

**TOPIC-BASED INFLUENCE COMPUTATION IN
SOCIAL NETWORKS UNDER RESOURCE
CONSTRAINTS**

A THESIS SUBMITTED TO
THE GRADUATE SCHOOL OF ENGINEERING AND SCIENCE
OF BILKENT UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR
THE DEGREE OF
MASTER OF SCIENCE
IN
COMPUTER ENGINEERING

By
Kaan Bingöl
June, 2015

TOPIC-BASED INFLUENCE COMPUTATION IN SOCIAL NETWORKS UNDER RESOURCE CONSTRAINTS

By Kaan Bingöl

June, 2015

We certify that we have read this thesis and that in our opinion it is fully adequate, in scope and in quality, as a thesis for the degree of Master of Science.

Assoc. Prof. Dr. Hakan Ferhatosmanođlu(Advisor)

Assoc. Prof. Dr. Buđra Gedik

Prof. Dr. İsmail Hakkı Toroslu

Approved for the Graduate School of Engineering and Science:

Prof. Dr. Levent Onural
Director of the Graduate School

ABSTRACT

TOPIC-BASED INFLUENCE COMPUTATION IN SOCIAL NETWORKS UNDER RESOURCE CONSTRAINTS

Kaan Bingöl

M.S. in Computer Engineering

Advisor: Assoc. Prof. Dr. Hakan Ferhatosmanoğlu

June, 2015

As social networks are constantly changing and evolving, methods to analyze dynamic social networks are becoming more important in understanding social trends. However, due to the restrictions imposed by the social network service providers, the resources available to fetch the entire contents of a social network are typically very limited. As a result, analysis of dynamic social network data requires maintaining an approximate copy of the social network for each time period, locally. We study the problem of dynamic network and text fetching with limited probing capacities, for identifying and maintaining influential users as the social network evolves. We propose an algorithm to probe the relationships (required for global influence computation) as well as posts (required for topic-based influence computation) of a limited number of users during each probing period, based on the influence trends and activities of the users. We infer the current network based on the newly probed user data and the recent version of the network maintained locally. Additionally, we propose to use link prediction methods to further increase accuracy of our network inference. We employ PageRank as the metric for influence computation. We illustrate how the proposed solution maintains accurate PageRank scores for computing global influence, and topic-sensitive weighted PageRank scores for topic-based influence. The latter relies on a topic-based network constructed via weights determined by semantic analysis of posts and their sharing statistics. We evaluate the effectiveness of our algorithms by comparing them with the true influence scores of the full and up-to-date version of the network, using data from the micro-blogging service Twitter. Results show that our techniques significantly outperform baseline methods (80% higher accuracy for network fetching and 77% for text fetching) and are superior to state-of-the-art techniques from the literature (21% higher accuracy).

Keywords: Evolving Social Networks, Data Probing, Network Inference.

ÖZET

KAYNAK KISITLAMALARI ALTINDA SOSYAL AĞLAR ÜZERİNDE KONU TABANLI ETKİ HESAPLAMASI

Kaan Bingöl

Bilgisayar Mühendisliği, Yüksek Lisans

Tez Danışmanı: Doç. Dr. Hakan Ferhatosmanoğlu

Haziran, 2015

Sosyal ağlar sürekli değişirken ve gelişirken, dinamik yapıdaki bu ağların analizi için gerekli metotların önemi de sosyal eğilimleri anlamak açısından artmaktadır. Fakat sosyal ağ servis sağlayıcıları tarafından uygulanan kısıtlamalar nedeniyle, bir sosyal ağın topolojik durumu ve paylaşımlarıyla birlikte tüm içeriğini toplamak için mevcut kaynaklar yetersiz kalmaktadır. Sonuç olarak, değişken sosyal ağ verisinin analizi; verinin yaklaşık bir kopyasını yerel olarak muhafaza etmeyi ve zaman içinde yenilemeyi gerektirir. Biz, sosyal ağ geliştikçe, ağ üzerindeki etkili kişileri belirlemek ve zaman içerisinde takip etmek adına; hem ağ hem de metin verisinin sınırlı kaynaklar altında toplanması problemi üzerinde çalışıyoruz. Sınırlı sayıda kullanıcının her zaman aralığı için ilişkilerini (genel etki hesaplaması için gereklidir) ve metin paylaşımlarını (konu tabanlı etki hesaplaması için gereklidir) toplamak için kullanıcıların etki eğilimlerini ve eylemlerini göz önünde bulunduran bir algoritma öneriyoruz. Yeni toplanmış kullanıcı verisini ve lokal olarak sürdürülen ağın en son versiyonunu temel alarak; güncel ağ yapısını çıkarıyoruz. Buna ek olarak, ağ çıkarsama metodumuzun doğruluğunu daha da artırmak adına, bağlantı önerme algoritmaları kullanıyoruz. "PageRank puanı" nı genel etki hesaplaması için ölçü olarak belirledik. Önerdiğimiz çözümlerin, genel etki için "PageRank skorları" nı ve konu tabanlı etki için paylaşım metinlerinin anlamsal analizleri ile paylaşım istatistikleri harmanlanarak ve belirlenen ağırlıklar kullanılarak oluşturulan konu temelli ağlar üzerinden hesaplanan ağırlıklı "PageRank skorları" nı, nasıl yüksek bir doğrulukla yenileyerek sürdürdüğünü gösterdik. Algoritmalarımızın etkinliğini ölçmek adına sonuçlarımızı, bir mikroblog servisi olan "Twitter" üzerinden toplanan veriler ile oluşturan ağın, tam ve en güncel hali üzerinden hesaplanan gerçek etki skorları ile karşılaştırdık. Elde edilen sonuçlara göre; önerdiğimiz teknikler temel referans tekniklerini önemli bir ölçüde geri bırakırken (ağ çekimi için %80, metin çekimi için %77 daha doğrudur), literatürdeki en gelişmiş tekniklerden de daha iyi bir performans

göstermiştir (%21 daha doğrudur).

Acknowledgement

I would like to gratefully and sincerely thank Assoc. Prof. Dr. Hakan Ferhatosmanođlu for his guidance, patience, understanding and support during my studies. His mentorship was priceless for me.

I would also like to gratefully and sincerely thank Assoc. Prof. Dr. Buđra Gedik for his assistance and guidance with his excellent knowledge during my studies.

I am grateful to Prof. Dr. İsmail Hakkı Toroslu for kindly accepting to be in my committee and, reading and reviewing this thesis.

I would like to thank Bahaeddin Eravcı and Dr. Çađrı Özgenç Etemođlu for all of their ideas and collaboration to improve this thesis.

I am also grateful to my office friends for their knowledge, friendship and support.

Finally, and most importantly, I would like to thank my family and friends for their unending encouragement, tolerance and support.

Contents

- 1 Introduction** **1**
 - 1.1 Contributions 2
 - 1.2 Outline 4

- 2 Related Work** **5**
 - 2.1 Social Networks and Influence 5
 - 2.2 Topic-Based Influence 7
 - 2.3 Evolving Social Networks 8
 - 2.4 Network Inference 8

- 3 Problem Definition** **10**

- 4 Overall System Architecture** **14**
 - 4.1 Social Network Data Collection 15
 - 4.2 Score Analysis 16

- 5 Dynamic Data Fetching Methods 20**
 - 5.1 Analysis of PageRank Change 20
 - 5.2 Dynamic Network Fetching using Influence Past 22
 - 5.2.1 Dynamic Tweet Fetching using Topic-Based Influence Past 27
- 6 Experiments and Results 29**
 - 6.1 Data Sets 29
 - 6.1.1 Evaluation of Dynamic Network Fetching 31
 - 6.1.2 Evaluation of Dynamic Tweet Fetching 32
 - 6.1.3 Experimental Results and Discussion 33
- 7 Conclusion 43**

List of Figures

4.1	Overall system architecture.	14
4.2	Influence past of a user	16
5.1	A sample graph for analysis.	21
5.2	G-WG method, probe the global network and probe the tweet sets for each topic of interest.	28
5.3	WG-WG method, probe the networks and the tweet sets for each topic of interest.	28
6.1	In-edge distributions of the original network (on the left) and the pruned network (on the right).	31
6.2	Performance of Change Probing.	34
6.3	Performance of Change Probing as a function of time.	35
6.4	Performance of Round-Robin Change Probing.	36
6.5	Performance of Round-Robin Change Probing as a function of time.	37
6.6	Comparison of the probing strategies.	38

6.7	Comparison of the Probing strategies with respect to time.	39
6.8	Accuracy of the link prediction algorithm.	39
6.9	Performance of RR Change with inference.	40
6.10	Performance of Change Probing for dynamic tweet fetching.	42

List of Tables

6.1	Estimated influential accounts.	40
6.2	Top-10 topic relevance ratios for $G-WG$ and $WG-WG$ for dynamic tweet fetching.	42

Chapter 1

Introduction

Analysis of social networks have attracted significant research attention in recent years due to the popularity of online social networks among users and the vast amount of social network data publicly available for analysis. Applications of social network analyses are abound, such as influential user detection, community detection, information diffusion, network modeling, user recommendation, to name a few.

Influential user detection is a key social analysis used for opinion mining, targeted advertising, churn prediction, and word-of-mouth marketing. Social networks are dynamic and constantly evolving via user interactions. Accordingly, the influence of users within the network are also dynamic. Beyond the current influence of users, tracking the influence trends provides greater insights for deeper analysis. By combining the patterns of the past with the current information, comprehensive analysis on customers, marketing plans, and business models can be performed more accurately. For example, forecasting future user influences can be used to detect ‘rising stars’, who can be employed in upcoming on-line advertisement campaigns.

1.1 Contributions

We address the problem of identifying and tracking influential users in dynamic social networks under real-world data acquisition resource limits. The current approaches for influence analysis mostly assume that the graph structure is static, or even when it is dynamic, the data is completely known and reside in a local database. However, in many cases, analysts are third-party clients and do not own the data. They cannot keep the data completely fresh as changes happen, since it is typically gathered from a service provider with limitations on resources or even on the amount of data provided. Third-party data acquisition tools access the data via rate-limited APIs, which constraint the fetching capacity of clients. These externally enforced limits prevent the collection of entire up-to-date data within a predetermined period. To this end, we present an effective solution to rate-limited fetching of evolving network relations and user posts. Our system maintains a local, partially fresh copy of the data (e.g. relations, status, tweets) and calculates influence scores based on inferred network and text data. The proposed solution probes limited number of active users whose influence scores are changing significantly within the network. By combining previous and the newly probed network data, we are able to infer the current network accurately. The local network copy is maintained while consuming resources within allowed limits, and at the same time, influence values of the users are computed as accurately as possible.

While computing and maintaining influence scores, we consider both global and topic-based influence. Active and influential users mostly affect the general opinion with respect to their topics of authority. For instance, a company marketing sports goods will be interested in locating users who have high influence in sports, rather than the global community. While this leads us to include topic-based analyses for our problem setting, general influence scores of users are still of interest as well. For instance, a politician would prefer a broader audience and identify a list of globally influential users to promote her cause. In our system, we utilize both global and topic-centric networks and compute global as well as topic-based influences.

To demonstrate the effectiveness of our solutions, we use Twitter [1]. Twitter is a good fit for research on dynamic user influence detection due to its large user base

and highly dynamic user activity. One can collect two-way friendship relations as well as one-way follow, re-tweet and favorite relations via publicly available Twitter API. The APIs have well-defined resource limits [2], which motivates the need for our probing algorithms. We calculate PageRank [3] on the Twitter network as the influence score for the users. To generate topic-based influence scores, we adapt the weighted PageRank [4], and adjust the initial scores and transition probabilities based on topic relevance scores of the users. The topic relevance scores are computed based on user posts, using text mining techniques, as well as their re-tweet and favorite counts.

To further improve the accuracy of our network inference, we perform link prediction using trends on user relationships. The proposed solution shows increased accuracy on Twitter data when compared with other methods from the literature. Estimated network structure is shown to be very close to the actual up-to-date network, with respect to influential users. The proposed solutions address not only the limitations of data fetching via public APIs, but also local processing when the resources are limited to fetch the entire data. We summarize our major contributions as follows:

- We estimate global and topic-based influence of users within a dynamic social network. For topic-based influence estimation, we construct topic-based networks via semantic analyses of tweets and the use of re-tweet and favorite statistics for the topic of interest.
- We propose efficient algorithms for collecting dynamic network and text data, under limited resource availability. We leverage both latest known user influence values, as well as the past user influence trends in our probing strategy. We further improve our probing techniques by applying link predictions methods.
- We evaluate our proposed algorithms and compare results to several alternatives from the literature. The experimental results for relationship fetching show that the proposed algorithms perform 80% better than the baseline methods, and 21% better than the state-of-the-art method from the literature in terms of mean squared error. For tweet fetching methods used for topic-based influence detection, our algorithms perform 77% better than the alternative baselines in terms of the Jaccard similarity measure.

1.2 Outline

The rest of the chapters is organized as follows. Chapter 2 gives background material on related work. Chapter 3 describes the resources constraint problem for data collection. Chapter 4 describes the overall system architecture and presents influence estimation techniques. Chapter 5 explains algorithms and strategies proposed for the network and text fetching problems. Chapter 6 discusses evaluation results obtained from experiments run on real data. Chapter 7 concludes the thesis.

Chapter 2

Related Work

2.1 Social Networks and Influence

Increases in the popularity of social networks and the availability of public data acquisition tools for them have put social networks on the spotlight of both academic and industrial research. Influential user estimation problem is studied by many researchers following a wide variety of different methodologies. Within this context, some studies introduce centrality measures in order to reflect influence of users. [5] introduces several definitions:

- *Degree centrality* picks users who are located at the center of a network, in the sense that they are connected with many other users.
- *Betweenness centrality* picks users who are located on the path between many nonadjacent users. Since such users connect many users, they should have a greater control within the network.
- *Closeness or distance centrality* picks users who are close to all users in network. The concept of closeness is defined by short average distance. The idea behind this measure is if a user interact with many others quickly, he should be influential.

For viral marketing applications, [6] develops methods for computing network influence from collaborative filtering databases by using heuristics in a general descriptive probabilistic model of influence propagation. [7] addresses a similar problem by studying the linear threshold and independent cascade models, and [8] presents a simple greedy algorithm for maximizing the social influence in a general model, termed the decreasing cascade model:

- In *the linear threshold model*, they assign edge weights as $\frac{c_{v,u}}{d_v}$, where $c_{v,u}$ is the number of edges that exist between node v and u , and d_v is the degree of user v .
- In *the independent cascade model*, a uniform probability $p_v(u)$ is assigned to edges between users so that users v has chance of $p_v(u)$ to affect user u .
- In *the decreasing cascade model*, they assign probability $p_v(u, S)$ where S denotes the set of user v 's neighbors that already tried to affect v and failed. It is the success probability of node u given that u affects v after v 's affected neighbors failed to affect. This model is a generalization of *the independent cascade model*.

In more recently published work, [9] presents a novel methodology for selecting users to maximize the influence spread. [10] uses *maximum influence in-arborescence* (MIIA) based greedy algorithms, which significantly improve scalability. [11] compares different types of influence measures and discusses the findings. [12] applies statistical tests in order to distinguish user influence from correlation, and [13] investigates conformity influence on social networks. [14] uses a greedy approach for the influence maximization problem and proposes efficient degree discount heuristics. [15] studies the determination of influence probabilities for edges by examining the past behavior of users. [16, 17] study the problem of finding rising stars in co-author networks based on mutual influence and other features.

2.2 Topic-Based Influence

Recently, researchers have studied extracting textual information associated with social networks. [18] studies topic modeling in social networks and proposes a solution for text mining on the network structure. [19] introduces the topic-based social influence problem. Their proposed model takes the result of any predefined topic modeling of a social network and constructs a network representing topic-based influence propagation. Distributed learning algorithms are used for this purpose, which leverage the Map-Reduce concept, thus, their methodology scales well for networks with millions of edges. [20] combines heterogeneous links and textual content for each user in order to mine topic-based influence.

Another recent study [21] uses a PageRank-like measure to find influential accounts on Twitter. They extend PageRank by using topic-specific probabilities in their random surfer model. Although their method is similar to ours, their influence measure utilizes the number of posts made on a specific topic. However, this is an indirect measure that cannot reliably capture influence. Therefore, we use topic distributions of user posts along with their sharing statistics (retweets and mentions in Twitter), which provides robust results, as it takes into account the real impact of posts. [22] conducts empirical study of different topic modeling strategies based on standard *Latent Dirichlet Allocation* (LDA) [23] the *Author-Topic Model* (AT model) [24]. [25] proposes joint probabilistic models of influence and topics. Their methodology performs a topic sampling over textual contents and tracks the topic snapshots over time. [26] uses re-tweets in measuring popularity and proposes machine learning techniques to predict popularity of the Twitter posts. [27, 28, 29] propose solutions for predicting popularity of online content. [30] studies the topic-aware influence maximization problem. Within this context, in this work we introduce a new method that combines topic-based analyses of posts with their sharing popularity for the purpose of topic-based influential user estimation.

2.3 Evolving Social Networks

Dynamic graph analysis has also attracted a lot of attention recently. In order to maintain dynamic networks, [31, 32, 33, 34, 35] propose algorithms for determining web crawling schedules. [36] studies the microscopic evolution of social networks. [37] studies incremental PageRank on evolving graphs. Researches also have investigated probing strategies for analyzing evolving social networks. [38] proposes influence proportional probing strategies for the computation of PageRank on evolving networks and [39] uses a probing strategy to capture observed image of the network by maximizing a performance gap function. [40, 41] study sampling over social networks. However, these studies only focus on current image of a network in their probing strategies. In contrast, we propose a method which also considers evolution of the probing metrics, so that the network could be probed more effectively.

2.4 Network Inference

In the context of network inference, [42] proposes representations for structural uncertainty and use directed graphical models and probabilistic relational models for link structure learning. However, their methodologies are not scalable. [43, 44, 45] use time evolving graph models for social network estimation. They apply time-varying dynamic Bayesian networks for modeling evolving network structures. [46] shows that third-parties can reach a user's information by searching a few friends. [47] develops a scalable algorithm to infer influence and diffusion network, assuming all users influence their neighbors with equal probabilities in the network. [48] removes this assumption and addresses the more general problem by formulating a maximum likelihood problem and guarantee the optimality of the solution. [49] proposes a linear model for the evolution of diffusion over time and [50] proposes the idea of diffusion centrality. [51, 52] studies a different problem related to network inference. Different from these works, we use friendship weighting method in order to infer link structures, similar to [53, 54, 55]. However, we use friendship weights only to infer edges between

users. Moreover, one can also use more informative features such as content-based influential effects. [56] studies diffusion of tweets throughout the Twitter network. This kind of technique could also be used in order to estimate impact of posts.

Chapter 3

Problem Definition

Our goal is to determine and maintain top-m influential users in the network, under a constrained probing setting. Among various methods to calculate a user's influence in the network, we have chosen PageRank based methods, since PageRank is well understood and used widely in the literature for various network structures [21, 57]. While computing influence, PageRank naturally considers the number of followers a user has, but more importantly it considered the topological place of the user within the network. Therefore, we assume that a user's influence in the network corresponds to its PageRank score. As a result, the top-m influential user determination problem turns into identifying the top-m users with the highest PageRank scores.

PageRank score calculation requires having access to all the relationships present between the users of the network. This means that we need to have the complete network data to compute exact PageRank scores. Moreover, if the network is dynamic, the calculation needs up-to-date network data for each time step in order to perform accurate influence analysis.

Our system continuously collects social network data (relations, tweets, re-tweets, etc.) via the publicly available Twitter API. Twitter enforces certain limitations on data acquisition using the Twitter APIs. There are different limitations for different types of data acquisition requests:

- *Relations*: 15 calls per 15 minutes, where each call is for retrieving a user’s relations. Moreover, if the user has more than $5K$ followers, we need an extra call for each additional $5K$ followers. This means that we can update relations with a maximum rate of 1 user per minute ($R_{rel} = 1$ user/min).
- *Tweets*: 180 calls per 15 minutes, where each call is for retrieving a user’s tweets. Moreover, if the user has more than $3.2K$ tweets, we need an extra call for each additional $3.2K$ tweets. This means that we can update tweets with a maximum rate of 12 users per minute ($R_{tw} = 12$ user/min).

Assuming that we update the network with a period of P days, we need the following condition to hold, in order to be able to capture the entire network of relations:

$$\text{Number of Users} \leq R_{rel} \cdot P \cdot 1440 \tag{3.1}$$

For getting the recent tweets of the users in the network, we need:

$$\text{Number of Users} \leq R_{tw} \cdot P \cdot 1440 \tag{3.2}$$

One can easily calculate that for a network as small as $250K$ users, we need 174 days to update the complete network in the best case¹. This analysis shows that the rate limits hinder the timeliness of the data collection process, which in turn affects the timeliness of the calculation process to find and track influential users in the network. Furthermore, Twitter is a highly dynamic network that evolves at a fast rate, which means that not refreshing the network frequently will result in significant degradation in the accuracy of the influence scores. Current resource limits prohibit the system to collect the network data in a reasonable period of time. Therefore, the evolving network’s relationships and the tweet sets are not fully observable at every analysis time step.

To overcome this limitation, we propose to determine a small subset of users during each data collection period, whose information is to be updated. This data collection process, which does not violate the rate limits of the API, is sufficient to maintain an

¹Assuming all users have $\leq 5K$ followers, so that they could be fetched with one call.

approximate network with a reasonable data collection period, while at the same time providing good accuracy for the influence scores.

We apply the concept of *probing* for efficient fetching of the dynamic network and the user tweet sets. We denote a network at time t as $G_t = \{V_t, E_t\}$, where V_t is the set of users and $E_t \subset V_t \times V_t$ is the set of edges representing the follower relationship within the network. In other words, $(u, v) \in E_t$ means that the user $u \in V_t$ is following the user $v \in V_t$. Our model uses an evolving set of networks in time, represented as $\{G_t \mid 0 \leq t \leq T\}$. However, we assume that we have fully² observed the network only at time $t = 0$. G_t where $t > 0$, can only be observed partially by probing. At each time period, we use an algorithm to determine a subset of k users and probe them via API calls. We then update the existing local network with the new information obtained from the probed users. In other words, we maintain a partially observed network G'_t , which is potentially different than the actual network G_t . Larger k values ($0 \leq k \leq |V_t|$) bring the partial network G'_t closer to the actual network G_t . However, using large k values is not feasible due to rate limits outlined earlier. Our probing strategy should select a relatively small number of users to probe, so that the data collection process can be completed within the period P (as determined by Eq. 3.1). Furthermore, these probed users should bring the most value in terms of performing accurate influence detection.

Dynamic Network Fetching Problem Definition: We assume that complete network information is available only at time 0, i.e., G_0 is known. The problem is defined as determining a subset of users of size k at time t , denoted by $U_t \subset V_t$ s.t. $|U_t| = k$, by analyzing the local graph G'_{t-1} . The system will update the relationships of the users included in this subset to construct the local network at time t , that is G'_t . Specifically, this new network G'_t is constructed by replacing the relationships of the users in G'_{t-1} with the newly fetched relationships from the probing of the users in U_t . We aim to choose U_t such that the influence scores of the estimated network G'_t will be as close as possible to the true scores of the real network G_t . The final objective is to estimate the PageRank scores $PR'_v(t), \forall v \in G_t$ as accurately as possible, using partial knowledge

²The initial probing of the network can be accelerated via the use of multiple cooperating fetchers. However, this is clearly not a sustainable and feasible approach for continued probing of the network, as it requires large number of accounts, which are subject to bot detection and suspension.

about G_{t-1} , that is G'_{t-1} .

In order to track topic-specific influence scores of the users, we analyze their latest tweets. One needs to collect predetermined amount of tweets for all of the users to be able to compute exact influence scores. However, due to the rate limitations (as determined in Eq. 3.2), we cannot fetch all the tweets within the desired period. Instead of retrieving tweets of every user, we determine a subset of users so that by collecting tweets of this subset, the topic scores of the users will be as close to the true scores as possible. We denote the tweet set at time t as T_t . We again assume that we have observed this set fully only at time 0, that is T_0 is known. The other snapshots can only be observed partially by probing. I.e., we locally maintain partial tweet sets T'_t , where $t > 0$.

Dynamic Tweet Fetching Problem Definition: Given the tweets T_0 of all users in the network at time 0, the problem is defined as determining a subset of users of size k at time t , denoted by $U_t \subset V_t$ s.t. $|U_t| = k$, by analyzing the tweet set T'_{t-1} and local graph G'_{t-1} . By collecting tweets of the users included in U_t , we construct an approximate tweet set T'_t and update the topic-based network accordingly. The final objective is to estimate the topic-based influence scores of the users in the network as accurate as possible. Thus, the goal is to pick the subset U_t , so as to maximize the accuracy of the influence scores computed on the estimated topic-based network.

Chapter 4

Overall System Architecture

In this chapter, we briefly describe our system architecture, depicted in Figure 4.1, and the basic workflow of the system.

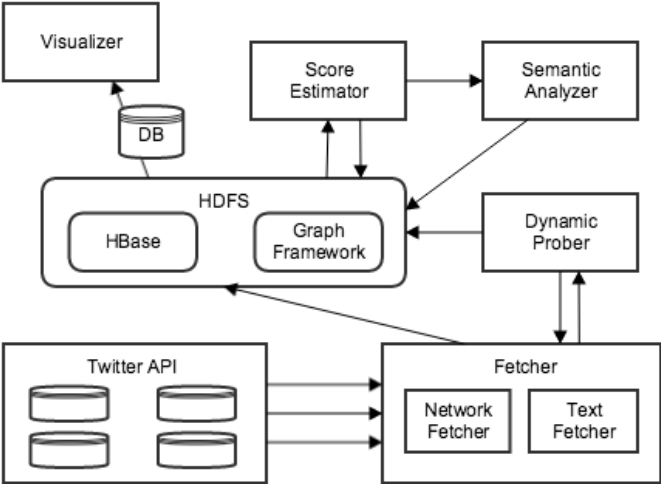


Figure 4.1: Overall system architecture.

4.1 Social Network Data Collection

We use the Twitter network and tweets to analyze user influence. A Twitter network is a directed, unweighted graph where the nodes represent users and the edges denote follower relationships in Twitter. When a user u follows a user v , there is mutual influence between them, which has an effect on both users' influence scores. In order to construct our network, we first determine a small set of users called the *core seeds*. For illustration, we started with some popular Turkish Twitter accounts including newspapers, TV channels, politicians, sport teams, and celebrities. Second, we collect one-hop relations of the core seeds and add the unique users to a set called the *main seeds*. We iterate once more to collect one-hop relations of the main seeds with a filter to avoid unrelated and inactive users. This filter has three conditions: *a)* a user must have at least five followers, *b)* a user must have at least one tweet within the last three months, and *c)* the tweet language of a user must be Turkish. As a result of this process, we have determined our *seed users* set, which includes approximately 2.8 million unique users. In the final step of the data collection phase, we acquire the relations of the seed users to determine G_0 , that is the social network graph at time 0. Furthermore, we collect tweets of the seed users in order to construct T_0 , that the tweet set at time 0.

We implemented the proposed methods using a distributed system with HBase and HDFS serving as the database and file system backends. The system consists of six main parts: *a)* local copy of the social network data on HDFS, *b)* data fetcher, *c)* dynamic prober, *d)* score estimator, *e)* semantic analyzer, and *f)* visualizer. Data fetcher component, as the name implies, fetches the data (network relations and tweets) via rate-limited Twitter APIs, periodically. Dynamic prober makes a dynamic probing analysis, decides which users are going to be fetched and notifies data fetcher to bring the information, accordingly. Score estimator calculates users' influence and the related parameters of the proposed algorithms, which are essential parts of the probing method. Semantic analyzer performs keyword extraction and calculates the related parameters for constructing topic-based networks. Finally, visualizer provides a graphical user interface for result analysis.




Figure 4.2: Influence past of a user

4.2 Score Analysis

We calculate influence scores of users based on their relationships and the overall impact of their tweets in the network. We analyze topic activities of the users from their tweets and determine topic-sensitive user influence scores. Overall, we are using two types of scores, namely *global influence* and *topic-based influence*, which can be interpreted together for a more detailed analyses.

Global Influence Score. This score is a measure of the user’s overall influence within the network. For this purpose we use the personalized PageRank algorithm. PageRank value $PR_v(t)$ at time t for a user $v \in G_t$ directly corresponds to the global influence score of it and will be used interchangeably throughout the thesis.

Figure 4.2 illustrates the evolving nature of the influence score by showing the global and topic-based influence scores history of a user, which is selected by our algorithm as one of the most important users that should be probed during the first collection period. This is the official account of the president of the Republic of Turkey. Besides the account’s high impact, we observe that its influence also varies significantly over time, which further justifies the need to probe this account frequently. A reason of the variation in influence score is that the time period shown in the figure matches with the elections for the Presidency (10 August 2014). After becoming the

new president, the account's influence has further increased. During this period, it is always selected as a top user to be probed by our proposed approach. This is intuitive, as it is a popular account with a high change in the influence scores over time.

Topic-Based Influence Score. The system calculates topic-based influence scores representing user activity and impact on a specific topic. We perform semantic analysis on user tweets by taking re-tweets and favorite numbers into consideration as well. A re-tweet (RT) is a re-posting of someone else's tweet, which helps users quickly share a tweet that they are influenced by or like. A favorite (FAV) is another feature that represents influence relation between users, wherein one user can mark a tweet by another user as a favorite. These two features are helpful to estimate influence of an individual tweet. Since Twitter is a micro-blogging platform, users are generally tweeting on specific topics. While many tweets are mostly conversational and reflect self-information [58, 59], some are being used for information sharing, which is important in harvesting knowledge. RTs and FAVs are effective in separating relevant and irrelevant tweets. Therefore, we use them in our topic weight analysis to estimate influence value of a tweet on a specific topic.

Topic-based network construction process consists of three main phases: *a*) keyword extraction on tweets, *b*) correlation of keywords with topic dictionaries, and *c*) weight calculation.

In the first phase, keywords are extracted from the tweets by using information retrieval techniques, including word stemming and stop word elimination. The output from this phase is a keyword analyzed tweet corpus for each individual user and the related histogram which captures the frequencies of the related keywords (K). These corpora are further analyzed in the second phase.

We have created a keyword dictionary (D_j) for each topic (C_j), in order to score tweets against topics. As part of each dictionary, we have assigned normalized weights to words, representing their topic relevance. In the second phase, using the weights from the dictionaries and the users' keyword histograms, we obtain the normalized raw topic scores of users for each one of the topics.

In the third phase, we calculate a value called the RT-FAV total for each user, which

is the summation of the number of re-tweets and favorites received by a user's tweets. We then scale the normalized raw topic score with the RT-FAV total for each user per topic of interest. The final results are used as the in-edge weights of the users on each topic, when forming the topic-based network.

Once the topic-based network construction is complete, we execute the weighted PageRank [4] algorithm and the resulting PageRank values of users, denoted by $WPR_v(t)$ at time t for $v \in G_t$, is assigned as their topic-based influence scores.

Due to the nature of the PageRank algorithm, some of the globally influential users also turn out to be highly influential for most or all of the topics. These users have a lot of followers and they are also followed by some of the influential accounts of the specific topics, which cause them to score high for topic-based analysis as well. Therefore, they can get high topic-based influence scores even if they do not actively tweet about the topic itself. To eliminate this effect, we apply one more level of filtering to remove these globally effective accounts from the topic-sensitive influence lists. In particular, if the number of tweets a user posted that are related with the topic at hand is less than a predefined percentage, e.g., %40¹, of the total number of tweets posted by the user, then the user is discarded for that topic. This filtering process significantly reduces the noise level in the analysis.

As a result, for each topic, we construct a weighted network in which an edge $((u, v))$ represents the amount of topic-specific influence a user (v) has on a follower user (u). Thus, the results of weighted PageRank algorithm gives us the overall topic-influence scores on the network.

Figure 4.2 also shows the topic-based score history of the official account of the president of the Republic of Turkey. According to our analysis, %80 of the account's topic activity is related to politics. Since it could not pass our applied activity filter on other topic categories, the system only calculates its topic influence scores for politics. We can see from the figure that the change on the topic-based scores are more dramatic compared to the global scores. This is intuitive, as they are depending on users' tweets and sharing statistics. A user might be very active on some weeks about a specific

¹Note that a tweet can be related to zero or more topics.

topic so that his influence on the topic might increase dramatically. Likewise, when he posts something important, he might get high sharing rates. On the other hand, when he just posts regular things which are not shared via others, his influence on the topic might decrease quickly.

Chapter 5

Dynamic Data Fetching Methods

In this chapter, we introduce our algorithms for probing in dynamic social networks. In order to efficiently determine a subset of vertices to probe, we develop heuristics for both dynamic network fetching and dynamic tweet fetching problems given in Chapter 3.

5.1 Analysis of PageRank Change

In this section, we give a theoretical analysis of how the changes in the network affect the PageRank values of the vertices. PageRank value of a specific vertex v is given as follows:

$$PR(v) = \alpha \sum_{\forall (u,v) \in E_{in}(v)} \frac{PR(u)}{|E_{out}(u)|} + \frac{1 - \alpha}{n}, \quad (5.1)$$

where $PR(v)$ denotes the PageRank value, $E_{in}(v)$ denotes the in-edge set, and $E_{out}(v)$ denotes the out-edge set for v .

Figure 5.1 shows an example network, which will be used to demonstrate the effects of network changes on PageRank values.




Figure 5.1: A sample graph for analysis.

Assume that an edge (u, v) is added due to the evolving nature of the network. Here, we analyze the effect of this addition on the PageRank values of the out neighbors of u . We see that the PageRank value of v is as follows per Equation 5.1:

$$\begin{aligned} PR^{new}(v) &= \alpha \left(\sum_{\forall(i,v) \in E_{in}(v)} \frac{PR(i)}{|E_{out}(i)|} + \frac{PR(u)}{|E_{out}(u)| + 1} \right) + \frac{1 - \alpha}{n} \\ &= PR(v) + \alpha \frac{PR(u)}{|E_{out}(u)| + 1} \end{aligned}$$

PageRank values of out neighbors of u other than v , such as w , are impacted as follows:

$$\begin{aligned} PR(w) &= \alpha \left(\sum_{\forall(i,w) \in E_{in}(w) \setminus (u,w)} \frac{PR(i)}{|E_{out}(i)|} + \frac{PR(u)}{|E_{out}(u)|} \right) + \frac{1 - \alpha}{n} \\ PR^{new}(w) &= \alpha \left(\sum_{\forall(i,w) \in E_{in}(w) \setminus (u,w)} \frac{PR(i)}{|E_{out}(i)|} + \frac{PR(u)}{|E_{out}(u)| + 1} \right) + \frac{1 - \alpha}{n} \\ PR^{new}(w) &= PR(w) - \alpha \frac{PR(u)}{|E_{out}(u)| \cdot (|E_{out}(u)| + 1)} \end{aligned}$$

These effects are the immediate responses on the vertices that are considered. These residual PageRanks will ripple out to all the vertices in all the paths from v and w in each iteration of the PageRank algorithm. But the effect will decrease as the residuals will be divided by the number of outgoing edges for each vertex visited. We will analyze the effects of the first iteration of the algorithm to simplify the problem and to get a general feel of the change in PageRank values. Considering expected value of

$\overline{E_{out}} = E[|E_{out}(u)|]$ as the average out-degree for vertices, the differential PageRanks are given as follows:

$$\nabla PR(v) = \alpha \frac{PR(u)}{\overline{E_{out}}} \quad (5.2)$$

$$\nabla PR(w) = -\alpha \frac{PR(u)}{\overline{E_{out}}^2} \quad (5.3)$$

We can see from Equations 5.2 and 5.3 that we should select the vertices, say u , with the following properties for accurate G'_t and $PR'_u(t)$ estimations:

- vertices with high PageRank values ($PR(u)$);
- vertices whose PageRank values change over time;
- vertices with high out-degrees ($E_{out}(u)$);
- vertices whose out-degrees change over time.

PageRank, when computed until the values converge in steady state, considers both incoming and outgoing edges. The parameters related to out-degree values are intrinsically taken into account when PageRank is computed. Hence, in our dynamic fetching approach, we focus only on PageRank values and their changes to cover all the cases listed above.

5.2 Dynamic Network Fetching using Influence Past

We aim to probe a subset, U_t , update the edges incident on vertices in U_t to form G'_t , and calculate PageRank values $PR'_v(t)$, $\forall v \in G_t$. In order to determine this subset, we use a time series of past PageRank values for a vertex v , named the *influence past* of v . Formally, we have $IP_v = [\dots, PR'_v(t-2), PR'_v(t-1)]$.

In our strategy for determining U_t , we consider the vertices whose PageRank values change considerably over time. In order to quantify this *change* for a vertex v , we are calculating the standard deviation of the time series IP_v , that is:

$$Change_v = \sigma_{IP_v} = \sqrt{Var(PR'_v)} \quad (5.4)$$

Choosing the best vertices to probe can be performed by calculating a score that is a linear combination of the PageRank value and the change in PageRank values, as given in Equation 5.5. Here, α parameter balances the importance of the two aspects. We assume that influence past that contains at least two data points is available for every user, in order to calculate the score changes.

$$Score(v) = (1 - \alpha)PR'_v(t - 1) + \alpha Change_v \quad (5.5)$$

After the selection of the users with respect to the ranking of $Score(v)$, we probe their current relations and form G'_t .

Round-Robin & Change Probing. Change Probing could cause the system to focus on a particular portion of the network and may discard the changes developing in other parts. This is because the probing scores of some vertices will be stale and as a result these vertices may consistently rank below the top-m, despite changes in their real scores. This bias could end up accumulating errors in the influence scores of these vertices and start to have an impact on the entire network. Therefore, we propose to use Change Probing together with Round-Robin Probing, in which users are probed in a random order with equal frequency. In this way, we aim to probe every vertex at least once within a specific period P . Round-Robin Change algorithm probes some portion of the network randomly and marks all probed users. Thus, any probed users are not probed randomly again, until all users are probed at least once within P . In this method, we control the balance between change vs. random selection by using a parameter $\beta \in [0, 1]$. In particular, we choose $\beta * k$ users to probe with Change Probing and $(1 - \beta) * k$ users with Round-Robin Probing.

Network Inference. Since we are able to fetch data only for a limited number of users, there is a high probability that other users in the network have changed their connections as well. To take these possible changes into account, we have also incorporated

link prediction into our solution, based on neighbor properties. Link prediction algorithms assign a score to an edge (u, v) based on their neighbors, denoted as Γ_u and Γ_v . The basic idea behind these scores is that the two vertices u and v are more likely to connect via an edge if Γ_u and Γ_v are similar, which is intuitive. Considering social networks, two people are likely to be friends if they have a lot of common friends. There are different scores used in the literature, including the common neighbors, Jaccard's coefficient, Adamic/Adar, and Resource Allocation Index (RA). We have adapted RA as part of our approach since it is found more successful on a variety of experimental studies on real life networks [60]. RA is founded on the resource allocation dynamics of complex networks and gives more weight to common neighbors that have low degree. For an edge (u, v) between any two vertices u and v , RA is defined as follows:

$$RA_{u,v} = \sum_{w \in \Gamma_u \cap \Gamma_v} \frac{1}{degree(w)}, \quad (5.6)$$

where Γ_v is the neighbors of v

The *RA* score, $RA_{u,v}$ for the edge (u, v) , is proportional to the probability of an edge being formed between the vertices u and v in the future. Based on this, we rank all the calculated *RA* scores. Since the edges in our network are not defined probabilistically and are defined deterministically as existent or non-existent, we need to determine how many of these scored edges should be selected. Therefore, we define a growth rate, E_g , which is the average change in the number of edges ($|E|$) between snapshots of the network after excluding the changes due to U_t . After calculating *RA* scores for all possible new edges, we choose E_g edges with the highest scores. Using this method, we add new connections to the current graph, to finally have the estimated graph G'_t . The pseudo code of the network inference based probing algorithm we use to select k vertices to probe is given in Algorithm 1.

ALGORITHM 1: Algorithm for Dynamic Network Fetching

Input: $G'_{t-1}, IP, PR'(t-1), \alpha, \beta \in [0, 1], k$ **Output:** G'_t *// Fetch network***for all** $v \in V_t$ **do**

$$\sigma_{IP_v} = \sqrt{\text{Var}(PR'_v)}$$

$$\text{Score}(v) = (1 - \alpha)PR'_v(t-1) + \alpha \cdot \sigma_{IP_v}$$

end for $U_t \leftarrow \emptyset$ **while** $|U_t| \leq k \cdot \beta$ **do**

$$v \leftarrow \text{argmax}_{v \in V_{t-1}} \text{Score}(v)$$

$$U_t \leftarrow U_t \cup \{v\}, V_{t-1} \leftarrow V_{t-1} \setminus \{v\}$$

end while**while** $|U_t| \leq k$ **do**

$$v \leftarrow \text{randomly choose from } V_{t-1}$$

$$U_t \leftarrow U_t \cup \{v\}, V_{t-1} \leftarrow V_{t-1} \setminus \{v\}$$

end whileProbe U_t for relationships, Form G'_t *// Infer network*Calculate $RA_{u,v}, \forall (u,v) \in \tilde{E} = V_t \times V_t$ **for** E_g **times do**

$$(u,v) \leftarrow \text{argmax}_{(u,v) \in E_t} RA_{u,v}$$

$$E_t \leftarrow E_t \cup \{(u,v)\}$$

end forOutput G'_t

ALGORITHM 2: Dynamic tweet fetching
via G - WG

Input: $T_{t-1}^{j'}$, TIP^j ,
 $WPR^{j'}(t-1)$, $\alpha, \beta \in [0, 1]$, k
Output: $T_t^{j'}$
for all C_j **do**
 for all $v \in V_{t-1}^j$ **do**
 $\sigma_{TIP_v} = \sqrt{Var(TPR_v')}$
 $Score^j(v) = (1 -$
 $\alpha)WPR_v^{j'}(t-1) + \alpha \cdot \sigma_{TIP_v^j}$
 end for
 $U_t^j \leftarrow \emptyset$
 while $|U_t^j| \leq k \cdot \beta$ **do**
 $v \leftarrow \operatorname{argmax}_{v \in V_{t-1}^j} Score^j(v)$

 $U_t^j \leftarrow U_t^j \cup \{v\}$,
 $V_{t-1}^j \leftarrow V_{t-1}^j \setminus \{v\}$
 end while
 while $|U_t^j| \leq k$ **do**
 $v \leftarrow$ randomly choose from
 V_{t-1}^j
 $U_t^j \leftarrow U_t^j \cup \{v\}$,
 $V_{t-1}^j \leftarrow V_{t-1}^j \setminus \{v\}$
 end while
 Probe U_t^j for tweets, Form $T_t^{j'}$
 Output $T_t^{j'}$
end for

ALGORITHM 3: Dynamic network and tweet
fetching via WG - WG

Input: $WG_{t-1}^{j'}$, $T_{t-1}^{j'}$, TIP^j ,
 $WPR^{j'}(t-1)$, $\alpha, \beta \in [0, 1]$, k
Output: $T_t^{j'}$, $WG_t^{j'}$
for all C_j **do**
 for all $v \in V_{t-1}^j$ **do**
 $\sigma_{TIP_v} = \sqrt{Var(TPR_v')}$
 $Score^j(v) = (1 -$
 $\alpha)WPR_v^{j'}(t-1) + \alpha \cdot \sigma_{TIP_v^j}$
 end for
 $U_t^j \leftarrow \emptyset$
 while $|U_t^j| \leq k \cdot \beta$ **do**
 $v \leftarrow \operatorname{argmax}_{v \in V_{t-1}^j} Score^j(v)$

 $U_t^j \leftarrow U_t^j \cup \{v\}$,
 $V_{t-1}^j \leftarrow V_{t-1}^j \setminus \{v\}$
 end while
 while $|U_t^j| \leq k$ **do**
 $v \leftarrow$ randomly choose from
 V_{t-1}^j
 $U_t^j \leftarrow U_t^j \cup \{v\}$,
 $V_{t-1}^j \leftarrow V_{t-1}^j \setminus \{v\}$
 end while
 Probe U_t^j for relationships, Form
 $WG_t^{j'}$
 Probe U_t^j for tweets, Form $T_t^{j'}$
 Output $WG_t^{j'}$, $T_t^{j'}$
end for

5.2.1 Dynamic Tweet Fetching using Topic-Based Influence Past

Our dynamic tweet fetching solution makes use of the weighted PageRank values and comprises of two two steps. First, we infer the evolving relationships of the network using the methods explained earlier in the previous section. This way we can track and estimate the changing relationships. Second, we select a subset of users to fetch their tweet data. Specifically, we aim to probe a subset, U_t , collect their tweets, and update the edge weights for the users in U_t ; all in order to form $WG_t^{j'}$ for a given topic C_j . We then compute weighted PageRank values to find $WPR_v^{j'}(t), \forall v \in WG_t^{j'}$ for a given topic C_j . To select the subset of users in U_t , we use a time series of the past weighted PageRank values, named the *topic-based influence past* of v . Formally, we have $TIP_v = [\dots, WPR_v^{j'}(t-2), WPR_v^{j'}(t-1)]$. This is performed independently for all topics of interest, $\{C_j\}$.

In this process, there two different evolving components: *a*) relationships among users (network) and *b*) topic-weights (tweets). Depending on an use-case, those two components could be maintained together or independently from each other. Therefore, we employ two different approaches in order to track the topic-based influence scores of the network:

- Use the global network parameters for network fetching and the topic-sensitive network parameters for tweet fetching. This is named as the *G-WG* method (Figure 5.2), where global G_t is used for network fetching, and topic-sensitive WG_t is used for tweet fetching.
- Use the topic-sensitive network parameters for both network and tweet fetching. This is named as the *WG-WG* method (Figure 5.3).

The first approach, *G-WG*, is useful for cases where globally influential users are tracked, but with minimal additional resources, topic-based influential users are to be determined as well. This might be the only viable option if the bandwidth is not enough for selecting and updating the vertices separately for each topic, especially if the number of topics is high.

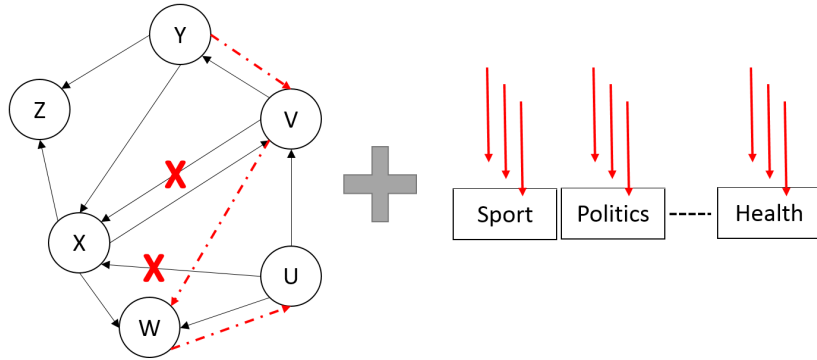


Figure 5.2: G-WG method, probe the global network and probe the tweet sets for each topic of interest.

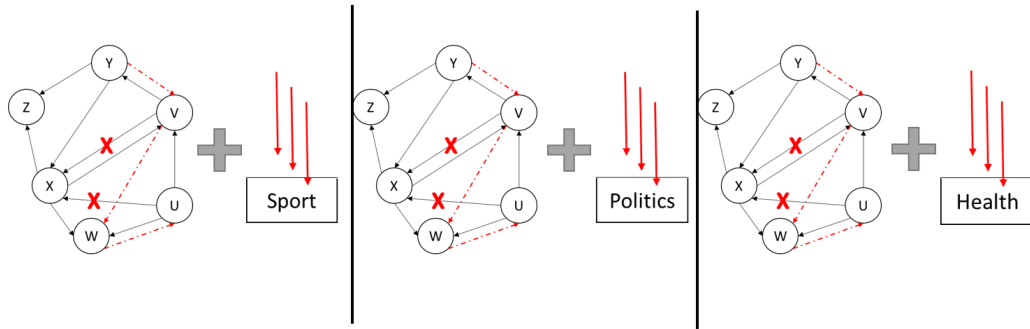


Figure 5.3: WG-WG method, probe the networks and the tweet sets for each topic of interest.

For the second approach, *WG-WG*, we construct separate networks WG^j for each topic and evolve them separately. We update each network at the end of a probing period, using the new tweets fetched to track the most influential vertices for each topic C_j . The high-level algorithms for the *G-WG* and *WG-WG* methods are given in Algorithms 2 and 3, respectively.

Chapter 6

Experiments and Results

In this chapter, we present the experimental setup and the results of our performance evaluation for the proposed algorithms. We also present experiments analyzing the sensitivity of the parameters used in the algorithms.

6.1 Data Sets

We have collected data using the public Twitter API, as described in Chapter 4. Twitter API calls are restricted by rate limit windows. These windows represent 15 minute intervals and the allowed number of calls within each window can vary with respect to the call type. Our system makes two different calls, *a*) “GET followers/ids”, which returns the followers list of the specified user, and *b*) “GET statuses/user_timeline”, which returns the most recent Tweets of the specified user. For the first call type, we are allowed to make 15 calls per rate limit window. Every call can return up to 5K followers. For the users who have more than 5K followers, we have to make multiple calls accordingly. For the second type, we are allowed to make 1804 calls per limit window. Every call can return 3.2K¹ tweets of the queried user. Details of the calls are also presented in Chapter 3 with the accompanying analysis.

¹With paging. Every page contains 200 tweets.

We have collected the network between the end of August 2014 and the beginning of January 2015, with a period of 15-20 days. As a result, we have obtained 11 snapshots of the Turkish users’ network with progressing timestamps. We have collected the relations of 2.8 million users, which amounts to a total of 310 million edges on average. We took the first snapshot as the initial network to calculate the probing scores (see Eq. 5.5) and the rest of the snapshots were used as ground truth for the evaluation of the probing algorithms. For the topic-based influence estimation, we have also collected the tweets of our seed users in the same period. We constructed a dataset formed of 11 snapshots containing 5.5 billion² tweets in total. We take the first snapshot as the initial tweet set as in the case of the relationship network analysis. From this data, we have built up the topic weighted networks and calculated probing scores (see Eq. 5.5), accordingly.

In our probe simulation module, we fetch the connections of the users we have selected for probing, from the real network G_t at time t . We then update these connections (adding new ones and deleting old ones) on the previously observed network G'_{t-1} at time $t - 1$, in order to obtain the estimated network G'_t at time t . Finally, we compare the influence estimation results from the observed network G'_t with the ones from the real network G_t . Same procedure is also applied for the tweet sets.

In order to include extensive number of experiments in our evaluation, we focused on the top $250K$ influential users and restricted the network on which the scores are computed to the network formed by these users.

Figure 6.1 shows the in-edge distribution of the original and the pruned network. Both follow a power-law distribution. Impact of the pruning process on the network structure seems to be minimal and has not created any anomalies in the analysis. We also pruned the tweet list according to the same top $250K$ influential users, which reduced the total size of the tweet sets to $200M$.

²This number includes re-tweets and duplicate tweets as well. In the collection phase, we are fetching last 200 tweets of the users without checking whether or not they exist in the local tweet sets or they are re-tweets. Because this possible checks also require extra API calls. However, re-tweets are not considered in the analysis.




Figure 6.1: In-edge distributions of the original network (on the left) and the pruned network (on the right).

6.1.1 Evaluation of Dynamic Network Fetching

We have implemented several algorithms to compare the performance of the proposed techniques. The details of the algorithms used are given below:

NoProbe and **Random Probing**. These are two baseline algorithms. *NoProbe* algorithm assumes that the network does not change over time and uses the fully observed network at time $t = 0$ for all time points without performing any probing. It represents the worst case scenario for the dynamic network fetching problem. The second baseline algorithm is *Random Probing* algorithm which randomly chooses k users to probe with uniform probability.

MaxG. As described in [39], users are probed with a probability proportional to the “performance gap”, which is defined as the predicted difference between the results of the approximate solution and the real solution. Briefly, the method incrementally probes users which will bring the largest difference in the results. The method assumes that the influence of a specific user is related to the output of the *degree discount heuristic*. Although their influence determination function is different than ours, we use the MaxG algorithm for performance evaluation of our proposed algorithms.

Priority Probing. As described in [38], this algorithm chooses users to probe according to a value proportional to their priorities. Priority of a node is defined as the value of its PageRank score. For every iteration of the method, if a node is not probed, the current PageRank value is added to its priority and if the node is probed, its priority is reset to 0.

Change Probing. This is our first proposed method, which chooses k users to probe with value proportional to their scores, as computed by Eq. 5.5. The network is then constructed via Alg. 1.

Round-Robin & Change Probing. This is our second proposed method, which chooses $\beta \cdot k$ users to probe with Change Probing and $(1 - \beta) \cdot k$ users with Round-Robin Probing. When $\alpha = 0$ in Eq. 5.5 for the Change Probing part, the method becomes similar to [38]. The difference is that Priority Probing increases the probe possibility of a node by its PageRank value in every step if it is not probed, so that at some point the probe possibility becomes 1.

We evaluate performance by comparing the quality of the influential users found by each approach with that of the ideal case. For this purpose, we use two different evaluation measures:

- Jaccard similarity between the correct and estimated top- m most influential users lists.
- The mean squared error of the PageRank scores.

6.1.2 Evaluation of Dynamic Tweet Fetching

We evaluate the performance of the proposed tweet fetching technique with two baselines algorithms, namely *NoProbe* and *Random Probing*. The details of these baselines and our proposed method are given below:

NoProbe. This algorithm assumes that the tweet set does not change over time and use the fully observed tweet set at time $t = 0$ for all time points without any probing. This method represents the worst case scenario for the dynamic tweet fetching problem.

Random Probing. This algorithm randomly chooses k users to collect tweets with uniform probability at each time step.

Round-Robin & Topic Change Proportional Probing. This is the algorithm we proposed, which greedily chooses k users to collect tweets with value proportional to their scores describe in Eq. 5.5. Differently from the network fetching method, scores are

calculated by using WPR_v^j for the topic C_j , instead of PR_v .

6.1.3 Experimental Results and Discussion

This section compares and discusses the performance of the proposed network and tweet probing methods with the state-of-the-art and baseline methodologies using experiments executed on real datasets. We also provide an empirical interpretation of the calculated topic-based influence scores.

6.1.3.1 Experimental Setup

As indicated in Eqs. 3.1 and 3.2, given the resource limits permitted by the service providers, one cannot probe a significant portion of the network. We have executed our experiments with different probing capacities and used 0.001%, 0.01%, 0.1% and 1% of the network as the size of the probe set. For the analysis of the effect of the α parameter used in Change Probing, we set: *a*) $\alpha = 0$, meaning PageRank proportional scores are used; *b*) $\alpha = 0.5$, meaning equally weighted PageRank and influence past scores are used; *c*) $\alpha = 1$, meaning only influence past scores are used. For the Round-Robin Change algorithm we tested the ratio parameter β with three values, which control the random selection: 0.4, 0.6, and 0.8.

6.1.3.2 Change Probing Performance w.r.t. α

Figure 6.2 depicts the performance of Change Probing algorithm for the Jaccard similarity measure. As expected, Change Probing algorithm significantly outperforms NoProb algorithm. For the optimization of the α parameter, we test Change Probing algorithm under three different α configurations:

- Using Average mean squared errors (MSE), $\alpha = 0.5$ setting performs 8% better than $\alpha = 0$ setting and 19% better than $\alpha = 1$ setting. Overall, it performs 83%

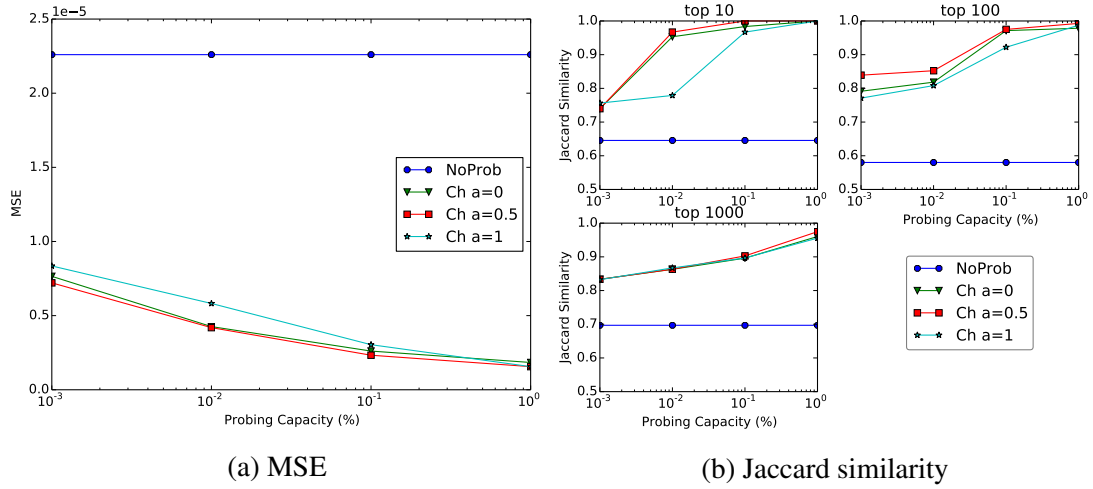


Figure 6.2: Performance of Change Probing.

better than NoProbing.

- Using the Jaccard distance measure, $\alpha = 0.5$ setting is 3% better than $\alpha = 0$ setting and 5% better than $\alpha = 1$ setting. In the overall case, $\alpha = 0.5$ outperforms NoProbe by 43%. We also note that as the probing capacity increases, performance of the Change Probing algorithm becomes less dependent on the setting of α .

We also illustrate the change in error as the network evolves, in order to see how the performance of different algorithms are affected as the seed network data ages. Figures 6.3a and 6.3b show the performance of Change Probing as a function of time for the mean squared error (MSE) and Jaccard similarity metrics, respectively. We observe that NoProb has an increasing error as time passes. Change Probing gives a more robust and stable performance with respect to time. This is mainly because as the number of past influence points increases, the algorithm can estimate the influence variability of the users more accurately, which compensates the deteriorating effect of aging of the baseline network data. Since $\alpha = 0.5$ outperforms the other cases, we use $\alpha = 0.5$ configuration in the subsequent experiments with other algorithms. We also note that y-axis contains relatively small values because the PageRank values are normalized. We have assumed NoProb algorithm as the reference point for normalization.

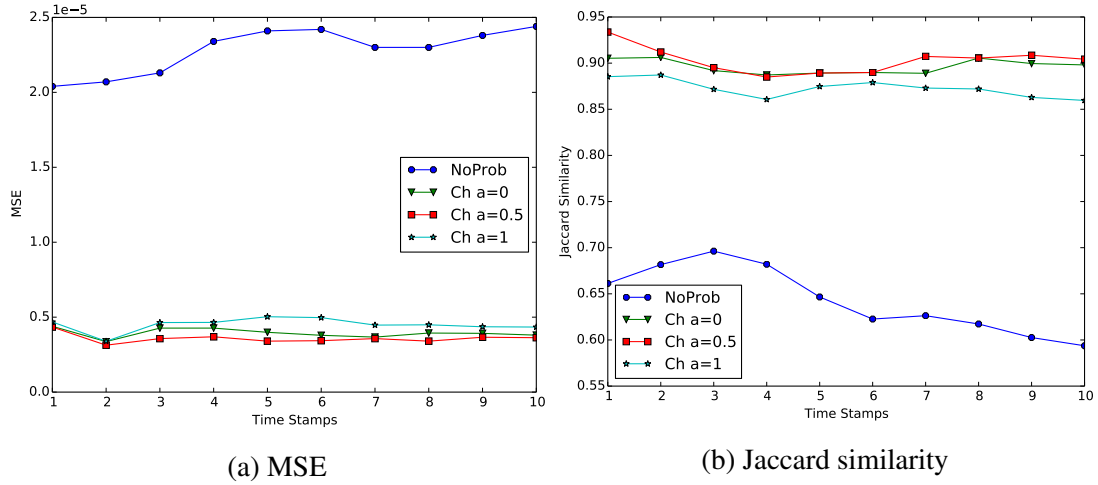


Figure 6.3: Performance of Change Probing as a function of time.

6.1.3.3 RR Change Probing Performance w.r.t. β

Figure 6.4 shows the performance results for the Round-Robin Change (RRCh) Probing algorithm under different round-robin ratios. We use the Change Probing algorithm (with $\alpha = 0.5$ setting) as the baseline reference point.

We observe that the RRCh algorithm performs poorly for small probing capacities, such as 0.001% and 0.01%. Randomness impacts the performance more with smaller number of probed users, since we are not able to probe the influential users with great influential power, thus lowering the performance. For MSE, $\beta = 0.8$ configuration performs 7% better than $\beta = 0.6$ and 12% better than $\beta = 0.4$. For the Jaccard similarity measure, it is 2% better than $\beta = 0.6$ and 7% better than $\beta = 0.4$. Although, it performs worse than Change Probing in the short term, it reaches the performance of Change Probing in the long term, as show in in Figures 6.5a and 6.5b. Moreover, it guarantees the probing of every node within a time frame, preventing the system to focus on only a limited section of the network and missing other regional changes that might accumulate and start to affect the network in the global sense. We would have seen this phenomenon more explicitly if the number of snapshots were larger, which was the case in [39]. The results are slightly better when the ratio is set to $\beta = 0.8$. Therefore, we choose to use this algorithm (with $\alpha = 0.5$ and $\beta = 0.8$ configurations) instead of Change Probing for the comparison with others in the following sections.

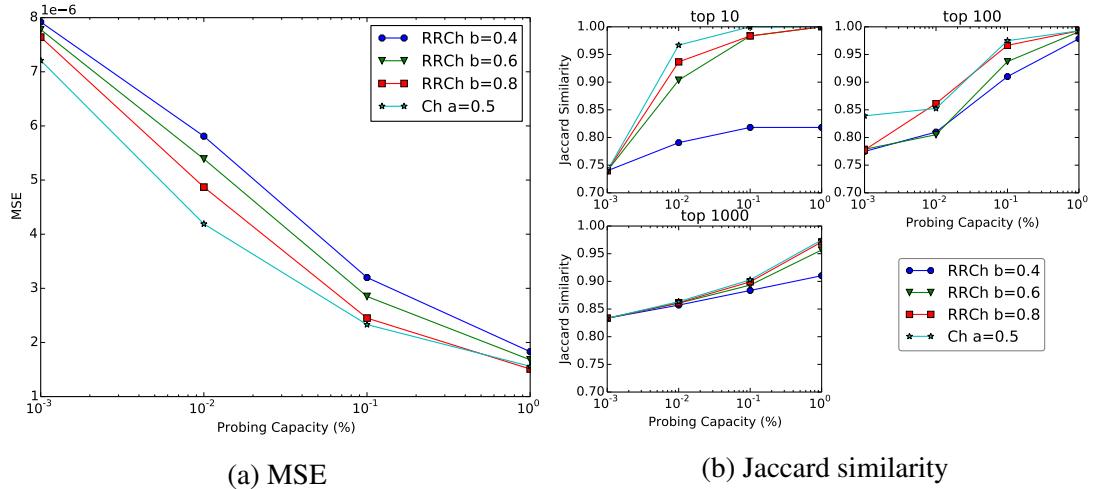


Figure 6.4: Performance of Round-Robin Change Probing.

6.1.3.4 Comparison with the State-of-the-Art

Figure 6.6 compares the performance of RR Change method (with $\alpha = 0.5$ and $\beta = 0.8$ settings) against the baselines and the state-of-the-art methods from the literature. RR Change achieves better results for all performance measures used for comparison in the thesis. It reduces MSE by 21% (see Figure 6.6a) when compared to Priority Probing and 50% when compared to the MaxG method. Priority Probing suffers especially for low probing capacity cases, since the priority of a user is set to 0 after probing. A probed user can regain its priority very late in the process, which prevents it to track quick changes in the scores of the highly influential users. Therefore, after probing an important user in terms of influence, that user is not being probed for some time, even if the influence of the user is changing very fast. RR Change always probes β portion of the users according to their influence impact and change over time, so that the important users are in the probe set in each time step.

Overall, the proposed method gives 81% higher performance than the baseline algorithms for the MSE measure. As seen in Figure 6.6b, RR Change shows better results for the top-m set similarities as well. It is 5% better than Priority Probing and 11% better than MaxG method on average. The performance difference is reaching up-to 18%. RR Change performs 35% better against baselines when Jaccard similarity is considered on average. It is 40% better than the NoProbe and difference reaches up to

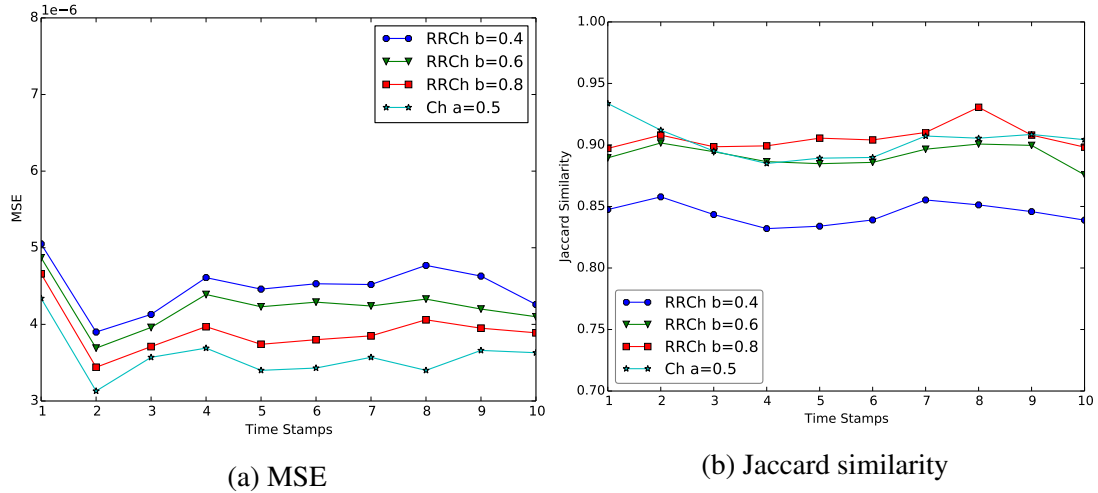


Figure 6.5: Performance of Round-Robin Change Probing as a function of time.

55% for top-100 case. Since it also considers the change in the influence over time, it is also able to preserve its accuracy while the performance of other methods degrade over time (see Figures 6.7a and 6.7b).

6.1.3.5 Evaluation of the Network Inference Method

To assess the prediction quality of the link prediction algorithm, we plotted the histogram of the edges proposed by RA index that has really occurred in the real network. This is shown in Figure 6.8. The histogram indicates the accuracy of the RA index used for network inference. The edges that were determined by the prediction algorithm as more likely to happen were found to be existent in the future network with a higher probability. However, when we analyzed the incorrectly predicted edges, we have observed that the algorithm predicts links between users who are unlikely to follow each other in real life. For example, the algorithms predict an edge between two pop stars since they have many common neighbors. However, they would not follow each other because they are main competitors. Furthermore, some of these users not willing to follow anybody at all. Link prediction algorithms typically do not consider these facts in social networks. This indicates a weakness of the “mechanical” link prediction algorithms on social networks. In addition to indexes which they use to calculate similarities between users, they should also consider the tendency of the users

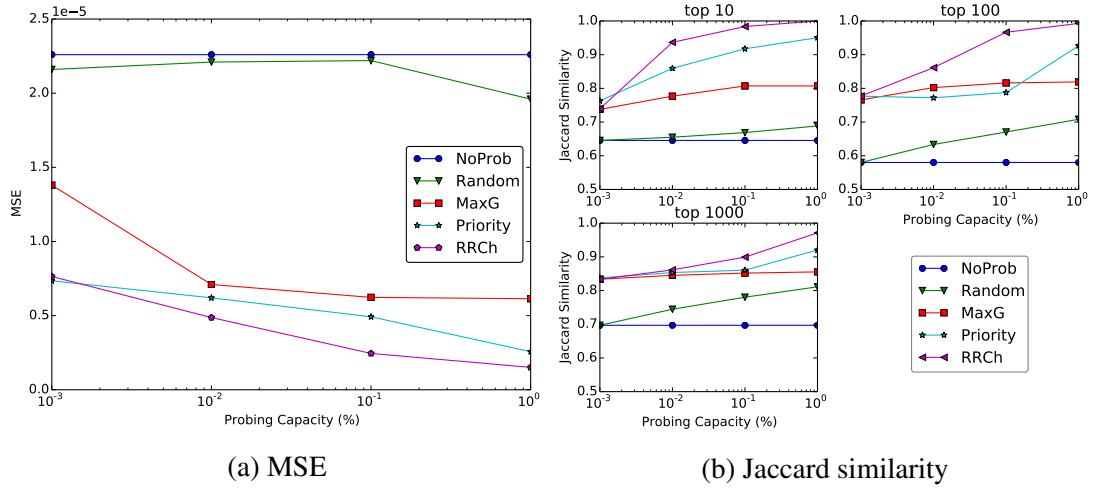


Figure 6.6: Comparison of the probing strategies.

to make new connections. Therefore, we apply a filtering process in order to determine users who are likely to follow somebody and we add the predicted edges only to those selected users.

As a result, we improve the RR Change method by 3% for MSE and 2% for the set similarities on average. Figure 6.9 compares the performance of our inference method against the baselines, the state-of-the-art methods and the RR Change method. Especially, it increases the performance of RR Change for the lower capacities e.g., 0.001% and 0.01%. In Figure 6.9b, we observed 7% improvement on the top-10 jaccard similarities for 0.001% and 0.01% probing capacities.

6.1.3.6 Evaluation of the Topic Influence Estimation

We evaluated the influence of users with respect to four different topics: *a)* Politics, *b)* Sport, *c)* Health, and *d)* Cultural and Art Activities. This section provides a qualitative discussion about the accounts which were found to be influential by the proposed methods. Table 6.1 shows the accuracy of topic relevance of the top-10 users found by the system for the specific topics.

For the evaluation of the results, we performed a small survey containing 10 people. We asked participants to evaluate the users with respect to their topic relevance

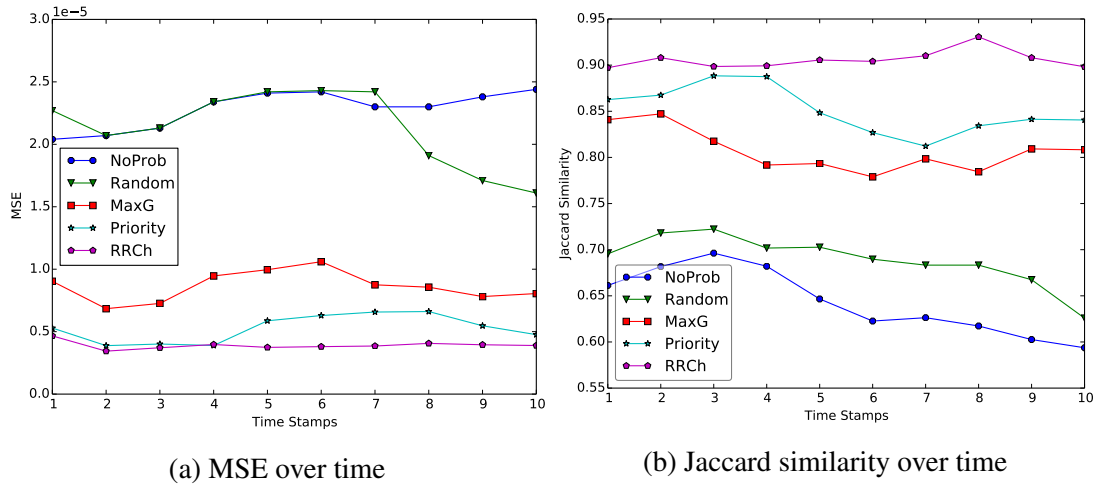


Figure 6.7: Comparison of the Probing strategies with respect to time.




Figure 6.8: Accuracy of the link prediction algorithm.

and their influence on the topic. In order to identify influence of a user, we asked participants to mark one of the following categories: *a*) very influential (1), *b*) influential (.5), *c*) not influential (0). We used the results of the survey to provide an evaluation of the selected users for the Turkish Twitter network, on a per-topic basis.

For the topic Politics, the results are very accurate for top-10. We have observed that the dictionaries constructed for each topic has a big impact on the results. For example, we observe that the dictionary constructed for Politics topic contains many keywords that are related only with politics without any ambiguity. These keywords have increased the performance of the semantic analysis, which in turn increased the

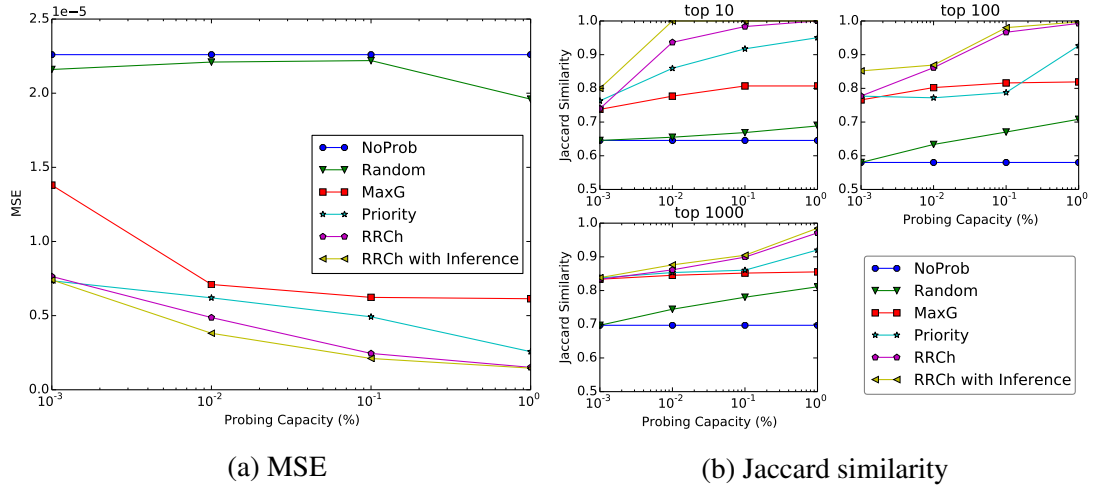


Figure 6.9: Performance of RR Change with inference.

Topics	Topic Relevance
<i>Politics</i>	<i>10 out of 10</i>
<i>Sport</i>	<i>8.5 out of 10</i>
<i>Health</i>	<i>4 out of 10</i>
<i>Cultural and Art Activities</i>	<i>9 out of 10</i>

Table 6.1: Estimated influential accounts.

accuracy of the topic-based network influence analysis. Top-10 list contains the president of Turkish Republic (RT_Erdogan), the chairman of one of the opposition parties (kilicdarogluk), and the mayor of the capital city (06melihgokcek). It is fair to assume that these users, who give political messages on their tweets and who have lots of followers, should be in the top-10 influential list on Turkish Politics topic.

The influential accounts for the Sport topic were the biggest sport clubs of Turkey (GalatasaraySK, Fenerbahce) and one of the highest rating sport channel (ntvspor). Their tweets were mostly related about the sport competitions, news from clubs, etc. They have a lot of followers who actively pay attention to what they tweet. Thus, they achieve high RT and Fav statistics, which shows that they have a big impact on their followers. It is very reasonable that they are top influential accounts on this topic.

As intuitively expected, the influential accounts for the Health topic are mostly doctor associations and governmental authorities. One of the accounts is Republic

of Turkey Ministry of Health (*saglikbakanligi*), which mainly tweets about hospitals, doctors and health regulations. Its follower numbers can be considered as relatively high and is followed by other influential accounts. Since its tweets have critical news potential, it has considerable number of RTs about the health topic. The other two are doctor associations (*YYD_tr*, *istabip*). They are followed by many doctors, which also have some potential impact on the Health topic. In this topic, accurate relevance ratio is relatively low because the constructed dictionary for this topic is not specific enough, causing errors in semantic analyses that propagates to the latter phase of influence estimation.

The Cultural and Art Activities topic includes users which tweet about movies, art, books, history, etc. The top-10 influential users are perfectly matched with the keywords. *CMYLMZ* is very famous Turkish comedian, actor and producer. He also has one of the highest follower numbers in the Turkish Twitter network. *AtlasTarihDergi* is a history magazine tweeting mainly about historical events and information which has considerable amount of followers and RTs. The third user (*Siirler_sokakta*) shares street poems and mottos, and it's posts receive many RTs and Favs.

6.1.3.7 Evaluation of Dynamic Tweet Fetching

We have used the same default parameter settings from the network fetching experiments to evaluate our proposed tweet fetching methods.

Figure 6.10 shows the performance of the RR Change method for dynamic tweet fetching. For the MSE measure, global network based *G-WG* method performs 78% better, and topic network based *WG-WG* method performs 40% better than the baselines, on average, respectively. In Figure 6.10b, we see that as the probing capacities increase, *G-WG* method achieves almost perfect similarity against the results obtained using the original network, for the top-10 influential users. For the top-1000 influential users experiment, it reaches close to 0.9 similarity. Together with *WG-WG* method, they quickly reach close to their top performance at around 1% capacity, except for the top-10 case. For the latter, *WG-WG* method does not enjoy the quality increase

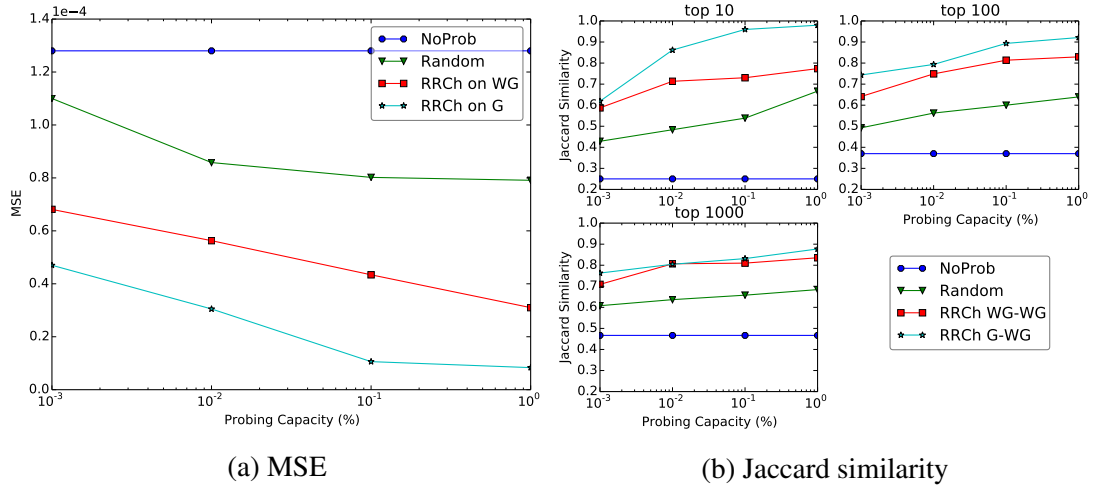


Figure 6.10: Performance of Change Probing for dynamic tweet fetching.

that the $G - WG$ method enjoys with increasing capacities. When we look at the Jaccard similarity based results, $G-WG$ achieves 77% better and $WG-WG$ achieves 65% better results than the baselines. Overall, the results show us that using the globally maintained network is more advantageous.

Although $G-WG$ method outperforms $WG-WG$ method when we compare the top-10 results for the two methods, they are similar in terms of the topic relevance of their top influential users. Table 6.2 shows the topic relevance ratios for the two methods. Top-10 selected users are found to be related with the topics of interest and are popular accounts in the topic area.

Topics	Topic Relevance	
<i>Politics</i>	<i>10 out of 10</i>	<i>10 out of 10</i>
<i>Sport</i>	<i>8 out of 10</i>	<i>9 out of 10</i>
<i>Health</i>	<i>5 out of 10</i>	<i>4 out of 10</i>
<i>Cultural and Art Activities</i>	<i>9 out of 10</i>	<i>9 out of 10</i>
	<i>G - WG method</i>	<i>WG - WG method</i>

Table 6.2: Top-10 topic relevance ratios for $G-WG$ and $WG-WG$ for dynamic tweet fetching.

Chapter 7

Conclusion

Restrictions applied by social network service providers limit the capabilities of evolving network analyses by third parties. Therefore, we have proposed algorithms to dynamically fetch network and text data from public API's under limited probing capacities. Our proposed solutions use the past influential trends of the users with their current statuses in order to predict their changes for the next fetching periods. In particular, we observed that high influential users and changeable influential trends effect the overall network structure the most. We leverage these two metrics across our algorithms. Experimental results showed that also considering past trends in the probing strategy increases the overall accuracy of observed networks compared to only using current metrics. Furthermore, we improved our probing strategies by inferring possible relations between users via link prediction algorithms. Due to computational ease, we choose weighting friendship based inference in our system among alternatives like bayes classification, bayes on friendship links and collective inference. Moreover, our inference method could be extended by applying more complex network evolution models in order to increase the accuracy of estimation probabilities.

We have also developed models for estimating global and topic-based influence in social networks. Our influence models consider the dynamic structure of the network for the influence estimation. For global influence, we considered the topological place

of users within the network. For topic-based influence, we further enhanced the traditional relation based graph model by assigning new features to relations. In particular, we used topological place of users together with their topic analysis and sharing statistics which helped us separating relevant and irrelevant contents of their posts. With this way, we could be able to specialize our social network model for topics of interest.

Bibliography

- [1] “Twitter.” <https://twitter.com/>, 2006.
- [2] “Twitter API rate limits.” <https://dev.twitter.com/rest/public/rate-limiting>, 2012.
- [3] L. Page, S. Brin, R. Motwani, and T. Winograd, “The pagerank citation ranking: Bringing order to the web,” 1999.
- [4] W. Xing and A. Ghorbani, “Weighted pagerank algorithm,” in *Communication Networks and Services Research, 2004. Proceedings. Second Annual Conference on*, pp. 305–314, IEEE, 2004.
- [5] S. Wasserman, *Social network analysis: Methods and applications*, vol. 8. Cambridge University Press, 1994.
- [6] P. Domingos and M. Richardson, “Mining the network value of customers,” in *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 57–66, ACM, 2001.
- [7] D. Kempe, J. Kleinberg, and É. Tardos, “Maximizing the spread of influence through a social network,” in *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146, ACM, 2003.
- [8] D. Kempe, J. Kleinberg, and É. Tardos, “Influential nodes in a diffusion model for social networks,” in *Automata, languages and programming*, pp. 1127–1138, Springer, 2005.
- [9] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and N. Glance, “Cost-effective outbreak detection in networks,” in *Proceedings of the 13th ACM*

SIGKDD international conference on Knowledge discovery and data mining, pp. 420–429, ACM, 2007.

- [10] W. Chen, C. Wang, and Y. Wang, “Scalable influence maximization for prevalent viral marketing in large-scale social networks,” in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1029–1038, ACM, 2010.
- [11] M. Cha, H. Haddadi, F. Benevenuto, and P. K. Gummadi, “Measuring user influence in twitter: The million follower fallacy,” *ICWSM*, vol. 10, no. 10-17, p. 30, 2010.
- [12] A. Anagnostopoulos, R. Kumar, and M. Mahdian, “Influence and correlation in social networks,” in *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 7–15, ACM, 2008.
- [13] J. Tang, S. Wu, and J. Sun, “Confluence: Conformity influence in large social networks,” in *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 347–355, ACM, 2013.
- [14] W. Chen, Y. Wang, and S. Yang, “Efficient influence maximization in social networks,” in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 199–208, ACM, 2009.
- [15] A. Goyal, F. Bonchi, and L. V. Lakshmanan, “Learning influence probabilities in social networks,” in *Proceedings of the third ACM international conference on Web search and data mining*, pp. 241–250, ACM, 2010.
- [16] X.-L. Li, C. S. Foo, K. L. Tew, and S.-K. Ng, “Searching for rising stars in bibliography networks,” in *Database Systems for Advanced Applications*, pp. 288–292, Springer, 2009.
- [17] A. Daud, R. Abbasi, and F. Muhammad, “Finding rising stars in social networks,” in *Database Systems for Advanced Applications*, pp. 13–24, Springer, 2013.
- [18] Q. Mei, D. Cai, D. Zhang, and C. Zhai, “Topic modeling with network regularization,” in *Proceedings of the 17th international conference on World Wide Web*, pp. 101–110, ACM, 2008.

- [19] J. Tang, J. Sun, C. Wang, and Z. Yang, “Social influence analysis in large-scale networks,” in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 807–816, ACM, 2009.
- [20] L. Liu, J. Tang, J. Han, M. Jiang, and S. Yang, “Mining topic-level influence in heterogeneous networks,” in *Proceedings of the 19th ACM international conference on Information and knowledge management*, pp. 199–208, ACM, 2010.
- [21] J. Weng, E.-P. Lim, J. Jiang, and Q. He, “Twitterrank: finding topic-sensitive influential twitterers,” in *Proceedings of the third ACM international conference on Web search and data mining*, pp. 261–270, ACM, 2010.
- [22] L. Hong and B. D. Davison, “Empirical study of topic modeling in twitter,” in *Proceedings of the First Workshop on Social Media Analytics, SOMA '10*, (New York, NY, USA), pp. 80–88, ACM, 2010.
- [23] D. M. Blei, A. Y. Ng, and M. I. Jordan, “Latent dirichlet allocation,” *the Journal of machine Learning research*, vol. 3, pp. 993–1022, 2003.
- [24] M. Rosen-Zvi, T. Griffiths, M. Steyvers, and P. Smyth, “The author-topic model for authors and documents,” in *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pp. 487–494, AUAI Press, 2004.
- [25] C. X. Lin, Q. Mei, J. Han, Y. Jiang, and M. Danilevsky, “The joint inference of topic diffusion and evolution in social communities,” in *Data Mining (ICDM), 2011 IEEE 11th International Conference on*, pp. 378–387, IEEE, 2011.
- [26] L. Hong, O. Dan, and B. D. Davison, “Predicting popular messages in twitter,” in *Proceedings of the 20th international conference companion on World wide web*, pp. 57–58, ACM, 2011.
- [27] G. Szabo and B. A. Huberman, “Predicting the popularity of online content,” *Commun. ACM*, vol. 53, pp. 80–88, Aug. 2010.
- [28] J. Yang and J. Leskovec, “Patterns of temporal variation in online media,” in *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, WSDM '11*, (New York, NY, USA), pp. 177–186, ACM, 2011.

- [29] J. Cheng, L. Adamic, P. A. Dow, J. M. Kleinberg, and J. Leskovec, “Can cascades be predicted?,” in *Proceedings of the 23rd international conference on World wide web*, pp. 925–936, International World Wide Web Conferences Steering Committee, 2014.
- [30] S. Chen, J. Fan, G. Li, J. Feng, K.-I. Tan, and J. Tang, “Online topic-aware influence maximization,” *Proceedings of the VLDB Endowment*, vol. 8, no. 6, pp. 666–677, 2015.
- [31] J. L. Wolf, M. S. Squillante, P. Yu, J. Sethuraman, and L. Ozsen, “Optimal crawling strategies for web search engines,” in *Proceedings of the 11th international conference on World Wide Web*, pp. 136–147, ACM, 2002.
- [32] J. Cho and H. Garcia-Molina, “Effective page refresh policies for web crawlers,” *ACM Trans. Database Syst.*, vol. 28, pp. 390–426, Dec. 2003.
- [33] J. Cho and H. Garcia-Molina, “Estimating frequency of change,” *ACM Transactions on Internet Technology (TOIT)*, vol. 3, no. 3, pp. 256–290, 2003.
- [34] S. Pandey and C. Olston, “User-centric web crawling,” in *Proceedings of the 14th International Conference on World Wide Web, WWW ’05*, (New York, NY, USA), pp. 401–411, ACM, 2005.
- [35] C. Olston and S. Pandey, “Recrawl scheduling based on information longevity,” in *Proceedings of the 17th International Conference on World Wide Web, WWW ’08*, (New York, NY, USA), pp. 437–446, ACM, 2008.
- [36] J. Leskovec, L. Backstrom, R. Kumar, and A. Tomkins, “Microscopic evolution of social networks,” in *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 462–470, ACM, 2008.
- [37] P. Desikan and N. Pathak, “Incremental pagerank computation on evolving graphs,” in *WWW*, pp. 10–14, 2005.
- [38] B. Bahmani, R. Kumar, M. Mahdian, and E. Upfal, “Pagerank on an evolving graph,” in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 24–32, ACM, 2012.

- [39] H. Zhuang, Y. Sun, J. Tang, J. Zhang, and X. Sun, "Influence maximization in dynamic social networks," in *Data Mining (ICDM), 2013 IEEE 13th International Conference on*, pp. 1313–1318, Dec 2013.
- [40] G. Valkanas, I. Katakis, D. Gunopulos, and A. Stefanidis, "Mining twitter data with resource constraints," in *Proceedings of the 2014 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT) - Volume 01, WI-IAT '14*, (Washington, DC, USA), pp. 157–164, IEEE Computer Society, 2014.
- [41] A. Nazi, Z. Zhou, S. Thirumuruganathan, N. Zhang, and G. Das, "Walk, not wait: Faster sampling over online social networks," *arXiv preprint arXiv:1410.7833*, 2014.
- [42] L. Getoor, N. Friedman, D. Koller, and B. Taskar, "Learning probabilistic models of link structure," *The Journal of Machine Learning Research*, vol. 3, pp. 679–707, 2003.
- [43] Z. Ghahramani, "Learning dynamic bayesian networks," in *Adaptive processing of sequences and data structures*, pp. 168–197, Springer, 1998.
- [44] L. Song, M. Kolar, and E. P. Xing, "Time-varying dynamic bayesian networks," in *Advances in Neural Information Processing Systems*, pp. 1732–1740, 2009.
- [45] J. H. Koskinen and T. A. Snijders, "Bayesian inference for dynamic social network data," *Journal of statistical planning and inference*, vol. 137, no. 12, pp. 3930–3938, 2007.
- [46] J. Bonneau, J. Anderson, R. Anderson, and F. Stajano, "Eight friends are enough: social graph approximation via public listings," in *Proceedings of the Second ACM EuroSys Workshop on Social Network Systems*, pp. 13–18, ACM, 2009.
- [47] M. Gomez Rodriguez, J. Leskovec, and A. Krause, "Inferring networks of diffusion and influence," in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1019–1028, ACM, 2010.
- [48] S. Myers and J. Leskovec, "On the convexity of latent social network inference," in *Advances in Neural Information Processing Systems*, pp. 1741–1749, 2010.

- [49] J. Yang and J. Leskovec, “Modeling information diffusion in implicit networks,” in *Proceedings of the 2010 IEEE International Conference on Data Mining, ICDM '10*, (Washington, DC, USA), pp. 599–608, IEEE Computer Society, 2010.
- [50] C. Kang, C. Molinaro, S. Kraus, Y. Shavitt, and V. Subrahmanian, “Diffusion centrality in social networks,” in *Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2012)*, pp. 558–564, IEEE Computer Society, 2012.
- [51] J. Yang and J. Leskovec, “Modeling information diffusion in implicit networks,” in *Data Mining (ICDM), 2010 IEEE 10th International Conference on*, pp. 599–608, IEEE, 2010.
- [52] M. G. Rodriguez, D. Balduzzi, and B. Schölkopf, “Uncovering the temporal dynamics of diffusion networks,” *arXiv preprint arXiv:1105.0697*, 2011.
- [53] B. Taskar, M.-F. Wong, P. Abbeel, and D. Koller, “Link prediction in relational data,” in *Advances in neural information processing systems*, p. None, 2003.
- [54] J.-P. Vert and Y. Yamanishi, “Supervised graph inference,” in *Advances in Neural Information Processing Systems*, pp. 1433–1440, 2004.
- [55] D. Liben-Nowell and J. Kleinberg, “The link-prediction problem for social networks,” *Journal of the American society for information science and technology*, vol. 58, no. 7, pp. 1019–1031, 2007.
- [56] B. Wang, C. Wang, J. Bu, C. Chen, W. V. Zhang, D. Cai, and X. He, “Whom to mention: expand the diffusion of tweets by@ recommendation on microblogging systems,” in *Proceedings of the 22nd international conference on World Wide Web*, pp. 1331–1340, International World Wide Web Conferences Steering Committee, 2013.
- [57] B. Bahmani, A. Chowdhury, and A. Goel, “Fast incremental and personalized pagerank,” *Proceedings of the VLDB Endowment*, vol. 4, no. 3, pp. 173–184, 2010.

- [58] M. Naaman, J. Boase, and C.-H. Lai, “Is it really about me?: Message content in social awareness streams,” in *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work, CSCW '10*, (New York, NY, USA), pp. 189–192, ACM, 2010.
- [59] R. Kelly, “Twitter study,” *San Antonio, TX: Pear Analytics*, 2009.
- [60] L. Lü and T. Zhou, “Link prediction in complex networks: A survey,” *Physica A: Statistical Mechanics and its Applications*, vol. 390, no. 6, pp. 1150–1170, 2011.