;

end

end

 Construct a vector $\theta(v_i)$ by taking the largest D values from $\{S(v_i, v_j) : v_j \in p_n \text{ and } p_n \in P(v_i)\}$;

end

Build a kd-tree index based on the Euclidean distance between any vectors $\theta(v_i)$ and $\theta(v_j)$;

Query the top- k similar vertices from the index for v .

Table 1. Time and space complexity for calculating top- k similar vertices for all vertices in a network. I — number of iterations, \bar{d} —average degree, f —feature number, D — vector dimension, T — maximal path length, l — path length, $SP_l(u, v)$ — the set of random walk paths on $G \times G$ of length l ending at the target vertex uv , β — decay factor, \vec{u}_q — steady state probability vector with v_i as the start vertex, \vec{l}_i — restart vector, A — the adjacency matrix of the network, $M_{i,j}$ — a matching between \mathcal{N}_i and \mathcal{N}_j .

Method	Equation	Time Complexity	Space Complexity
SimRank [24]	$S(v_i, v_j) = \frac{\beta}{ \mathcal{N}_i \mathcal{N}_j } \sum_{v_k \in \mathcal{N}_i} \sum_{v_h \in \mathcal{N}_j} S(v_k, v_h)$	$O(I V ^2 \bar{d}^2)$	$O(V ^2)$
TopSim [33]	$S(v_i, v_j) = \sum_{l=1}^T \sum_{p \in SP_l(v_i, v_j)} \beta^l w(p)$	$O(V T \bar{d}^T)$	$O(V + E)$
RWR [39]	$\vec{u}_i = (1 - \beta)A\vec{u}_i + \beta\vec{l}_i$	$O(I V ^2 \bar{d})$	$O(V ^2)$
RoleSim [27]	$S(v_i, v_j) = (1 - \beta) \max_{M_{i,j}} \frac{\sum_{(v_k, v_h) \in M_{i,j}} S(v_k, v_h)}{ \mathcal{N}_i + \mathcal{N}_j - M_{i,j} } + \beta$	$O(I V ^2 \bar{d}^2)$	$O(V ^2)$
ReFex [20]	$S_v(v_i, v_j) = \frac{1}{\ \theta(v_i) - \theta(v_j)\ }$	$O(V + I(f E + V f^2))$	$O(V + E f)$
Panther	$S(v_i, v_j) = \frac{\sum_{p \in P(v_i, v_j)} w(p)}{\sum_{p \in P} w(p)}$	$O(RTc + V \bar{d}T)$	$O(RT + V \bar{d})$
Panther _m	$S_M(v_i, v_j) = \frac{\sum_{p \in P_M(v_i, v_j)} w(p)}{\sum_{p \in P_M} w(p)}$	$O(RTc + V \bar{d}T)$	$O(V \bar{d})$
Panther _v	$S_v(v_i, v_j) = \frac{1}{\ \theta(v_i) - \theta(v_j)\ }$	$O(RTc + V \bar{d}T + V c)$	$O(RT + V \bar{d} + V D)$

distance between vectors as their similarities. However, the complexity of constructing features is determined by the iteration times and is also exponentially proportional to the iteration times. Table 1 shows the time and space complexity of different methods and our methods.

For our method Panther, its time complexity is determined by two main steps. The first step is random path sampling process. The time complexity of generating random paths is $O(RT \log \bar{d})$, where $\log \bar{d}$ is for randomly sampling a neighbor and can be simplified as a small constant c . Hence, the time complexity is $O(RTc)$. The second step is top- k similarity search process. The time complexity of calculating top- k similar vertices for all vertices is $O(|V|\bar{R}T + |V|\bar{M})$. The first part $O(|V|\bar{R}T)$ is the time complexity of calculating Eq. (7) for all pairs of vertices, where \bar{R} is the average number of paths that contain a vertex and is proportional to the average degree \bar{d} . The second part $O(|V|\bar{M})$ is the time complexity of searching top- k similar vertices based on a heap structure,

where \bar{M} represents the average number of co-occurred vertices with a vertex and is proportional to \bar{d} . Hence, the time complexity is $O(|V|\bar{d}T)$. The space complexity for storing paths and vertex-to-path index is $O(RT)$ and $O(|V|\bar{d})$, respectively. The time and space complexity of method Panther_m is the same as that of Panther.

The method Panther_v requires additional computation to build the kd-tree. The time complexity of building a kd-tree is $O(|V|\log|V|)$ and querying top- k similar vertices for any vertex is $O(|V|\log|V|)$, where $\log|V|$ is small and can be viewed as a small constant c . Additional space (with a complexity of $O(|V|D)$) is required to store $|V|$ D -dimension vectors.

8 EXPERIMENTS

8.1 Experimental Setup

In this section, we conduct various experiments to evaluate the proposed methods for top- k similarity search.

Datasets. We evaluate the proposed method on four different networks: Tencent, Twitter, Mobile, and Medicine.

Co-author [51]: The dataset is from AMiner.org⁵ and contains 2,092,356 papers. We extracted a weighted co-author graph from each of the following conferences from 2005 to 2013: KDD, ICDM, SIGIR, CIKM, SIGMOD, ICDE⁶. The weight associated with each edge is the number of papers collaborated by the two connected authors. We use the dataset to evaluate the sampling performance (i.e., the trade-off between the accuracy and efficiency performance) of Panther comparing to alternative sampling algorithms.

Tencent [57]: The dataset is from Tencent Weibo,¹ a popular Twitter-like microblogging service in China, and consists of over 355,591,065 users and 5,958,853,072 “following” relationships. The weight associated with each edge is set as 1.0 uniformly. This is the largest network in our experiments. We use it to evaluate the efficiency performance of Panther and Panther_v comparing to alternative similarity metrics.

Twitter [18, 22]: We collect the dataset by first selecting the most popular user on Twitter, i.e., “Lady Gaga”, and randomly selecting 10,000 of her followers, and then collecting all followers of these users. In the end, we obtained 113,044 users and 468,238 “following” relationships in total. The weight associated with each edge is also set as 1.0 uniformly. We use this dataset to evaluate the accuracy of Panther and Panther_v comparing to alternative similarity metrics.

Mobile [11]: We build the mobile network using call records from a mobile communication company within two weeks by treating each user as a vertex, and the communication between users as an edge. The resultant network consists of 194,526 vertices and 206,934 edges, with the weight associated with each edge as the number of calls. We also use this dataset to evaluate the accuracy of Panther and Panther_v comparing to alternative similarity metrics.

Medicine [9]: The dataset is a heterogeneous medical information network, which consists of 295,897 nodes and 727,931 edges. There are 9 types of vertices in the networks, including 258,030 drugs, 22,056 proteins, and so on. 12 types of edges are built between these vertex types. Please refer to [9] for details of the dataset. We use the dataset to evaluate the efficiency and accuracy of Panther_m.

Evaluation aspects. To quantitatively evaluate the proposed methods, we consider the following performance measurements:

Sampling performance: We apply our method Panther to the co-author networks to evaluate the sampling performance, i.e., the trade-off between accuracy and efficiency performance.

Efficiency performance: We apply our methods to the Tencent network to evaluate the computational time.

⁵<http://aminer.org/citation>

⁶Numbers of vertices/edges of different conferences are: KDD: 2,867/7,637, ICDM: 2,607/4,774, SIGIR: 2,851/6,354, CIKM: 3,548/7,076, SIGMOD: 2,616/8,304, ICDE: 2,559/6,668.

Accuracy performance: We use the number of common neighbors as ground truth to evaluate Panther on Twitter and Mobile networks. We apply Panther_v to find top- k structural hole spanners in Twitter and Mobile networks. We evaluate Panther_m using the task of link prediction on a medical information network.

Parameter sensitivity analysis: We analyze the sensitivity of different parameters in our methods: path length T , vector dimension D , and error-bound ϵ .

All codes are implemented in C++ and compiled using GCC 4.8.2 with -O3 flag. The experiments were conducted on a Ubuntu server with four Intel Xeon(R) CPU E5-4650 (2.70GHz) and 1T RAM.

Comparison methods. We compare with the following methods:

Edge Sampling [2]: Samples a sub-network by selecting each edge in the original network by probability p , and then enumerates all the T -paths in the sampled sub-network. Based on the sampled paths, we calculate path similarity according to Eq. (7).

Node Sampling [56]: Samples a path by iteratively selecting a neighbor of each vertex by probability q , until the length of a path is T . Based on the sampled paths, we calculate path similarity according to Eq. (7).

RWR [39]: Starts from v_i , iteratively walks to its neighbors with the probability proportional to their edge weights. At each step, it also has some probability to walk back to v_i (set as 0.1). The similarity between v_i and v_j is defined as the steady-state probability that v_i will finally reach at v_j . We calculate RWR scores between all pairs and then search the top- k similar vertices for each vertex.

TopSim [33]: Extends SimRank [24] on one graph G to finding top- k authoritative vertices on the product graph $G \times G$ efficiently.

RoleSim [27]: Refines SimRank [24] by changing the average similarity of all neighbor pairs to all matched neighbor pairs. We estimate RoleSim scores between all pairs and select the top- k similar vertices for each vertex.

ReFeX [20]: Defines local, egonet, and recursive features to capture structural characteristics. The local feature is the vertex degree. Egonet features include the number of within-egonet edges and the number of out-egonet edges. For weighted networks, they contain weighted versions of each feature. Recursive features are defined as the mean and sum value of each local or egonet feature among all neighbors of a vertex. In our experiments, we only extract recursive features once and construct a vector for each vertex by a total of 18 features. For fair comparison, to search top- k similar vertices, we also build the same kd-tree as that used in our method.

SLAP [9]: Is proposed to predict links in a medical information network. Given a pair of source and target vertices in a network, it employs a heap-based Dijkstra algorithm to find all the paths between the two vertices and sums up the scores of all the paths together. The score for each path is a combination of the path weight (i.e., the multiplication of the transition probabilities in the path) and the significance score of the meta-path the path belongs to (i.e., the expected mean and standard deviation of the weights of the paths that belong to the meta-path).

The codes of TopSim, RoleSim, ReFeX, and SLAP are provided by their authors. We used the fast versions of TopSim and RoleSim mentioned in their paper.

Implementation notes. In our experiments, we empirically set the parameters as follows: $c = 0.5$, $\delta = 0.1$, $T = 4$, $D = 50$ and $\epsilon = \sqrt{1/|E|}$. The optimal values of T , D and ϵ are discussed in § 8.5. We build the kd-tree using the toolkit ANN.⁷

⁷<http://www.cs.umd.edu/~mount/ANN/>

8.2 Sampling Performance

We use Precision@Top- k , where we vary k as 5 and 10, as the metric to measure how far the approximate estimation by Panther, is from the exact estimation, Eq. (1). Specifically, for each vertex v_i , given an approximate estimation of the top- k similar vertices, \widetilde{TS}_i , and the exact estimation of the top- k similar vertices, TS_i , Precision@Top- k is defined as:

$$\text{Precision@Top-}k = \frac{\sum_{i=1}^{|V|} |\widetilde{TS}_i \cap TS_i|}{k|V|} \quad (11)$$

To quantitatively evaluate the sampling performance of Panther, we compare Panther with two alternative sampling methods, Edge Sampling method and Node Sampling method. The comparison methods use the same way as Panther to calculate path similarity except that different methods sample paths in different ways. We show how Precision@Top- k and speed vary by varying the parameters in different methods. The experimental setting is as follows.

- We tune the parameters. For Panther, we try the error-bound ε as all the possible values in the range $\{0.0006, 0.001, 0.003, 0.006, 0.01, 0.03, 0.06, 0.1, 0.3, 0.6\}$; for Edge Sampling method, we try the edge sampling probability p as all the possible values in the range $\{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$; for Node Sampling method, we try the node sampling probability q with same values of p .
- For each configuration of the parameters ε , p , and q , we calculate Precision@Top- k over 10 independent sampling results.
- For each configuration of the parameters ε , p , and q , we record the average CPU time of 10 independent sampling process, and divide them by the CPU time of the exact estimation (i.e., Sampling-time/Exact-time).

Figure 5 and Figure 6 plot all the (Precision@Top- k , Sampling-time/Exact-time) pairs when varying the parameters ε , p , and q on 6 different co-author networks. From the results, we can see that for all the three sampling methods, Precision@Top- k increases quickly at the beginning, and then becomes almost stable. Precision@Top- k of Panther increases dramatically at the beginning and takes much shorter time to obtain a higher accuracy than the other two methods, thus performs best. This is because Panther actually absorbs the path weight (Eq. (6)) in the random sampling process, which exactly reserves the effect of path weight defined in path similarity (Eq. (1)), while Edge Sampling and Node Sampling methods do not consider path weight in their sampling processes. The experimental results demonstrate the superiority of the proposed sampling method.

8.3 Efficiency and Scalability Performance

In this subsection, we first fix $k = 5$, and evaluate the efficiency and scalability performance of different comparison methods using the Tencent dataset. We randomly extract different (large and small) versions of the Tencent networks. For TopSim and RoleSim, we only show the computational time for similarity search. For ReFeX, Edge Sampling, Node Sampling, Panther, and Panther _{v} , we also show the computational time used for preprocessing. Since Tencent is a homogeneous network, we show the time efficiency of Panther _{m} on a heterogeneous medical information network in § 8.4. Sampling probability p for Edge Sampling is set as 0.6, and q for Node Sampling is set as 0.3 according to the experimental results on co-author, Twitter, and Mobile datasets.

Table 2 lists statistics of the different Tencent sub-networks and the efficiency performance of the comparison methods. Clearly, our methods (both Panther and Panther _{v}) are much faster than the comparison methods. For example, on the Tencent6 sub-network, which consists of 443,070 vertices and 5,000,000 edges, Panther _{v} achieves a 390× speed-up, compared to the fastest (ReFeX) of all the comparative methods.

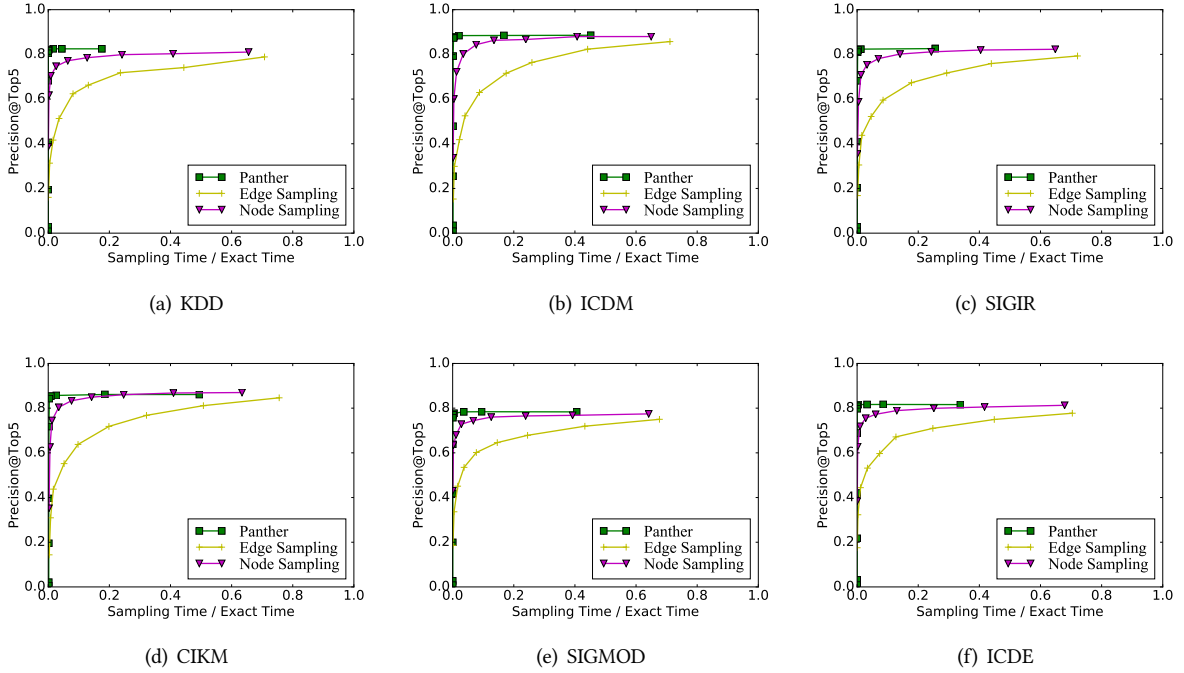


Fig. 5. The trade-off between accuracy (i.e., Precision@Top-5 at y-axis) and efficiency (i.e., Sampling-time/Exact-time at x-axis) performance by varying error-bound ε of Panther, edge sampling probability p of Edge Sampling method, and node sampling probability q of Node Sampling method on the co-author networks.

We can also see that RWR, TopSim and RoleSim cannot complete top- k similarity search for all vertices within a reasonable time when the number of edges increases to 500,000. ReFeX can deal with larger networks, but also fails when the edge number increases to 10,000,000. According to the results shown in Figure 5, 6, and 7, when setting sampling probability p for Edge Sampling as 0.6, and q for Node Sampling as 0.3, the accuracy performance can approach that of Panther, but according to the results in Table 2, Panther is much more efficient than the two baseline sampling methods. Our methods can scale up to handle very large networks with more than 10,000,000 edges. On average, Panther only needs 0.0001 second to perform top- k similarity search for each vertex in a large network.

8.4 Accuracy Performance

8.4.1 Performance of Panther. We evaluate how Panther can approximate the similarity based on common neighbors. The evaluation procedure is described as follows:

- (1) For each vertex u in the seed set S , generate top k vertices $\text{Top}_{A,k}(u)$ that are the most similar to u by the algorithm A .
- (2) For each vertex $v \in \text{Top}_{A,k}(u)$, calculate $g(u, v)$, where g is a coarse similarity measure defined as the ground truth. Define $f_{A,k} = \sum_u \sum_v g(u, v)$.
- (3) Let $f_{R,k}$ denotes the result of a random algorithm.

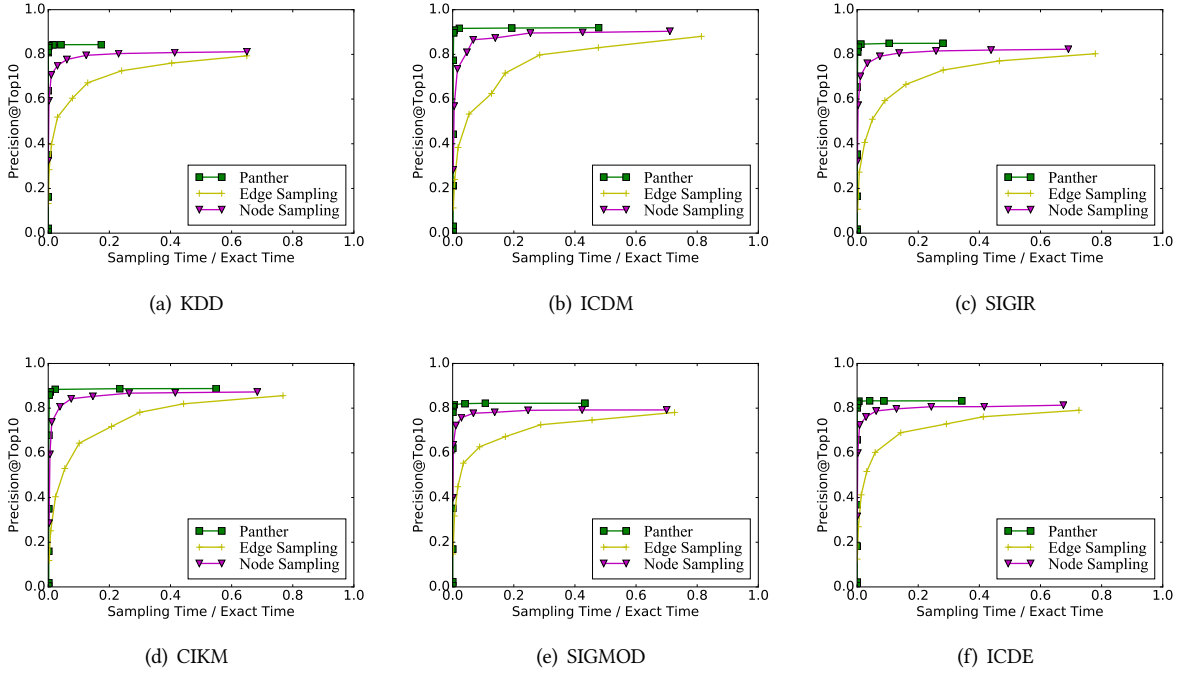


Fig. 6. The trade-off between accuracy (i.e., Precision@Top-10 at y-axis) and efficiency (i.e., Sampling-time/Exact-time at x-axis) performance by varying error-bound ε of Panther, edge sampling probability p of Edge Sampling method, and node sampling probability q of Node Sampling method on the co-author networks.

Table 2. Efficiency performance (CPU time) of comparison methods on different sizes of the Tencent sub-networks. The time before “+” denotes the time used for processing and the time after “+” denotes that used for top- k similarity search. “—” indicates that the corresponding algorithm cannot finish the computation within a reasonable time.

Sub-network	$ V $	$ E $	RWR	TopSim	RoleSim	ReFeX	Edge Sampling	Node Sampling	Panther	Panther _v
Tencent1	6,523	10,000	+7.79hr	+28.58m	+37.26s	3.85s+0.07s	14.16m+1.68hr	58.82s+1.12m	0.07s+0.26s	0.99s+0.21s
Tencent2	25,844	50,000	+>150hr	+11.20hr	+12.98m	26.09s+0.40s	—	14.22m+26.14m	0.28s+1.53s	2.45s+4.21s
Tencent3	48,837	100,000	—	+30.94hr	+1.06hr	2.02m+0.57s	—	1.50hr+1.67hr	0.58s+ 3.48s	5.30s+5.96s
Tencent4	169,209	500,000	—	+>120hr	+>72 hr	17.18m+2.51s	—	—	8.19s+16.08s	27.94s+24.17s
Tencent5	230,103	1,000,000	—	—	—	31.50m+3.29s	—	—	15.31s+30.63s	49.83s+22.86s
Tencent6	443,070	5,000,000	—	—	—	24.15hr+8.55s	—	—	50.91s+2.82m	4.01m+1.29m
Tencent7	702,049	10,000,000	—	—	—	>48hr	—	—	2.21m+6.24m	8.60m+6.58m
Tencent8	2,767,344	50,000,000	—	—	—	—	—	—	15.78m+1.36hr	1.60hr+2.17hr
Tencent9	5,355,507	100,000,000	—	—	—	—	—	—	44.09m +4.50hr	5.61hr +6.47hr
Tencent10	26,033,969	500,000,000	—	—	—	—	—	—	4.82hr +25.01hr	32.90hr +47.34hr
Tencent11	51,640,620	1,000,000,000	—	—	—	—	—	—	13.32hr +80.38hr	98.15hr +120.01hr

(4) Finally, we define the score for algorithm A as $\text{score}(A, k) = \frac{f_{A,k} - f_{R,k}}{|S| \times k}$, which represents the improvement of algorithm A over a random method.

Specifically, we define $g(u, v)$ as the number of common neighbors between u and v on each dataset.

Figure 7 shows the performance of Panther evaluated on the ground truth of common neighbors on Twitter and Mobile networks. Some baselines such as RWR and RoleSim are ignored on the datasets, because they cannot

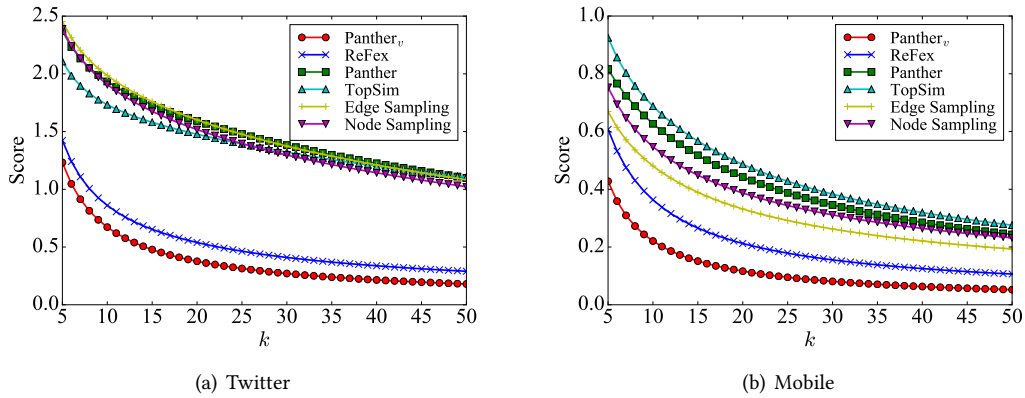


Fig. 7. Performance of approximating common neighbors on the Twitter and Mobile networks.

complete top- k similarity search for all vertices within a reasonable time. It can be seen that Panther performs better than any other methods on both the datasets. Panther and ReFex perform worst as they are not devised to address the similarity between near vertices. Our method Panther performs as good as TopSim, the top- k version of SimRank, and the two baseline sampling methods, Edge Sampling ($q = 0.6$) and Node Sampling ($p = 0.3$), because they are all based on the principle that two vertices are considered structurally equivalent if they have many common neighbors in a network. However, according to our previous analysis, TopSim, Edge Sampling and Node Sampling perform much slower than Panther.

8.4.2 Performance of Panther_v. We consider an application of Top- k structural hole spanner finding to evaluate the performance of Panther_v. The theory of structural holes [8] suggests that, in social networks, individuals would benefit from filling the “holes” between people or groups that are otherwise disconnected. The problem of finding top- k structural hole spanners was proposed in [36], which also shows that 1% of users who span structural holes control 25% of the information diffusion (retweeting) in Twitter.

Structural hole spanners are not necessarily connected, but they share the same structural patterns such as local clustering coefficient and centrality. Thus, the idea here is to feed a few seed users to Panther_v, and use it to find other structural hole spanners. For evaluation, we use network constraint [8] to obtain the structural hole spanners in Twitter and Mobile, and use this as the ground truth. Then we apply different methods—Panther_v, ReFex, Panther, TopSim, Edge Sampling ($p = 0.6$), and Node Sampling ($q = 0.3$)—to retrieve top- k similar users for each seed. If an algorithm can find another structural hole spanner in the top- k returned results, it makes a correct search. We define $g(u, v) = 1$, if both u and v are structural hole spanners, and $g(u, v) = 0$ otherwise.

Figure 8 shows the performance of comparison methods for finding structural hole spanners in different networks. Panther_v achieves a consistently better performance than the comparison methods by varying the value of k . TopSim, Panther, and the two baseline sampling methods seem inapplicable to this task. This is reasonable, as the underlying principle of them is to find vertices with more connections to the query vertex.

8.4.3 Performance of Panther_m. We use the same task and dataset in [9] to evaluate the efficiency and effectiveness performance of Panther_m. The task is to predict links between drugs and proteins. In a medical information network, due to the investigation difficulty and latency in clinical field experiments, few reliable links are constructed and thus the existing links in the network are usually quite sparse. For example, the OMIM dataset used in [37] contains 3,209 diseases and 8755 genes, but only 3,954 gene-disease associations. Therefore,

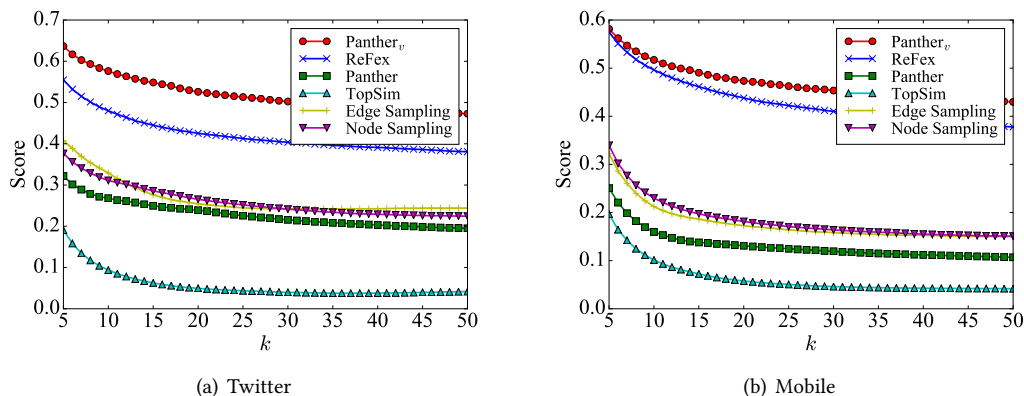


Fig. 8. Performance of mining structural hole spanners on the Twitter and Mobile networks.

automatically mining and discovering the potential associations between medical entities can effectively enrich the links between medical entities, and meanwhile provide useful hypothesis for new clinical experiments.

Predicting links in a medical information network has been extensively studied. Some research builds links between two entities two-hops away from each other based on the number of their common neighbors[17]. Other research work further extends the two-hop away associations to more distant associations. For example, Chen et al. [9] consider the paths with different lengths between two entities and restrict the maximal length of paths as 3. However, their method is inefficient, taking about several days to finish on such a dataset.

Setting. We collect the ground truth from a public database PubChem BioAssay,⁸ which contains bioactivity screens of chemical substances described in PubChem Substance. For example, there are active or inactive bioactivities between some drugs and proteins in the database. We first choose the pairs of drugs and proteins with activity records as candidate links, and then treat the active interactive pairs as positive links and the inactive ones as negative links. Finally, we remove the links existing in the medical information network. The resulting ground truth contains in total 36,254 positive links and 343,043 negative links.

For each drug in the ground truth, we use $Panther_m$ to estimate top- k similar vertices, and treat the links between the drug and its top- k similar proteins as the predicted positive links, and others as negative links. Specifically, we first enumerate all kinds of meta-paths with length no longer than 4. The paths longer than 4 are considered too long to predict significant links. And then for each drug, we estimate top- k similar vertices following each selected meta-path, and aggregate the similarities of different meta-paths together. Please refer to [9] for details of the meta-paths. We compare with the baseline SLAP [9] and our method Panther.

Results. Figure 9 shows the ROC curves and AUC of comparison methods for predicting links in a medical information network. In the result, k is set as 500 and error-bound ϵ is set as 1×10^{-4} . We see that SLAP performs better than $Panther_m$. The results are reasonable, because SLAP is an exact solution, which enumerate all the possible paths between two vertices, while $Panther_m$ only samples a subset in the whole path space. However, $Panther_m$ takes only 1 hour, while SLAP takes about 58 hours. $Panther_m$ achieves a more than 50X speed-up, compared to SLAP. We also conduct Panther in the network. The results show that Panther performs much worse than $Panther_m$, because Panther ignores the type of paths, which may introduce additional noises.

⁸<http://www.ncbi.nlm.nih.gov/pcassay/>

We then reduce the value of error-bound ϵ to investigate whether the performance of Panther_m can be improved. The results are shown in Figure 10(a). We see that with the decreasing of error-bound, the performance of Panther_m approaches to that of SLAP. When error-bound is reduced to 6×10^{-5} , AUC is improved to 0.66, with a time cost of 2.5 hours, which still achieves 20X speed-up, compared to SLAP.

Furthermore, we change the value of k to see whether the performance is sensitive to k . In this experiment, we fix ϵ as 1×10^{-4} . Figure 10(b) shows AUC when changing the value of k from 1 to 1000. We can see that AUC increases and then becomes stable when k is 500. This indicates that some positive instances in the ground truth are not ranked quite high in the top similar list. Thus more positive instances can be recalled when increasing k . We further study several vertices ranked before the positive instances in the ground truth. Figure 11 shows a case study of the paths between protein AURKB and drug 5312137,⁹ where AURKB is ranked 46th by our algorithm and is included in the ground truth. Figure 1(b) presents the paths between protein MET and drug 5312137, where MET is ranked 8th by our algorithm but is not in the ground truth. We plot the paths with length no longer than 3 from drug 5312137 to the predicted proteins. The number of paths from drug 5312137 to protein MET is 42, which is more than the 24 paths from drug 5312137 to protein AURKB. Moreover, from the figure we can see that there are three meta-paths, i.e., $Drug \xrightarrow{bind} Protein \xrightarrow{bind} Drug \xrightarrow{bind} Protein$, $Drug \xrightarrow{bind} Protein \xrightarrow{bind} Protein \xrightarrow{bind} Protein$ and $Drug \xrightarrow{bind} Protein \xrightarrow{bind} Pathway \xrightarrow{bind} Protein$ from drug 5312137 to protein MET, where only one meta-path, i.e., $Drug \xrightarrow{bind} Protein \xrightarrow{bind} Drug \xrightarrow{bind} Protein$ from drug 5312137 to protein AURKB. The meta-paths from drug 5312137 to protein MET are richer than those to protein AURKB. The results indicate that the proteins ranked before the positive instances in the ground truth are more closely related to the source drug, although many top ranked proteins are not included in the ground truth.

Finally, we compare the performance of different meta-paths. Specifically, we set ϵ as 1×10^{-4} . For each meta-path, we evaluate its top- k ($k = 500$) similar vertices based on the ground truth. Table 3 shows 10 meta-path examples, with the evaluated scores of AUC no less than 0.5. We can see that meta-path $Drug \xrightarrow{bind} Protein \xrightarrow{bind} Drug \xrightarrow{bind} Protein$ is most significant compared with all the other meta-paths. The meta-path may be explained as a drug can bind to a protein which shares another compound/drug with another protein. Most of the meta-paths are actually insignificant and can be ignored when predicting links.



8.5 Parameter Sensitivity Analysis

We now discuss how different parameters influence the accuracy performance of our methods.

Effect of Path Length T . Figure 12 shows the accuracy performance of Panther evaluated by common neighbors by varying T as 1, 4, 9, 19, 49 and 99, i.e., varying the number vertices in a path as 2, 5, 10, 20, 50, 100. A too small $T (< 5)$ would result in inferior performance and when increasing its value up to 5, it almost becomes stable.

Effect of Vector Dimension D . Figure 13 shows the accuracy performance of Panther_v for mining structural hole spanners by varying the vector dimension D as 2, 5, 10, 20, 50 and 100. Generally speaking, the performance gets better when D increases and it remains the same after D gets larger than 50. This is reasonable, as Panther_v reflects the distribution of a vertex linking to the other vertices. Thus, the higher the vector dimension is, the better the representation will be. While the performance will be stable when the dimension exceeds a threshold.

⁹The URIs are <http://chem2bio2rdf.org/uniprot/resource/gene/MET>, <http://chem2bio2rdf.org/uniprot/resource/gene/AURKB> and http://chem2bio2rdf.org/pubchem/resource/pubchem_compound/5312137.

Fig. 9. ROC curve and AUC of Panther_m.Fig. 10. AUC of Panther_m by varying (a) ε and (b) k .

Effect of Error Bound ε . Figure 14 shows the accuracy performance of Panther evaluated by common neighbors and that of Panther_v for mining structural hole spanners on the Tencent networks with different scales by varying error-bound ε from 0.06 to 0.0001. We see that when the ratio $\frac{|E|}{(1/\varepsilon)^2}$ ranges from 5 to 20, scores of Panther are almost convergent on all the datasets. And when the ratio $\frac{|E|}{(1/\varepsilon)^2}$ ranges from 0.2 to 5, the scores of Panther_v are almost convergent on all the datasets. Thus we can reach the conclusion that the value of $(1/\varepsilon)^2$ is almost linearly positively correlated with the number of edges in a network. Therefore we can empirically set $\varepsilon = \sqrt{1/|E|}$.

9 CONCLUSION

In this paper, we propose a sampling method to quickly conduct top- k similarity search on large networks. The algorithm is based on the idea of *random path*, and is to solve neighborhood similarity in homogeneous networks.

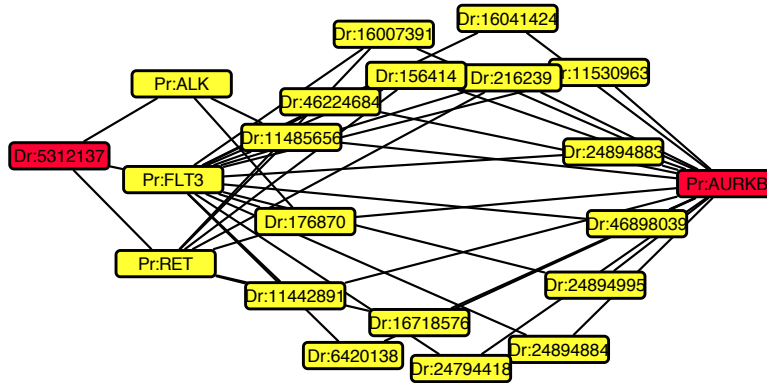


Fig. 11. The paths between drug 5312137 and protein AURKB. The vertices in red color are source and target vertices. The prefix "Dr" in the vertex label denotes Drug, "Pr" denotes Protein, and "Pa" denotes Pathway. .

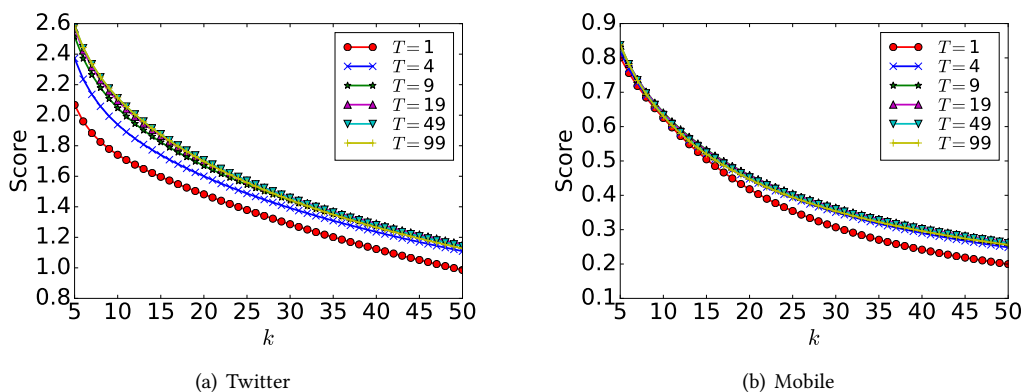
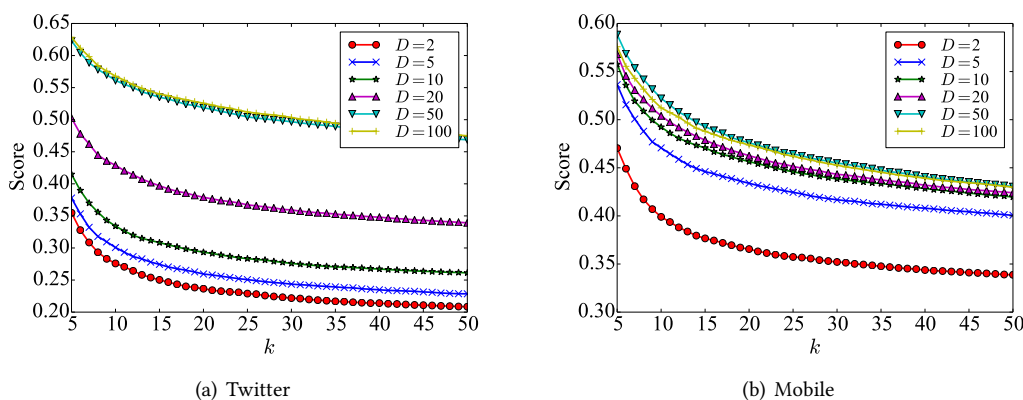
Table 3. AUC of different Meta-paths.

Meta-path	AUC
$Drug \xrightarrow{bind} Protein \xrightarrow{bind} Drug \xrightarrow{bind} Protein$	0.64
$Drug \xrightarrow{bind} Protein \xrightarrow{bind} GeneOntology \xrightarrow{bind} Protein$	0.57
$Drug \xrightarrow{bind} Protein \xrightarrow{bind} Protein \xrightarrow{bind} Protein$	0.56
$Drug \xrightarrow{bind} Protein \xrightarrow{bind} Pathway \xrightarrow{bind} Protein$	0.53
$Drug \xrightarrow{bind} Substructure \xrightarrow{bind} Drug \xrightarrow{bind} Protein$	0.51
$Drug \xrightarrow{bind} Substructure \xrightarrow{bind} Drug \xrightarrow{bind} Protein$	0.51
$Drug \xrightarrow{bind} Substructure \xrightarrow{bind} Drug \xrightarrow{bind} Protein$	0.51
$Drug \xrightarrow{bind} Protein \xrightarrow{bind} Tissue \xrightarrow{bind} Protein$	0.50
$Drug \xrightarrow{bind} Protein \xrightarrow{bind} Tissue \xrightarrow{bind} Protein$	0.50
$Drug \xrightarrow{bind} ChemicalOntology \xrightarrow{bind} Drug \xrightarrow{bind} Protein$	0.50

One extended method is to solve neighborhood similarity in heterogeneous networks, and another extended method is presented to enhance the structure similarity when two vertices are completely disconnected. We provide theoretical proofs for the error-bound and confidence of the proposed algorithm. We perform an extensive empirical study and show that our algorithm can obtain top- k similar vertices for any vertex in a network approximately $300\times$ faster than state-of-the-art methods. We also build a prototype system of recommending similar authors to demonstrate the effectiveness of our proposed method.

ACKNOWLEDGMENTS

The authors thank Pei Lee, Laks V.S. Lakshmanan, Jeffrey Xu Yu; Ruoming Jin, Victor E. Lee, Hui Xiong; Keith Henderson, Brian Gallagher, Lei Li, Leman Akoglu, Tina Eliassi-Rad, Christos Faloutsos for sharing codes of the comparison methods. We thank Tina Eliassi-Rad for sharing the datasets. We thank Gang Fu, Bin Chen, Ying

Fig. 12. Effect of T on the performance of Panther.Fig. 13. Effect of D on the performance of Panther_v.

Ding and Abhik Seal for sharing the datasets and baseline results. The work is supported by the National Basic Research Program of China (2014CB340506, 2014CB340402), National Natural Science Foundation of China (61631013, 61561130160, 61532021), National Social Science Foundation of China (13&ZD190), National Key Research&Develop Plan (2016YFB1000702), a research fund supported by MSRA, the Royal Society-Newton Advanced Fellowship Award, and the Fundamental Research Funds for the Central Universities, and the Research Funds of Renmin University of China(15XNLQ06, 17XNLF09).

REFERENCES

- [1] Charu C Aggarwal, Yuchen Zhao, and S Yu Philip. 2011. Outlier detection in graph streams. In *ICDE'11*. 399–409.
- [2] Nesreen K. Ahmed, Nick Duffield, Jennifer Neville, and Ramana Kompella. 2014. Graph Sample and Hold: A Framework for Big-graph Analytics. In *KDD'14*. 1446–1455.
- [3] Kazuo Aoyama, Kazumi Saito, Hiroshi Sawada, and Naonori Ueda. 2011. Fast approximate similarity search based on degree-reduced neighborhood graphs. In *KDD'11*. 1055–1063.
- [4] Ricardo Baeza-Yates, Berthier Ribeiro-Neto, and others. 1999. *Modern information retrieval*. Vol. 463. ACM press.
- [5] Vincent D Blondel, Anahí Gajardo, Maureen Heymans, Pierre Senellart, and Paul Van Dooren. 2004. A measure of similarity between




Fig. 14. Effect of ϵ on the performance of Panther and Panther_v on different sizes of Tencent networks.

- graph vertices: Applications to synonym extraction and web searching. *SIAM Rev.* 46, 4 (2004), 647–666.
- [6] Luciana S Buriol, Gereon Frahling, Stefano Leonardi, Alberto Marchetti-Spaccamela, and Christian Sohler. 2006. Counting triangles in data streams. In *PODS'06*. 253–262.
- [7] Ronald S Burt. 1990. Detecting role equivalence. *Social Networks* 12, 1 (1990), 83–97.
- [8] Ronald S Burt. 2009. *Structural holes: The social structure of competition*. Harvard university press.
- [9] Bin Chen, Ying Ding, and David J Wild. 2012. Assessing drug target association using semantic linked data. *PLoS Computational Biology* 8, 7 (2012), e1002574.
- [10] Graham Cormode and S Muthukrishnan. 2005. Space efficient mining of multigraph streams. In *PODS'05*. 271–282.
- [11] Yuxiao Dong, Yang Yang, Jie Tang, Yang Yang, and Nitesh V. Chawla. 2014. Inferring User Demographics and Social Strategies in Mobile Social Networks. In *KDD'14*. 15–24.
- [12] Nick Duffield, Yunhong Xu, Liangzhen Xia, Nesreen Ahmed, and Minlan Yu. 2017. Stream Aggregation Through Order Sampling. *arXiv preprint arXiv:1703.02693* (2017).
- [13] William Feller. 2008. *An introduction to probability theory and its applications*. Vol. 2. John Wiley & Sons.
- [14] Linton C Freeman. 1977. A set of measures of centrality based on betweenness. *Sociometry* (1977), 35–41.
- [15] Linton C Freeman. 1979. Centrality in social networks conceptual clarification. *Social Networks* 1, 3 (1979), 215–239.
- [16] Yasuhiro Fujiwara, Makoto Nakatsuji, Hiroaki Shiokawa, Takeshi Mishima, and Makoto Onizuka. 2013. Efficient Ad-hoc Search for Personalized PageRank. In *SIGMOD'13*. 445–456.
- [17] Kwang-II Goh, Michael E Cusick, David Valle, Barton Childs, Marc Vidal, and Albert-László Barabási. 2007. The human disease network. *PNAS* 104, 21 (2007), 8685–8690.
- [18] Wentao Han, Xiaowei Zhu, Ziyang Zhu, Wenguang Chen, Weimin Zheng, and Jianguo Lu. 2016. A comparative analysis on Weibo and Twitter. *Tsinghua Science and Technology* 21, 1 (2016), 1–16.
- [19] Keith Henderson, Brian Gallagher, Tina Eliassi-Rad, Hanghang Tong, Sugato Basu, Leman Akoglu, Danai Koutra, Christos Faloutsos, and Lei Li. 2012. Rolx: structural role extraction & mining in large graphs. In *KDD'12*. 1231–1239.
- [20] Keith Henderson, Brian Gallagher, Lei Li, Leman Akoglu, Tina Eliassi-Rad, Hanghang Tong, and Christos Faloutsos. 2011. It's who you know: graph mining using recursive structural features. In *KDD'11*. 663–671.
- [21] Paul W Holland and Samuel Leinhardt. 1981. An exponential family of probability distributions for directed graphs. *J. Amer. Statist. Assoc.* 76, 373 (1981), 33–50.
- [22] John Hopcroft, Tiancheng Lou, and Jie Tang. 2011. Who Will Follow You Back? Reciprocal Relationship Prediction. In *CIKM'11*. 1137–1146.
- [23] Paul Jaccard. 1901. Étude Comparative de la Distribution Florale dans une Portion de Alpes et du Jura. *Bulletin de la Société Vaudoise des Sciences Naturelles* 37 (1901), 547–579.
- [24] Glen Jeh and Jennifer Widom. 2002. SimRank: a measure of structural-context similarity. In *KDD'02*. 538–543.
- [25] Madhav Jha, C Seshadhri, and Ali Pinar. 2013. A space efficient streaming algorithm for triangle counting using the birthday paradox. In *KDD'13*. 589–597.
- [26] Madhav Jha, C Seshadhri, and Ali Pinar. 2015. Path Sampling: A Fast and Provable Method for Estimating 4-Vertex Subgraph Counts.

- In *WWW'15*. 495–505.
- [27] Ruoming Jin, Victor E Lee, and Hui Hong. 2011. Axiomatic ranking of network role similarity. In *KDD'11*. 922–930.
- [28] Nadav Kashtan, Shalev Itzkovitz, Ron Milo, and Uri Alon. 2004. Efficient sampling algorithm for estimating subgraph concentrations and detecting network motifs. *Bioinformatics* 20, 11 (2004), 1746–1758.
- [29] Leo Katz. 1953. A new status index derived from sociometric analysis. *Psychometrika* 18, 1 (1953), 39–43.
- [30] M. M. Kessler. 1963. Bibliographic coupling between scientific papers. *American Documentation* 14, 1 (1963), 10–25.
- [31] Angelos Kremyzas, Norman Jaklin, and Roland Geraerts. 2016. Towards social behavior in virtual-agent navigation. *Science China Information Sciences* 59, 11 (2016), 112102.
- [32] Mitsuru Kusumoto, Takanori Maehara, and Ken-ichi Kawarabayashi. 2014. Scalable similarity search for SimRank. In *SIGMOD'14*. 325–336.
- [33] Pei Lee, Laks VS Lakshmanan, and Jeffrey Xu Yu. 2012. On top-k structural similarity search. In *ICDE'12*. 774–785.
- [34] EA Leicht, Petter Holme, and Mark EJ Newman. 2006. Vertex similarity in networks. *Physical Review E* 73, 2 (2006), 026120.
- [35] Yongsu Lim and U Kang. 2015. Mascot: Memory-efficient and accurate sampling for counting local triangles in graph streams. In *SIGKDD'15*. 685–694.
- [36] Tiancheng Lou and Jie Tang. 2013. Mining Structural Hole Spanners Through Information Diffusion in Social Networks. In *WWW'13*. 837–848.
- [37] Nagarajan Natarajan and Inderjit S Dhillon. 2014. Inductive matrix completion for predicting gene–disease associations. *Bioinformatics* 30, 12 (2014), i60–i68.
- [38] Mark EJ Newman. 2006. Finding community structure in networks using the eigenvectors of matrices. *Physical Review E* 74, 3 (2006), 036104.
- [39] Jia-Yu Pan, Hyung-Jeong Yang, Christos Faloutsos, and Pinar Duygulu. 2004. Automatic multimedia cross-modal correlation discovery. In *KDD'04*. 653–658.
- [40] Aduri Pavan, Kanat Tangwongsan, Srikanta Tirthapura, and Kun-Lung Wu. 2013. Counting and sampling triangles from a graph stream. *VLDB* 6, 14 (2013), 1870–1881.
- [41] Mahmudur Rahman and Mohammad Al Hasan. 2013. Approximate triangle counting algorithms on multi-cores. In *Big Data, 2013 IEEE International Conference on*. 127–133.
- [42] Mahmudur Rahman, Mansurul Bhuiyan, and Mohammad Al Hasan. 2012. Graft: an approximate graphlet counting algorithm for large graph analysis. In *CIKM'12*. 1467–1471.
- [43] Matteo Riondato and Evgenios M. Kornaropoulos. 2014. Fast Approximation of Betweenness Centrality Through Sampling. In *WSDM'14*. 413–422.
- [44] Ryan A Rossi and Nesreen K Ahmed. 2015. Role Discovery in Networks. *IEEE TKDE* (2015).
- [45] Ryan A Rossi, Rong Zhou, and Nesreen K Ahmed. 2017. Estimation of Graphlet Statistics. *arXiv preprint arXiv:1701.01772* (2017).
- [46] Purnamrita Sarkar and Andrew W. Moore. 2010. Fast Nearest-neighbor Search in Disk-resident Graphs. In *KDD'10*. 513–522.
- [47] Atish Das Sarma, Sreenivas Gollapudi, and Rina Panigrahy. 2011. Estimating pagerank on graph streams. *Journal of the ACM (JACM)* 58, 3 (2011), 13.
- [48] Chuan Shi, Xiangnan Kong, Yue Huang, and Philip S. Yu. 2014. Hetesim: A general framework for relevance measure in heterogeneous networks. *TKDE'14* 26, 10 (2014), 2479–2492.
- [49] Henry Small. 1973. Co-citation in the scientific literature: A new measure of the relationship between two documents. *Journal of the American Society for Information Science* 24, 4 (1973), 265–269.
- [50] Yizhou Sun, Jiawei Han, Xifeng Yan, Philip S Yu, and Tianyi Wu. 2011. Pathsim: Meta path-based top-k similarity search in heterogeneous information networks. *VLDB'11* (2011), 992–1003.
- [51] Jie Tang, Jing Zhang, Limin Yao, Juanzi Li, Li Zhang, and Zhong Su. 2008. ArnetMiner: Extraction and Mining of Academic Social Networks. In *KDD'08*. 990–998.
- [52] Hanghang Tong, Christos Faloutsos, and Jia-Yu Pan. 2006. Fast random walk with restart and its applications. In *ICDM'06*. 613–622.
- [53] Charalampos E Tsourakakis. 2014. Toward Quantifying Vertex Similarity in Networks. *Internet Mathematics* 10, 3-4 (2014), 263–286.
- [54] Vladimir N Vapnik and A Ya Chervonenkis. 1971. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications* 16, 2 (1971), 264–280.
- [55] Ingo Wald and Vlastimil Havran. 2006. On building fast kd-trees for ray tracing, and on doing that in $O(N \log N)$. In *Interactive Ray Tracing 2006, IEEE Symposium on*. 61–69.
- [56] Sebastian Wernicke. 2006. Efficient Detection of Network Motifs. *TCBB* 3, 4 (2006), 347–359.
- [57] Yang Yang, Jie Tang, Cane Wing-ki Leung, Yizhou Sun, Qicong Chen, Juanzi Li, and Qiang Yang. 2014. RAIN: Social Role-Aware Information Diffusion. In *AAAI'14*. 367–373.
- [58] Xiong Yun, Yangyong Zhu, and S. Yu Philip. 2015. Top-k similarity join in heterogeneous information networks. *TKDE'15* 27, 6 (2015), 1710–1723.
- [59] Jing Zhang, Jie Tang, Cong Ma, Hanghang Tong, Yu Jing, and Juanzi Li. 2015. Panther: Fast top-k similarity search on large networks.

39:30 • J. Zhang et al.

In *KDD'15*. 1445–1454.

Received February 2007; revised March 2009; accepted June 2009