

Chapter 1

Link Prediction in Highly Fractional Data Sets

Michael Fire, Rami Puzis, and Yuval Elovici

Abstract Extremist organizations all over the world increasingly use online social networks as a communication media for recruitment and planning. As such, online social networks are also a source of information utilized by intelligence and counter terror organizations investigating the relationships between suspected individuals. Unfortunately, the data mined from open sources is usually far from being complete due to the efforts of suspected and known terrorists to hide their relationships. One of the methods used to uncover missing information in social networks is referred to as link prediction. We use link prediction methods solely based on network structure analysis to infer hidden relationships among individuals and investigate their effectiveness in fractional datasets. Experiments performed on a number of closed communities extracted from organizational and public social networks show that structural link prediction retains its effectiveness even when large parts of the original social network are hidden.

1.1 Introduction

In recent years, online social networks have grown in scale and variability and offer individuals with similar interests the possibility of exchanging ideas and networking. On the one hand, social networks create new opportunities to develop friendships, share ideas, and conduct business. On the other hand, they are also an effective media tool for plotting crime and organizing extremists groups around the

Michael Fire and Yuval Elovici

Department of Information Systems Engineering, Ben-Gurion University of the Negev and Telekom Innovation Laboratories at Ben-Gurion University of the Negev, e-mail: {elovici, mickyfi}@bgu.ac.il

Rami Puzis

University of Maryland Institute for Advanced Computer Studies (UMIACS), e-mail: puzis@umd.edu

world. Online social networks, such as Facebook, Google+, and Twitter are hard to track due to their massive scale and increased awareness of privacy. Criminals and terrorists strive to hide their relationships, especially those that can associate them with a executed terror act.

A large portion of recent research in social network analysis has been targeted at identifying these hidden relationships in social networks. These research endeavours are usually referred to as *link prediction* methods. These methods are able to detect existing social ties that have not been established in a particular social network [1, 2, 3, 4, 5, 6]. In the security and counter-terrorism domains link prediction can assist in identifying hidden groups of terrorists or criminals [2]. However, link prediction is also useful for civil applications, such as friend-suggestion mechanisms embedded in online social networks. For example, in bioinformatics, link prediction can be used to find interactions between proteins [7], and in e-commerce, it can help build recommendation systems [8].

Several different methods for solving the link prediction problem have been proposed in recent years. These days, the majority of solutions are based on supervised machine learning methods, such as Bayesian probabilistic models, relational Bayesian networks, and linear algebraic methods. Further details on these approaches can be found in a thorough survey written by Zaki and Hasan [9]. We briefly describe a selected few in Section 1.2.2.

We focus on link prediction methods based on machine learning classifiers trained on a set of topological features. We use a set of well-known features, such as connectivity features, intersection and union of the friends groups, *Jaccard's-Coefficient* [10], *Preferential-Attachment* score [2, 6], and the *Friends-Measure*, introduced in [11]. The latter is a variation of the Katz measure [12] and estimates how well the friends of two users know each other.

Many link prediction methods are applied to large social networks where most of the links are assumed to be known. This is a reasonable assumption when the main objective is to predict ties that are not yet established in the social network. However, terrorist social networks mined from open sources are typically small and very partial due to the efforts of their subjects to obfuscate their activity.

In this study, we investigate the effects of dataset partiality on the effectiveness of link prediction by gradually reducing the number of visible links in the test networks. The effectiveness of classifiers was evaluated using six different types of data sets: a group of Facebook users sharing the same employer, a researchers community from Academia.edu [13], Friends and Family SMS messages social network [14], Students Cooperation Network [15], AnyBeat social network ¹, and the Profiles in Terror (PIT) dataset [16]. Information on these social networks is presented in Section 1.3. Evaluation results presented in Section 1.5 demonstrate that classification quality (in terms of Area Under the ROC Curve (AUC)) degrades with the number of visible links. Nevertheless, even a small fraction of visible links (5%-20%) helps in solving the link prediction problem with chances significantly higher than random.

¹ <http://www.anybeat.com>

The remainder of the paper is organized as follows. In Section 1.2, we review previous studies on terrorist’s social networks and link prediction. In Section 1.3, we describe the datasets used in this study. Experimental setup and the features extracted from the structure of social networks are described in Section 1.4.1. We present the experimental results in Section 1.5, and conclusions in Section 1.6.

1.2 Background

1.2.1 *Social Networks of Terrorists*

Over the last two decades, social networks have been studied fairly extensively in the general context of analyzing interactions between people and determining the important structural patterns of such interactions [17]. In the previous decade, even before September 11, 2001, social network analysis was recognized as a tool for fighting the war against criminal organizations in an age where there is no well-defined enemy with a formal hierarchical organization [18]. Moreover, after the September 11, 2001 events, social network analysis became a well-known mainstream tool to help the fight against terror [19].

Several studies have analyzed terror organization social networks based on graph structural features. In the winter of 2002, Valdis [20] studied Al-Qaeda’s network structural properties by collecting publicly available data on the Al-Qaeda hijackers. Rothenberg [21] conjectured on the structure of the al Qaeda network based on public media sources. After the Madrid bombing, in March 11, 2004, Rodriguez [22] used public sources to construct and study the terrorists’ network. He showed that the terror organization network included mainly weak ties that are hard to detect. In 2004, Sageman [23] used various public sources, mostly records of trials, to collect and analyze 400 terrorist biographies. He discovered that 88 percent of the terrorists had friendship or family bonds to the Jihad. In 2005, Basu [24] studied terrorists’ organization in India. He used social network analysis, such as the betweenness measure, to identify major groups of terrorists and key players. In 2010, Wiil et al. [25] studied a recent Denmark terror plan. By using data mining techniques, they were able to construct, from public sources, the social network of David Coleman Headley, one of the terror plan conspirators.

Attempts to reconstruct the social networks of terrorists requires a significant effort spent on mining the Web for publicly available information and free text analysis. This typically results in the ability to obtain small networks only with a high likelihood of missing information. In this study, we attempt to predict links inside social networks where a substantial amount of the network’s links data were missing. A similar idea was studied by Dombroski et al. [26]. Their study examined the possibilities of using the inherent structures observed in social networks to make predictions of networks using limited and missing information.

1.2.2 Link Prediction

In this study, we focus on link prediction methods based on supervised machine learning algorithms. These methods were first introduced by Liben-Nowell and Kleinberg in 2003 [1], who used graph topological features in a study on five co-authorship networks, each containing several thousands of vertices. In 2006, Hasan et al. [2] increased the scale of analyzed networks to hundreds of thousands of nodes by analyzing the DBLP and BIOBASE co-authorship networks. Supervised learning was also applied by many other researchers to solve the link prediction problems, for example in [3, 4, 27]. Initially, the proposed link prediction solutions were tested on bibliographic or on co-authorship data sets [1, 2, 3, 4].

In 2009, Song et al. used matrix factorization to estimate the structural similarity between profiles in online social network services, such as Facebook and MySpace [5]. In 2011, the IJCNN social network challenge [28] inspired several publications on link prediction using topological network analysis. These publications proposed and evaluated different methods for predicting links in social networks. Narayanan et al. won the challenge by using a method that combined machine-learning algorithms with de-anonymization [29]. Cukierski et al. [6] took second place by extracting 94 distinct features for each one of the of several thousands of vertex pairs in the training data and analyzing it with the Random Forest algorithm. Recently, Fire et al. presented a method for predicting links inside communities; their methods used supervised learning ensemble classifiers constructed by using small training sets only which consisted of several hundreds of examples [13].

When no information on the social network besides its structure is available, it is crucial to define and calculate the features that are as informative as possible. On the one hand, large networks containing millions of vertices and links pose a scalability challenge and require the use of easy to compute topological features extracted from the neighbourhoods of the tested vertices. For example, Facebook has more than 845 million registered users and each month many new users are added [30]. On the other hand, small networks, such as the terrorist datasets available today (e.g., [20, 31]), pose a different challenge. They contain too few links and vertices to construct a large training set. Moreover, these networks are much more prone to noise than their huge counterparts because the existence or absence of every link may significantly change the values of the extracted features. In this study, we take the latter challenge to the extreme and evaluate structural link prediction methods while gradually removing random links from organizational, group, and terrorist affiliation networks.

1.3 Social Network Datasets

In this study, we apply link prediction classifiers to six labeled social network datasets (see Table 1.1), namely, Profile in Terror (PIT) [31], AnyBeat network², a group of Facebook³ users, Academia.edu [11, 32], Friends and Family study [33], and Students' Cooperation Social Network [15].

Profile in Terror (PIT) [16] is a data set that captures intelligence information extracted from publicly available sources collected by the MIND Lab at UMD⁴. This data set contains 851 labeled relationships among terrorists and was previously used in multi-label link classification in [31]. Each relationship can have one or more labels from: colleague, congregate, contact, or family. In this study, we disregard the relationship labels and treat this data set as a flat network which contains 840 links and 244 vertices. Figure 1.1 depicts the graph topology⁵.




Fig. 1.1 Profile in Terror social network

AnyBeat "AnyBeat is an online community; a public gathering place where you can interact with people from around your neighborhood or across the world"¹. AnyBeat is a relatively new social network in which members can log in without using their real name and members can follow any other member in the network. In this study, we evaluated our algorithm on a major part of the network's topology,

² <http://www.anybeat.com>

³ <http://www.facebook.com>

⁴ <http://www.mindswap.org/>

⁵ All the social networks figures in this paper were created by Cytoscape software [34]

which was obtained using a dedicated web crawler. The topology contained 12,645 vertices and 67,053 links (see Figure 1.2). Among the networks studied in this study, AnyBeat is outstanding in its size. We included a network of more than ten thousand vertices in order to shade the results of the study and focused mainly on small partially visible social networks. As presented in Section 1.5, link prediction appears to be significantly easier for this network.




Fig. 1.2 AnyBeat social network

A Facebook group of co-workers (Co-Workers). Facebook is a website and social networking service that was launched in February 2004. As of December 2011, Facebook has more than 845 million registered users [30]. Facebook users may create a personal profile, add other users as friends, and interact with other members. Friendship ties in Facebook are reciprocal, therefore, we refer to the underlying group’s social network as undirected. We extracted a community of co-workers who, according to their Facebook profile, worked for the same well-known high-tech company. The graph representing the co-workers’ community network contains 165 vertices and 726 links and was obtained using a web crawler in the beginning of February 2012 (see Figure 1.3).

Ivy League University (Researchers). Academia.edu is a platform for academics to share and follow research underway in a particular field or discipline [35]. Members upload and share their papers with other researchers in over 100,000 fields and categories. An Academia social network member may choose to follow any of the network’s members; hence, the directed nature of the links within this network. We evaluated our classifiers for a small community of researchers who, according to their Academia.edu profiles, belonged to the same Ivy League University. The




Fig. 1.3 Facebook coworker community social network

researchers' community network graph contained 207 nodes and 702 links (see Figure 1.4) and was obtained using a web crawler.




Fig. 1.4 Academia.edu researchers community social network

Friends and Family (F&F). The Friends and Family dataset contains rich data signals gathered from the smart-phones of 140 adult members of a young-family residential community. The data were collected over the course of one year [33]. We evaluated our classifiers on a social network that was constructed based on SMS messages sent and received by the members. The SMS messages social network directed graph contained 103 nodes and 281 links (see Figure 1.5).





Fig. 1.5 Friends and Family SMS messages social network

Students' Cooperation Social Network (Students). The students' cooperation social network was constructed from data collected during a "Computer and Network Security" course; a mandatory course taught by two of this paper's authors at Ben-Gurion University [15]. The social network contains data collected from 185 participating students from two different departments. The course's social network was created by analyzing the implicit and explicit cooperation among the students while doing their homework assignments. The Students' Cooperation graph contained 185 nodes and 311 links (see Figure 1.6).

**Fig. 1.6** Students' Cooperation Social Network**Table 1.1** Datasets

Network	Directed	Vertices	Links	Positive ex- amples	Visible net- works	Features
PIT	No	244	840	420	200	8
AnyBeat	Yes	12,645	67,053	25,000	40	12
Co-Workers	No	165	722	361	200	8
Researchers	Yes	207	702	351	200	12
F&F	Yes	103	281	140	200	12
Students	No	185	311	155	200	8

1.4 Methods and Experiments

1.4.1 Experimental Setup

The goal of structural link prediction is to identify a set of hidden links within a social network's structure by analyzing the topology of the known (visible) network. As a first step in our experiments, we preprocessed each one of the social networks described in Section 1.3. The preprocessing process is slightly different for the AnyBeat networks due to its size.

We will refer to the unmodified social networks as *original* in the rest of this paper. We randomly removed a portion of links from each original network to generate a *visible* network. The number of removed links ranged from 0% to 95% in steps of

5%. For each one of these twenty fractions we chose several random sets of removed links that spawn several different visible networks of the same size.; ten random sets for the small networks and two random sets for the AnyBeat network. Overall, 200 *visible* networks were created for each one of the small *original* networks and 40 *visible* networks were created for AnyBeat. Our goal was to predict the links in the *original* network by extracting features from the *visible* network topology.

Next, in order to evaluate the machine learning classifiers, we generated training and testing data sets from each visible network with the original network treated as the ground truth. Note that due to privacy profiles and imperfect data acquirement methods, the original networks may have missing links as well. We will disregard these links and assume acquaintances between individuals that are linked together in the original networks only.

Every instance in the dataset represents a possible candidate link. The target attribute of each instance is binary, indicating the existence or absence of a link in the original social network. A set of structural features was extracted from the *visible* part of the corresponding social networks. As opposed to the 54 features discussed in [32], we have only seven and eleven of the most efficient features for undirected and directed networks, respectively in this study. These features are briefly described in Section 1.4.2.

In order to create a balanced data set for training the classifiers, we included the same number of positive and negative examples. Positive examples are vertex pairs that are connected in the *original* network. Negative examples include the same number of vertex pairs, but this time, random pairs of vertices not connected in the *original* network. The method for selecting the negative examples in this study corresponds to the *easy* dataset, as described in [32]. Note that the features describing each positive or negative example were extracted from the structure of the *visible* networks. Table 1.1 summarizes the datasets used in this study.

WEKA [36], a popular suite of machine learning software written in Java and developed at the University of Waikato, New Zealand, was used as the machine learning platform for this study. We used a WEKA’s J48 classifier with ten minimum objects that showed effective performance in past experiments. We evaluated our results using a 10-folds cross validation method.

Next, we describe the set of features extracted from the visible social network graphs. The features used in this study are a subset of a more extensive set of features investigated in [32].

1.4.2 Feature Extraction

This section describes the features extracted from the social network structure in order to build our link prediction classifiers. The extracted features are based primarily on the Friends-features subset, as suggested by Fire et al. [32].

Let $G = \langle V, E \rangle$ be the graph representing the structure of a social network. Links in the graph are denoted by $e = (u, v) \in E$ where $u, v \in V$ are vertices in the

graph. Our goal is to construct classifiers capable of computing the likelihood of $(u, v) \in E$ or $(u, v) \notin E$ for every two vertices $u, v \in V$. To achieve this goal, we extracted the following features for each pair, (u, v) , in our datasets.

1. **Vertex degree:** Let $v \in V$ be some vertex, we can define the neighborhood of v by:

$$\Gamma(v) := \{u | (u, v) \in E \text{ or } (v, u) \in E\}$$

If G is a directed graph, we can also define the following neighborhoods:

$$\Gamma_{in}(v) := \{u | (u, v) \in E\}$$

$$\Gamma_{out}(v) := \{u | (v, u) \in E\}$$

We define the degree features for directed and undirected networks as the sizes of the respective neighborhoods:

$$degree(v) := |\Gamma(v)| \quad (1.1)$$

$$degree_{in}(v) := |\Gamma_{in}(v)| \quad (1.2)$$

$$degree_{out}(v) := |\Gamma_{out}(v)| \quad (1.3)$$

The degree features measure the number of friends v has. If we look at a directed graph, such as Academia.edu, the meaning of the degree feature is how many other members of the community v follows (out-degree), and how many members of the community follow v (in-degree).

2. **Common-Friends:** Let $u, v \in V$ be a pair of vertices in the network; we define the common friends of u and v to be all the vertices in the network that are friends of both u and v . Formally, common friends is the size of the intersection of the respective neighborhoods:

$$Common-Friends(u, v) := |\Gamma(v) \cap \Gamma(u)| \quad (1.4)$$

The Common-Friends feature was widely used in previous works for predicting links in different datasets [1, 4, 5, 6, 8, 11].

3. **Total-Friends:** Let $u, v \in V$ be a pair of vertices; we can define the number of distinct friends of u and v as the size of the union of the respective neighborhoods:

$$Total-Friends(u, v) := |\Gamma(u) \cup \Gamma(v)| \quad (1.5)$$

4. **Jaccard's-Coefficient:** Jaccard's-Coefficient is a well-known feature for link prediction [1, 4, 5, 6, 8, 11]. This feature, which measures the similarity among sets of nodes, is defined as the size of the intersection divided by the size of the union of the sample sets:

$$Jaccard's-Coefficient(u, v) := \frac{|\Gamma(u) \cap \Gamma(v)|}{|\Gamma(u) \cup \Gamma(v)|} \quad (1.6)$$

In our approach, this measure indicates whether two community members have a significant number of common friends regardless of their total number of friends. A higher value of the Jaccard's-Coefficient typically indicates a stronger connection between two nodes in the network.

5. **Preferential-Attachment-Score:** The Preferential-Attachment-Score indicates the likelihood of a new link to be formed between the vertices u and v , according to the preferential attachment model [37]. It is defined as the multiplication of the number of friends of u and v .

$$\text{Preferential-Attachment-Score}(u, v) := |\Gamma(u)| \cdot |\Gamma(v)| \quad (1.7)$$

The Preferential-Attachment score was used many times in past research, for examples see [2, 6, 11].

6. **Friends-Measure:** Let $u, v \in V$ be two vertices; the Friends-Measure of u and v is the extent to which their friends are interconnected. The higher the number of connections between u and v 's friends, the greater the chance that u and v know each other.

$$\text{Friends-Measure}(u, v) := \sum_{x \in \Gamma(u)} \sum_{y \in \Gamma(v)} \delta(x, y) \quad (1.8)$$

where $\delta(x, y)$ is defined as:

$$\delta(x, y) := \begin{cases} 1 & \text{if } x = y \text{ or } (x, y) \in E \text{ or } (y, x) \in E \\ 0 & \text{otherwise} \end{cases}$$

The Friends-Measure was first presented by Fire et al. [11].

1.5 Results

For each undirected and directed visible network, we extracted seven and eleven features, respectively, (see Section 1.4.2) and evaluated the specified machine-learning algorithms (see Section 1.4.1) using a 10-fold cross-validation approach. The AUC results of the J48 algorithm on all the networks where no edges were removed from the networks are presented in Figure 1.7.

Table 1.2 and Figure 1.8 present the AUC results of the J48 algorithm for each social network with various numbers of removed links. As expected, we noticed a degradation in the AUC values as more edges are removed from the network. The slope of the AUC as a function of the visible network size increases with the number of removed vertices. In the terrorists affiliation network (PIT), the AUC remains above 0.9, even when features in the dataset are computed according to 40% of the original links.

Another notable observation is that the classification results (with the exception of the Students' Cooperation network) are significantly above random, even when only 5%-10% of the *original* links are *visible* with respect to the computation of




Fig. 1.7 AUC Results - J48 with no removed links.

structural features. The Students' Cooperation network was based on links between students in a single class taught for a short period of time. Therefore, the network did not contain as many cliques as in the other social networks. We believe that the existence of large cliques, (as in PIT network) or densely connected parts in the network, makes it easier to predict ties randomly removed from the network.

A final interesting observation is that the best AUC values were obtained for the AnyBeat network. Link prediction in this network is much easier than in the smaller networks when significant portions of the network are removed. The AUC remains above 0.9, even when there are only 5% of visible links. Note that 5% of visible links in AnyBeat is 3,352, which is a few orders of magnitude higher than in the other networks.

Table 1.2 Effectiveness of link prediction in fractional networks

Network	Percent of visible links				
	100%	80%	60%	40%	20%
PIT	0.985	0.980	0.966	0.940	0.789
Co-Workers	0.893	0.858	0.820	0.791	0.738
Researchers	0.952	0.942	0.920	0.895	0.799
F&F	0.941	0.928	0.908	0.877	0.744
Students	0.926	0.877	0.841	0.762	0.629
AnyBeat	0.982	0.979	0.972	0.966	0.954

To obtain an indication of the usefulness of the various features, we analyzed their importance using Weka's Information Gain attribute selection algorithm. The InfoGain results for different datasets are presented in Tables 1.3 and 1.4. In order to check the effect of link removal on the usefulness of various features, we com-




Fig. 1.8 AUC as a function of the fraction of visible links for various social networks.

puted the InfoGain values for the original networks and for networks with only 5% of visible links. The results indicate that InfoGain of all features drops and the set of useful features changes when random links are removed from the dataset. Only the *Preferential-Attachment-Score* is useful for both collections of datasets. Nonetheless, the InfoGain values of connectivity degree features drop in Table 1.4 and they become more useful relative to the other features. This effect is even more apparent for the *Total-friends* feature that has the highest InfoGain value in Table 1.4. Yet, this value is more than two times lower than the respective value in Table 1.3. In contrast, we can clearly see that *Common-friends*, *Friends-Measure*, and *Jaccard's Coefficient* have high InfoGain values for networks where all links are visible and the lowest InfoGain values for networks with only 5% of visible links.

1.6 Conclusion

Today's terror acts are conducted by groups of individuals who link with facilitators and other groups. Together, such a network has the resources, the means, and the insights to execute these attacks [25]. It is important to establish connections between the participating parties in order to understand their group dynamics and effectively mitigate their activity [23].

In this chapter, we presented the results of a study on link prediction in small scale, highly fractional networks, such as terrorists networks obtained from an analysis of publicly available media. We employed machine learning classifiers trained on a set of features extracted from the network structure. The results indicate that, in contrast to large networks where effective link prediction can be performed even when the vast majority of links are hidden, small scale networks require at least

Table 1.3 InfoGain of the extracted features for various social networks.

	Common-Friends(u, v)			Friends-measure(u, v)		Jaccard's-coefficient(u, v)	Preferential-attachment-score(u, v)		Total-Friends(u, v)				
	$d(u)$	$dout(u)$							$d(v)$	$din(u)$	$din(v)$	$dout(v)$	
Students	0.59	0.02		0.58	0.60	0.08	0.30	0.02					
F&F	0.36	0.14	0.17	0.67	0.36	0.30	0.13	0.12	0.13	0.13	0.09		
Researchers	0.48	0.20	0.28	0.59	0.46	0.25	0.09	0.09	0.08	0.14	0.05		
Co-Workers	0.49	0.13		0.37	0.46	0.25	0.12	0.19					
PIT	0.74	0.11		0.74	0.75	0.26	0.18	0.11					
AnyBeat	0.36	0.41	0.35	0.65	0.37	0.70	0.63	0.38	0.37	0.34	0.34		
Grand Total	0.50	0.17	0.27	0.60	0.50	0.31	0.24	0.15	0.19	0.20	0.16		

Table 1.4 InfoGain of the extracted features for fractional social networks with 5% of visible links.

	Common-Friends(u, v)			Friends-measure(u, v)		Jaccard's-coefficient(u, v)	Preferential-attachment-score(u, v)		Total-Friends(u, v)				
	$d(u)$	$dout(u)$							$d(v)$	$din(u)$	$din(v)$	$dout(v)$	
Students	0.01	0.00		0.01	0.01	0.00	0.00	0.01					
F&F	0.01	0.02	0.03	0.03	0.01	0.06	0.03	0.03	0.01	0.02	0.02		
Researchers	0.01	0.03	0.03	0.03	0.01	0.03	0.05	0.04	0.00	0.02	0.02		
Co-Workers	0.01	0.01		0.01	0.01	0.03	0.04	0.05					
PIT	0.01	0.02		0.01	0.01	0.01	0.01	0.02					
AnyBeat	0.02	0.30	0.24	0.07	0.02	0.26	0.55	0.29	0.24	0.23	0.24		
Grand Total	0.01	0.06	0.10	0.03	0.01	0.06	0.11	0.07	0.08	0.09	0.10		

30%–50% visible links to get AUC values above 0.8. We also notice that networks containing large cliques, or densely connected parts, are less vulnerable to random removal of links.

In conclusion, as expected, our results show that Information Gain values of the structural features drop when most of the links are removed from the network. However, it should be noted that the decrease in the Information Gain is not uniform. Features that were the most useful for predicting links in the original network become the least useful when only 5% of the links are visible. These features are the number of common friends, the Friends-Measure [11], and Jaccard's Coefficient [10]. All these measures refer to the connections and similarities between vertices in

both neighborhoods of the two ends of the link being tested. When there are enough connections, these measures make sense, however, when only few links are visible, the most useful information is the number of acquaintances each vertex has.

Availability

An anonymous version of the Co-Workers, Researchers, Students, and the AnyBeat social network topologies crawled as a part of this study are available for other researchers to use, by contacting the authors of this paper.

References

1. D. Liben-Nowell and J. Kleinberg. The link-prediction problem for social networks. *Journal of the American society for information science and technology*, 58(7):1019–1031, 2007.
2. M. Al. Hasan, V. Chaoji, S. Salem, and M. Zaki. Link prediction using supervised learning. *SDM Workshop of Link Analysis, Counterterrorism and Security*, 2006.
3. Janardhan Rao Doppa, Jun Yu, Prasad Tadepalli, and Lise Getoor. Chance-constrained programs for link prediction. In *Proceedings of Workshop on Analyzing Networks and Learning with Graphs at NIPS Conference*, 2009.
4. Hially Rodrigues Sa and Ricardo B. C. Prudencio. Supervised learning for link prediction in weighted networks. *III International Workshop on Web and Text Intelligence*, 2010.
5. Han Hee Song, Tae Won Cho, Vacha Dave, Yin Zhang, and Lili Qiu. Scalable proximity estimation and link prediction in online social networks. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, IMC '09, pages 322–335, New York, NY, USA, 2009. ACM.
6. William J. Cukierski, Benjamin Hamner, and Bo Yang. Graph-based features for supervised link prediction. *International Joint Conference on Neural Networks*, 2011.
7. E.M. Airolidi, D.M. Blei, S.E. Fienberg, E.P. Xing, and T. Jaakkola. Mixed membership stochastic block models for relational data with application to protein-protein interactions. In *Proceedings of the international biometrics society annual meeting*, 2006.
8. Z. Huang, X. Li, and H. Chen. Link prediction approach to collaborative filtering. *Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*, 2005.
9. M. Al. Hasan and M. J. Zaki. *Social Network Data Analytics*. Springer, 2011.
10. P. N. Tan, M. Steinbach, and V. Kumar. *Introduction to Data Mining*. Addison Wesley, 2005.
11. M. Fire, L. Tenenboim, O. Lesser, R. Puzis, L. Rokach, and Y. Elovici. Link prediction in social networks using computationally efficient topological features. In *Privacy, Security, Risk and Trust (PASSAT), 2011 IEEE Third International Conference on and 2011 IEEE Third International Conference on Social Computing (SocialCom)*, pages 73–80. IEEE, 2011.
12. Leo Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43, March 1953.
13. Michael Fire, Gilad Katz, Lior Rokach, and Yuval Elovici. Links reconstruction attack using link prediction algorithms to compromise social networks privacy. *Security and Privacy in Social Networks*, 2012.
14. Y. Altshuler, M. Fire, Y. Elovici, and A. Pentland. How many makes a crowd? on the evolution of learning as a factor of community coverage. 2012.
15. M. Fire, G. Katz, Elovici Y., Shapria B, and Rokach L. Predicting student exam's scores by analyzing social network data. 2012.

16. Prithviraj Sen, Galileo Mark Namata, Mustafa Bilgic, Lise Getoor, Brian Gallagher, and Tina Eliassi-Rad. Collective classification in network data. *AI Magazine*, 29(3):93–106, 2008.
17. C.C. Aggarwal. *Social network data analytics*. Springer-Verlag New York Inc, 2011.
18. J. Arquilla and D.F. Ronfeldt. *Networks and networks: The future of terror, crime, and militancy*. Number 1382. Rand Corp, 2001.
19. S. Ressler. Social network analysis as an approach to combat terrorism: Past, present, and future research. *Homeland Security Affairs*, 2(2):1–10, 2006.
20. V.E. Krebs. Mapping networks of terrorist cells. *Connections*, 24(3):43–52, 2001.
21. R. Rothenberg. From whole cloth: Making up the terrorist network. *Connections*, 24(3):36–42, 2001.
22. J.A. Rodriguez. The march 11th terrorist network: In its weakness lies its strength. 2005.
23. M. Sageman. *Understanding terror networks*. Univ of Pennsylvania Pr, 2004.
24. A. Basu. Social network analysis of terrorist organizations in india. In *North American Association for Computational Social and Organizational Science (NAACSOS) Conference*, pages 26–28, 2005.
25. U.K. Wiil, N. Memon, and P. Karampelas. Detecting new trends in terrorist networks. In *Advances in Social Networks Analysis and Mining (ASONAM), 2010 International Conference on*, pages 435–440. IEEE, 2010.
26. M. Dombroski, P. Fischbeck, and K.M. Carley. Estimating the shape of covert networks. In *Proceedings of the 8th International Command and Control Research and Technology Symposium*, 2003.
27. J. Leskovec, D. Huttenlocher, and J. Kleinberg. Predicting positive and negative links in online social networks. In *Proceedings of the 19th international conference on World wide web*, pages 641–650. ACM, 2010.
28. Dirk Nachbar. IJCNN Social Network Challenge. <http://www.kaggle.com/c/socialNetwork/Data>, November 8 2010.
29. A. Narayanan, E. Shi, and B.I.P. Rubinstein. Link prediction by de-anonymization: How we won the kaggle social network challenge. In *Neural Networks (IJCNN), The 2011 International Joint Conference on*, pages 1825–1834. IEEE, 2011.
30. Facebook-Newsroom. <http://www.facebook.com>.
31. B. Zhao, P. Sen, and L. Getoor. Event classification and relationship labeling in affiliation networks. In *Proceedings of the Workshop on Statistical Network Analysis (SNA) at the 23rd International Conference on Machine Learning (ICML)*, 2006.
32. M. Fire, L. Tenenboim, R. Puzis, O. Lesser, L. Rokach, and Y. Elovici. Computationally efficient link prediction in variety of social networks. *submitted for publication in ACM TIST*.
33. N. Aharony, W. Pan, C. Ip, I. Khayal, and A. Pentland. Social fmri: Investigating and shaping social mechanisms in the real world. *Pervasive and Mobile Computing*, 2011.
34. P. Shannon, A. Markiel, O. Ozier, N.S. Baliga, J.T. Wang, D. Ramage, N. Amin, B. Schwikowski, and T. Ideker. Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome research*, 13(11):2498–2504, 2003.
35. Academia.edu. <http://www.academia.edu>.
36. Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. The weka data mining software: an update. *SIGKDD Explor. Newsl.*, 11:10–18, November 2009.
37. A.-L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.