A topology-based approach for followees recommendation in Twitter

Marcelo G. Armentano, Daniela L. Godoy and Analía A. Amandi
Universidad Nacional del Centro de la Provincia de Buenos Aires
CONICET, Consejo Nacional de Investigaciones Científicas y Técnicas
Argentina
{marmenta,dgodoy,amandi}@exa.unicen.edu.ar

Abstract

Nowadays, more and more users keep up with news through information streams coming from real-time micro-blogging activity offered by services such as Twitter. In these sites, information is shared via a followers/followees social network structure in which a follower will receive all the micro-blogs from the users he follows, named followees. Recent research efforts on understanding micro-blogging as a novel form of communication and news spreading medium have identified different categories of users in Twitter: information sources, information seekers and friends. Users acting as information sources are characterized for having a larger number of followers than followees, information seekers subscribe to this kind of users but rarely post tweets and, finally, friends are users exhibiting reciprocal relationships. With information seekers being an important portion of registered users in the system, finding relevant and reliable sources becomes essential. To address this problem, we propose a followee recommender system based on an algorithm that explores the topology of followers/followees network of Twitter considering different factors that allow us to identify users as good information sources. Experimental evaluation conducted with a group of users is reported, demonstrating the potential of the approach.

1 Introduction

Micro-blogging activity taking place in sites such as Twitter is becoming every day more important as real-time information source and news spreading medium. In the followers/followees social structure defined in Twitter a follower will receive all the micro-blogs from the users he follows, known as followees, even though they do not necessarily follow him back. In turn, re-tweeting allows users to spread information beyond the followers of the user that post the tweet in the first place

Studies conducted to understand Twitter usage [Java et al., 2007; Krishnamurthy et al., 2008] revealed that few users maintain reciprocal relationships with other users, which can be regarded as friends or acquaintances, while most of them behave either as information sources or information seekers. Users behaving as information sources tend to collect a large amount of followers as they are actually posting useful information or news. In turn, information seekers follow several users to obtain the information they are looking for and rarely post any tweet themselves.

Finding high quality sources among the expanding micro-blogging community using Twitter becomes essential for information seekers in order to cope with information overload. In this paper we present a topology-based followee recommendation algorithm aiming at identifying potentially interesting users to follow in the Twitter network. This algorithm explores the graph of connections starting at the target user (the user to whom we wish to recommend previously unknown followees), selects a set of candidate users to recommend and ranks them according to a scoring function that favors those users exhibiting the distinctive behavior of information sources.

Unlike other works that focus on ranking users according to their influence in the entire network [Weng et al., 2010; Yamaguchi et al., 2010], the algorithm we propose explores the follower/following relationships of the user up to a certain level, so that more personalized factors are considered in the selection of candidates for recommendation, such as the number of friends in common with the target user. Since only the topology of the social structure is used but not the content of tweets, this algorithm also differs from works exploiting user-generated content in Twitter to filter information streams [Chen et al., 2010; Phelan et al., 2009; Esparza et al., 2010] or to extract topic-based preferences for recommendation [Hannon et al., 2010].

The rest of this paper is organized as follows. Section 2 reviews related research in the area. Section 3 describes our approach to the problem of followee recommendation in Twitter. In Section 4 we present the experiments we performed to validate our proposal and in Section 5 we present and discuss the results obtained and in Section 6 we compared our results with a related approach. Finally, in Section 7, we discuss some aspects of our proposal and present our conclusions.

2 Related Work

The problem of helping users to find and to connect with people on-line to take advantage of their friend relationships has been studied in the context of traditional human social networks. For example, SONAR [Guy *et al.*, 2009] recommends related people in the context of enterprises by aggregating information about relationships as reflected in different sources within an organization, such as organizational chart relationships, co-authorship of papers, patents, projects and others. Chen et al. [Chen *et al.*, 2009] compared relationship-based and content-based algorithms in making people recommendations, finding that the first ones are better at finding known contacts whereas the second ones are stronger at discovering new friends. Weighted minimum-message ratio (WMR) [Lo and Lin, 2006] is a graph-based algorithm which generates a personalized list of friends in social network build according to the observed interaction among members. Unlike these algorithms that gathered social networks in enclosed domains, mainly starting from structured data (such as interactions, co-authorship relations, etc.), we propose a people recommendation algorithms that take advantage of Twitter social structure populated by massive, unstructured and user-generated content.

Understanding micro-blogging as a novel form of communication and news spreading medium has been one of the primary concerns of recent research efforts. Kwak *et al.* [2010] analyzed the topological characteristics of Twitter and its power for information sharing, finding some divergences between this follower/followees network and traditional human social networks: follower distribution exhibit a non-power-law (users have more followers than predicted by power-law), the degree of separation is shorter than expected and there is a low reciprocity (most followers in Twitter do not follow their followers back). Other works addressed the problem of detecting influential users as a method of ranking people for recommendation. In the previous study it was found that ranking users by the number of followers and by PageRank give similar results. However, ranking users by the number of re-tweets indicates a gap between influence inferred from the number of followers and that inferred from the popularity of user tweets. Coincidently, a comparison of in-degree, re-tweets and mentions as influence indicators carried out in [Cha *et al.*, 2010] concluded that the first is more related to user popularity, whereas influence is gained only through a concentrated effort in spawning re-tweets and mentions and can be hold over a variety of topics. TwitterRank [Weng *et al.*, 2010] tries to find influential twitterers by taking into account the topical similarity between users as well as the link structure, TU-Rank [Yamaguchi *et al.*, 2010] considers the social graph and the actual tweet flow and Garcia and Amatriain [2010] propose a method to weight popularity and activity of links for ranking users.

The influence rankings presented by studies on the complete Twittersphere have not direct utility for followee recommendation since people get connected for multiple reasons. We demonstrated with our experiments that indegree, which has proven to be a good representation of a user's influence in Twitter using only its topology (see for example [Kwak et al., 2010]) gives the worst results for followee recommendation since people that are popular in Twitter would not necessarily match a particular user interests (if a user follows accounts talking about technology, he/she would not be interest in Ashton Kutcher, one of the most influential Twitter accounts according to Kwak et al. [2010])

Recommendation technologies applied to Twitter have mainly focused on taking advantage of the massive amount of user-generated content as a novel source of preference and profiling information [Chen *et al.*, 2010, Phelan *et al.*, 2009, Esparza *et al.*, 2010]. In contrast, we concentrate in recommending interesting people to follow. In this direction, Sun *et al.* [2009] proposes a diffusion-based micro-blogging recommendation framework which identifies a small number of users playing the role of news reporters and recommends them to information seekers during emergency events. Closest to our work are the algorithms for recommending followees in Twitter evaluated and compared using a subset of users in [Hannon *et al.*, 2010]. Multiple profiling strategies were considered according to how users are represented in a content-based approach (by their own tweets, by the tweets of their followees, by the tweets of their followers, by the combination of the three), a collaborative filtering approach (by the IDs of their followees, by the IDs of their followers or a combination of the two) and two hybrid algorithms. User profiles are indexed and recommendations generated using a search engine, receiving a ranked-list of relevant Twitter users based on a target user profile or a specific set of query terms. Our work differs from this approach in that we do not require indexing profiles from Twitter users; instead a topology-based algorithm explored the follower/followee network in order to find candidate users to recommend.

The main difference between existent work and our work is that the mentioned approaches for followee recommendations, except for the approach presented in [Hannon *et al.*, 2010], were evaluated using datasets gathered from Twitter, with no assessment about the target user interest in the recommendations. In other words, the target user interest in a followee recommended that is not in the current list of the target user's followees cannot be assessed within these datasets in order to determinate the correctness of the recommendation. For this reason, the approach proposed in this work was evaluated with a controlled experiment with real users.

3 Followees Recommendations on Twitter

The algorithm we propose for recommending followees on Twitter consists in two steps: (1) we explore the target user's neighborhood in search of candidates and (2) we rank candidates according to different weighting features. These steps are detailed in Sections 3.1 and 3.2, respectively.

3.1 Finding candidates

The general idea of the algorithm we implemented is to suggest users that are in the neighborhood of the target user,

where the neighborhood of a user is determined from the follower/followee relations in the social network.

In order to find candidate followees to recommend to a target user U, we based our search algorithm on the following hypothesis: The users followed by the followers of U's followees are possible candidates to recommend to U. In other words, if a user F follows a user that is also followed by U, then other people followed by F can be interesting to U.

The rationale behind this hypothesis is that the target user is an information seeker that has already identified some interesting users acting as information sources, which are his/her current followees. Other people that also follows some of the users in this group (i.e. is subscribe to some of the same information sources) have interests in common with the target user and might have discover other relevant information sources in the same topics, which are in turn their followees.

This scheme is outlined in Figure 1 and can be resumed in the following steps:

1. Starting with the target user, we first obtain the list of users he/she follows, let's call this list S.

$$S = \bigcup_{s \in followees(U)} s$$

2. From each element in S we get its followers, let's call the union of all these lists L

$$L = \bigcup_{s \in S} followers(s)$$

3. Finally, from each element in L, we get its followees to obtain the list of possible candidates to recommend. Let's call the union of all these lists T.

$$T = \bigcup_{l \in L} followees(l)$$

4. Exclude from T those users that the target user already follows. Let's call the resulting list R.

$$R = T - S$$

Each element in R is a possible user to recommend to the target user. Notice that each element can appear more than once in R, depending on the number of times that each user appears in the followees or followers lists obtained at steps 2 and 3 above.

3.2 Weighting features

Once we find the list R of candidate recommendations for the target user, we explored different features to give a score to each unique user $x \in R$.

The first feature explored is the relation between the number of followers a user has with respect to the number of users the given user follows, as shown in Equation 1.

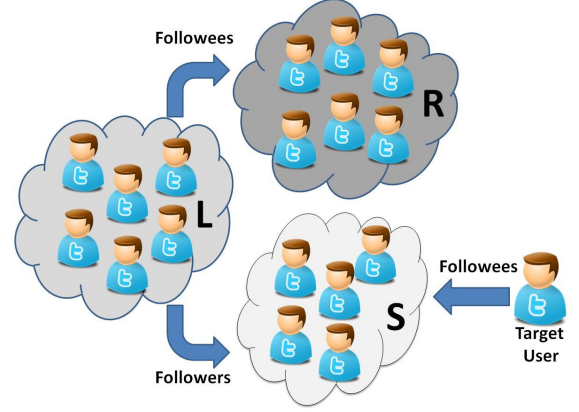$$w_f(x) = \frac{|followers(x)|}{|followees(x)|} \qquad (1)$$

Figure 1: Scheme for finding candidate recommendations

Since we seek for sources of information to recommend, we assume that this kind of users will have a lot of followers and that they will follow few people. If user x has no followees, then only the number of followers is considered without changing the significance of the weighting feature.

We use this metric as a baseline for comparison with other metrics. Our aim is to demonstrate that metrics for ranking popular users on Twitter are not good for ranking recommendations of users that a target user might be interest in following. In [Kwak *et al.*, 2010] it has been shown that the rankings of users that can be obtained by number of followers and by PageRank [Brin and Page, 1998] are very similar. We opted to use this factor as an estimator of the "importance" of a given user because the number of followers is a metric by far more easily to obtain that the user PageRank in a network with an order of almost 2 billion social relations.

The second feature explored corresponds to the number of occurrences of the candidate user in the final list of |R| candidates for recommendations, as shown in Equation 2.

$$w_o(x) = \frac{|\{i \in R / i = x\}|}{|R|} \qquad (2)$$

The number of occurrences of a given user x in this final list is an indicator of the amount of (indirect) neighbors that also have x as a (direct) connection itself.

The third feature we considered is the number of friends in common between the target user U and the candidate recommendation x:

$$w_c(x) = |followees(x) \cap followees(U)| \qquad (3)$$

Finally, we considered two combinations of these features: the average of the three features, and their product:

$$w_s(x) = \frac{1}{3}[w_o(x) + w_f(x) + w_c(x)] \qquad (4)$$

$$w_p(x) = w_o(x) * w_f(x) * w_c(x) \qquad (5)$$

It is worth noticing that the selection of these weighting features was not arbitrary. Our choice was based on a deep analysis of previous studies about Twitter and particular properties of this specific network that makes general link prediction approaches unsuitable. All the studies about the properties of the Twitter network agree in that there is a minimal overlap with the features available on other online social networks (OSNs).

4 Experiment setting

To evaluate the proposed algorithms, we have carried out a preliminary experiment using a group of 14 users. These users, 8 males and 6 females, were in the last years of their course of studies and were students of a Recommender Systems related course dictated at our University as an elective course during 2010. The students selected for the experiment were volunteers familiarized with Twitter.

During the first part of the course, we asked these users to create a Twitter account and to follow at least 20 Twitter users who publish information or news about a set of particular subjects of their interest. The general interests expressed by users ranged between diverse subjects such as technology, software, math, science, football, tennis, basket, religion, movies, journalists, government, music, cooking, shoes, TV programs and even other students in their faculty. Some users only concentrated on one particular subject while others distributed their followees among several topics.

Then, we used the user IDs of the user accounts created by the students as seeds to crawl a sub-graph of the Twitter network corresponding to three levels of both followee and follower relations, centered on each seed. The resulting dataset consisted on 1,443,111 Twitter users and 3,462,179 *following* relations already existing among them.

During the second part of the course, we provided these users with a desktop tool that allowed them to login to Twitter and ask for followees recommendations. Since the users who participated in the experiment were students of a "recommender systems" course, all of them had knowledge about concepts such as rankings and metrics. As part of a not compulsory practical exercise of the course they were motivated to discover which metric better ranked recommendation results and to write a brief report about the results they obtained for their particular case. The desktop application provided for this exercise allowed students to select the weighting feature by which they liked to rank recommendations, with no predefined order.

In all cases, 20 recommendations were presented to the users. Then, we asked the students to explicitly evaluate whether the recommendations were relevant or not according to the same topical criteria they have chosen to select their followees as information sources in the first place. For each recommendation in the resulting ranking the application showed the user name, description, profile picture and a link to the home page of the corresponding account. This link could be used to read the tweets published by the recommended user in the case that the information provided by the application was not enough to determine the student's interest in the recommendation. The question we asked students to ask themselves to determine whether a recommendation was relevant or not was "Would you have followed this recommended user in the first place (when selecting which users to follow in the first part of the experiment), if you had know this account?" For example, if a given student was interest in technology and he/she had not discovered the account @TechCrunch during his/her first selection of followees, that would be an interest recommendation because @TechCrunch tweets about news on technology.

5 Results

We first evaluated the performance of the proposed algorithm in terms of their overall precision in followees recommendation. *Precision* can be defined as the number of relevant recommendations over the number of recommendations presented to the user and it can be also computed at different positions in the ranking. For example, P@5 ("precision at five") is defined as the percentage of relevant recommendations among the first five, averaged over all runs. Figure 2 shows the precision achieved by the algorithm, averaged between all users, for each weighting feature at four different positions of the ranking: P@1, P@5, P@10 and P@20. The results of considering each feature separately and the two aggregations functions are showed in this figure.
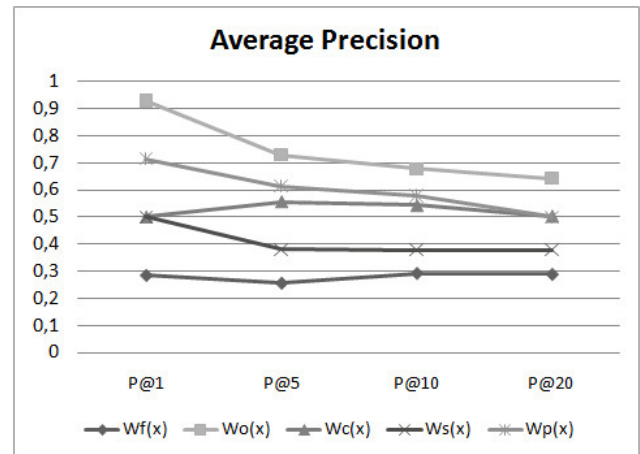

Figure 2: Average precision for each weighting feature

We can observe several interesting facts in the results presented in Figure 2. First, it results that $w_o(x)$, the weighting feature considering the number of occurrences of a user in the list of recommendations as gathered by the algorithm proposed, generates better precision scores than any other weighting feature explored. For this weighting feature we obtained a good recommendation in the first position of the ranking for 93% of the users. For longer ranking lists, preci-

sion decrease from 0.73 for P@5 to 0.64 for P@20, which we believe are all good results.

It is worth noticing that although we reported results up to P@20, recommendations lists tend to be shorter (frequently 5) in order to help the user to focus on the most relevant results. In these small lists the algorithm reached good levels of precision, recommending mostly relevant users.

The weighting feature considering the number of followers, $w_f(x)$, got the worst precision scores, with values under 30%. This fact reveals that this metric, although widely used in other approaches as mentioned in Section 2, is only good at measuring a user's general popularity in the entire Twitter network, but popularity does not necessarily translate into relevance for a particular user. Celebrities and politicians, such as Barack Obama (@barackobama), Lady Gaga (@ladygaga), Yoko Ono (@yokoono), and Tom Cruise (@tomcruise) were a common factor in the rankings of many users regardless their particular interests. Among other popular users suggested that in some cases met the user's interests were popular blogs and news media such as Mundo Geek (@mundo_geek), C5N (@C5N), El Pais, (@el_pais), Mundo Deportivo (@mundodeportivo), Red Hat News (@redhatnews) and Fox Sport LA (@foxsportslat).

A similar situation occurs with $w_c(x)$, the weighting feature considering the number of friends in common between the target user U and the candidate user to recommend to U. Although precision is better than $w_f(x)$ for every size of the recommendation lists, this weighting feature does not reach the performances obtained with $w_o(x)$. This result is expected since the fact that two users U and X share a friend Y does not necessarily means that X is a good information source.

We also found that $w_s(x)$ tends to perform poorly. This score is affected by the term corresponding to the relation between the number of followers and the number of followees, which in most cases is higher than the other terms involved. This factor highly affects the overall average among the three weighting features, causing a decrease in precision.

The second score which combines the three weighting features, $w_p(x)$, seems to overcome this problem since in this case each weighting feature is multiplied to obtain the final score. Nevertheless, celebrities and very popular Twitter user accounts also tend to appear at the top positions of the ranking diminishing the general precision again. However, the factor corresponding to $w_o(x)$ also makes good recommendations to appear interleaved with some popular users on Twitter.

Another interesting issue observed in the results presented in Figure 2 is that for both $w_f(x)$ and $w_c(x)$ precision tend to keep almost constant across different sizes in the list of recommendations and even with a slightly increment as the size of the recommendation set increases. This fact seems to contradict the definition of precision in the information retrieval sense which, by principle, should decrease as the number of recommendations increases. However, this behavior occurs because all $w_f(x)$, $w_s(x)$ and $w_c(x)$ does not

concentrate relevant recommendations in the top positions of the ranking. On contrary, we can observe that for $w_o(x)$ and $w_p(x)$ relevant recommendations tend to be clustered towards the top of the ranking.

Although precision measure gives a general idea of the overall performance of the presented weighting features, it is also very important to consider the position of relevant recommendations in the ranking presented to the user. Since it is known that users focus their attention on items at the top of a list of recommendations [Joachims, 2005], if relevant recommendations appear at the top of the ranking using one algorithm and at the bottom of the ranking using the other, the first algorithm will be perceived as better performing by users even though their general precision might be similar.

Discounted cumulative gain (DCG) is a measure of effectiveness used to evaluate ranked lists of recommendations. DCG measures the usefulness, or gain, of a document based on its position in the result list using a graded relevance scale of documents in a list of recommendations. The gain is accumulated from the top of the result list to the bottom with the gain of each result discounted at lower ranks. The premise of DCG is that highly relevant documents appearing lower in a list should be penalized as the graded relevance value is reduced logarithmically proportional to the position of the result. The DCG accumulated at a particular rank position k is defined as shown in Equation 6:

$$DCG@k = rel_1 + \sum_{i=2}^{k} \frac{rel_i}{\log_2 i} \qquad (6)$$

DCG is often normalized using an ideal DCG vector that has value 1 at all ranks. Figure 3 shows the normalized DCG obtained for both algorithms at four different positions of the ranking: nDCG@1, nDCG@5, nDCG@10 and nDCG@20.
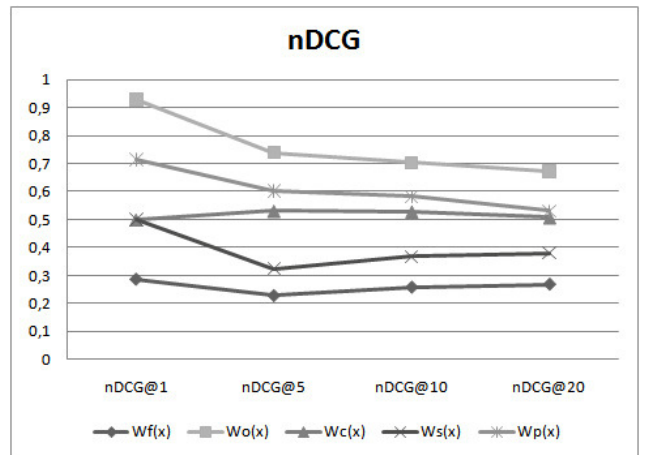

Figure 3: Normalized Discounted Cumulative Gain at rank k for each weighting feature

nDCG@1 is equivalent to P@1 by definition. Then, we can see that scoring users with $w_o(x)$ always positions relevant users above in the ranking than other weighting features, seconded by $w_p(x)$.

Success at rank k (S@k) is another metric commonly used for ranked lists of recommendations. The success at rank k is defined as the probability of finding a good recommendation among the top k recommended users. In other words, S@k is the percentage of runs in which there was at least one relevant user among the first k recommended users. Figure 4 shows the results we obtained for this metric with values of k ranging from 1 to 10.
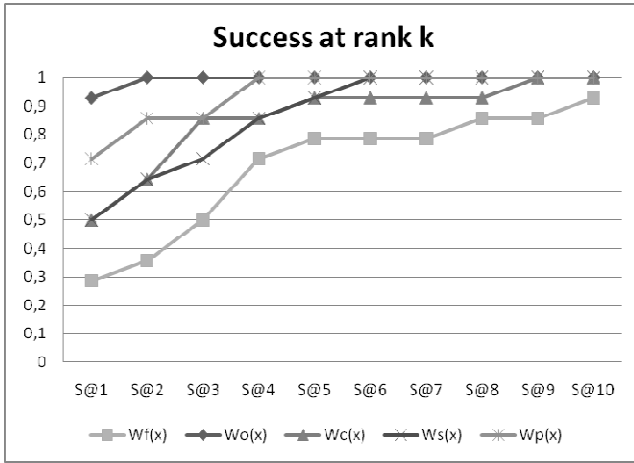

Figure 4: Success at rank k for each weighting feature

For S@k we can observe results equivalent to nDCG@k. Again, scoring users with $w_o(x)$ always positions relevant users above in the ranking than the other weighting features. The ranking according $w_p(x)$ allowed users to find a relevant recommendation always at the most at position 4 in the ranking, while for $w_s(x)$ we obtain success 1 at position 6. With this metric we can confirm that $w_f(x)$ and $w_c(x)$ are not good weighting factors by their own.

To study further the algorithm ability to rank followees for recommendation, we used *Mean Reciprocal Rank* (MRR), a metric that measures where in the ranking is the first relevant recommendation. If the first relevant recommendation is at rank r, then the MRR is 1/r. This measure averaged over all runs provides insight in the ability of the system to recommend a relevant user to follow in Twitter at the top of the ranking. Figure 5 plots the MMR measure for both proposed algorithms.

This metric gives us another view of which weighting feature generates better ranking of recommendations. We confirm that $w_o(x)$ always ranks users better than the other proposed weighting features, while ranking users by their "popularity" does not generate good recommendations.

The experiments presented make us believe that there is reason to be optimistic about the potential for a followee recommender for Twitter using the method described in Section 3.1 to obtain a list of candidates and simple ranking

them by the number of occurrences of each candidate in the list generated by this method. Among the advantages of this method when compared with content-based alternatives is that recommendations can be found quickly based on a simple analysis of the network structure, without considering the content of the tweets posted by the candidate user. Nevertheless, we also believe that combining the proposed method with an analysis of the content of the tweets posted by a user in the list of candidates can improve the precision of a followee recommender system, at the expense of computational performance.
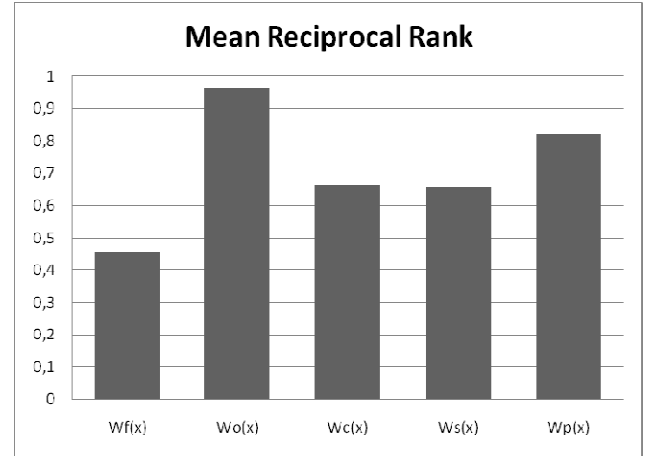

Figure 5: Mean Reciprocal Rank for each weighting feature

6 Comparison with related work

From the related work, the approach that we find more similar to ours (and the only one, up to our knowledge that experimented with real users in a controlled experiment) is *Twittomender*, proposed by Hannon et al. [2010]. Although the results presented in [Hannon *et al.*, 2010] are not fully comparable to the results presented in this article since different datasets were used, in this section we present a comparison about the precision reported for *Twittomender* and the precision obtained with our approach.

Twittomender create different indexes for all users in the dataset generated from different sources of profile information. Four of these indexes are content-based, modeling users by their own tweets, by the tweets of their followers, by the tweets of their followees and by a combination of the three. The three remaining strategies are topology-based and model users by the IDs of their followees, by the IDs of their followees and by a combination of both.

The strategy used for ranking users in the online experiment presented in [Hannon *et al.*, 2010] generates the seven rankings according to the different approaches described above and then generate a single ranking by merging those seven rankings. When merging the rankings they use a scoring function that is based on the position of each user in the recommendation lists. In this way users that are frequently present in high positions are preferred over users that are recommended less frequent or in lower positions.

Hannon et al. performed a live user trial with 34 users, reporting a precision of about 38.2% for k=5 and 33.8% for k=10. Table 1 summarizes the comparison between *Twittomender* and our system. Notice that precision values for *Twittomender* system are approximate because they were taken (and in some cases computed) from the graphics presented in the article.

It is worth noticing that although the number of volunteers who participated in *Twittomender* experiment is more than twice the number of volunteers who participated in our experiment, the number of Twitter users involved in our experiment is by far higher than the number of users in their database. Furthermore, *Twittomender* can only recommend users that are previously indexed. When a user is registered into the system, all his/her followees and followers profiles along with his/her own profile are indexed. Our work differs from this approach in that we do not require indexing profiles from Twitter users; instead a topology-based algorithm explores three levels of the follower/followee network in order to find candidate users to recommend.

7 Discussion and Conclusion

In this article we presented a simple but effective algorithm for recommending followees in the Twitter social network. This algorithm first explores the target user neighborhood in search of candidate recommendations and then sorts these candidates according to different weighting features: the relation between the number of followers and the number of followees, the number of occurrences of each candidate in the final list, the number of friends in common, and two combinations of the three features.

We evaluated the proposed algorithm with real users and we obtained satisfactory results in finding good followee recommendations. We found that considering just the overlapping users among the different lists of follower and followees explored by our crawling method gives better results than the other features considered. As expected, the indegree of a user is not a good feature for ranking followee recommendations. Considering the number of followers for ranking users put celebrities and popular Twitter accounts at the top of the list, but these recommendations are not necessarily interesting for a particular user. However, there are some interesting recommendations discovered by this feature, such as top bloggers who write about a particular subject or news media accounts.

Although the results reported seems promising, we are planning to repeat the experiment this year in order to involve more users in the experiment and obtain more statistical support for the results reported. Moreover, we are very optimistic about the potential improvements that we can obtain by extending the presented approach with content-based techniques. A natural extension of our approach in which we are currently working on is a hybrid algorithm that filters the candidate recommendations found with the topology-based method with a content-based analysis of the tweets posted by the users. In this new approach, a target user U is modeled with a vector of terms built from a content analysis of the tweets posted by U's followees. This vector is then compared with the vector of terms corresponding to each candidate recommendation and the similarity obtained is considered in the generation of the ranking.

The results reported in this article make us feel really enthusiastic about the potentials of Twitter for building recommender systems of sources of information.

	Twittomender "live-user" trial	Our approach
#seeds	34	14
#users	100,000	1,443,111
P@5	~38.2%	72.9%
P@10	~33.8%	67.9%
P@20	~26.9%	64.3%

Table 1: Comparative table between *Twittomender* and our approach

References

[Brin and Page, 1998] S. Brin and L. Page. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*. Volume 30, Issue 1-7, pages 107-117. 1998.

[Cha *et al.*, 2010] M. Cha, H. Haddadi, F. Benevenuto, and K. Gummadi. Measuring user influence in Twitter: The million follower fallacy. *In Proceedings of the 4th International Conference on Weblogs and Social Media (ICWSM'10)*, Washington DC, USA, 2010.

[Chen *et al.*, 2009] J. Chen, W. Geyer, C. Dugan, M. Muller, and I. Guy. Make new friends, but keep the old: recommending people on social networking sites. *In Proceedings of the 27th International Conference on Human Factors in Computing Systems*, pages 201–210, 2009.

[Chen *et al.*, 2010] J. Chen, R. Nairn, L. Nelson, M. Bernstein, and E. Chi. Short and tweet: experiments on recommending content from information streams. *In Proceedings of the 28th International Conference on Human Factors in Computing Systems (CHI'10),* pages 1185–1194, 2010.

[Esparza *et al.*, 2010] S. Garcia Esparza, M. P. O'Mahony, and B. Smyth. On the real-time web as a source of recommendation knowledge. *In Proceedings of the 4th ACM Conference on Recommender Systems (RecSys'10)*, pages 305–308, 2010.

[Garcia and Amatriain, 2010] R. Garcia and X. Amatriain. Weighted content based methods for recommending connections in online social networks. *In Workshop on Recommender Systems and the Social Web*, pages 68–71, Barcelona, Spain, 2010.

[Guy *et al.*, 2009] I. Guy, I. Ronen, and E. Wilcox. Do you know?: recommending people to invite into your social network. *In Proceedings of the 13th International Conference on Intelligent User Interfaces (IUI'09)*, pages 77–86, 2009.

[Hannon *et al.*, 2010] J. Hannon, M. Bennett, and B. Smyth. Recommending Twitter users to follow using content and collaborative filtering approaches. *In Proceedings of the 4th ACM Conference on Recommender Systems (RecSys'10),* pages 199–206, 2010.

[Java *et al*., 2007] A. Java, X. Song, T. Finin, and B. Tseng. Why we twitter: understanding micrblogging usage and communities. In *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis (WebKDD/SNA-KDD '07)*. ACM, New York, NY, USA, pages 56-65. 2007.

[Joachims *et al*., 2005] T. Joachims, L. Granka, B. Pan, H. Hembrooke, and G. Gay. Accurately interpreting click-through data as implicit feedback. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval (SIGIR '05)*. ACM, New York, NY, USA, 154-161. 2005.

[Krishnamurthy *et al*., 2008] B. Krishnamurthy, P. Gill, and M. Arlitt. A few chirps about twitter. *In Proceedings of the first workshop on Online social networks (WOSP '08)*. ACM, New York, NY, USA, pages 19-24. 2008.

[Kwak *et al.*, 2010] H. Kwak, C. Lee, H. Park, and S. Moon. What is Twitter, a social network or a news media? *In Proceedings of the 19th International Conference on World Wide Web (WWW'10)*, pages 591–600, 2010.

[Lo and Lin, 2006] S. Lo and C. Lin. WMR–A graph-based algorithm for friend recommendation. *In Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence (WI'06)*, pages 121–128, Washington, DC, USA, 2006.

[Phelan *et al.*, 2009] O. Phelan, K. McCarthy, and B. Smyth. Using Twitter to recommend real-time topical news. *In Proceedings of the 3rd ACM Conference on Recommender Systems (RecSys'09),* pages 385–388, 2009.

[Sun *et al.*, 2009] A. R. Sun, J. Cheng, and D. D. Zeng. A novel recommendation framework for micro-blogging based on information diffusion. *In Proceedings of the 19th Workshop on Information Technologies and Systems,* 2009.

[Weng *et al.*, 2010] J. Weng, E-P. Lim, J. Jiang, and Q. He. TwitterRank: finding topic-sensitive influential twitterers. *In Proceedings of the 3rd ACM International Conference on Web Search and Data Mining (WSDM'10),* pages 261–270, New York, NY, USA, 2010.

[Yamaguchi *et al.*, 2010] Y. Yamaguchi, T. Takahashi, T. Amagasa, and H. Kitagawa. TURank: Twitter user ranking based on user-tweet graph analysis. *In Web Information Systems Engineering*, volume 6488 of LNCS, pages 240–253, Hong Kong, China, 2010.