A Sentiment-Based Approach to Twitter User Recommendation

Davide Feltoni Gurini, Fabio Gasparetti,
Alessandro Micarelli and Giuseppe Sansonetti
Roma Tre University
Department of Engineering
Via della Vasca Navale 79
Rome, 00146 Italy
{feltoni,gaspare,micarel,gsansone}@dia.uniroma3.it

ABSTRACT

Nowadays, the emerging popularity of Social Web raises new application areas for recommender systems. The aim of a social user recommendation is to suggest new friends having similar interests. In order to identify such interests, current recommender algorithms exploit social network information or the similarity of user-generated content. The rationale of this work is that users may share similar interests but have different opinions on them. As a result, considering the contribution of user sentiments, can yield benefits in recommending possible friends to follow.

In this paper we propose a user recommendation technique based on a novel weighting function, we named *sentiment-volume-objectivity (SVO)* function, which takes into account not only user interests, but also his sentiments. Such function allows us to build richer user profiles to employ in the recommendation process than other content-based approaches. Preliminary results based on a comparative analysis show the benefits of the advanced approach in comparison with some state-of-the-art user recommender systems.

Categories and Subject Descriptors

H.3.3 [**Information Search and Retrieval**]: [Information Filtering]

General Terms

Algorithms,Experimentation

Keywords

User Recommendation, Twitter, Sentiment Analysis

1. INTRODUCTION

The growing popularity of social networks increases the availability of user sentiments, which has become a significant impact factor on buying decisions, brand reputations and public opinions. Furthermore, recommending pertinent news stories, documents, and users to follow, has long been a favourite domain for recommender systems research. Several new approaches harness real-time micro-blogging activity from services, such as Twitter[1], as the basis for identifying user preferences and filtering relevant contents to specific people. Recently, Twitter has become an interesting source of research activity as a result of the large amount of available user-generated data. In particular Twitter permits users to share a sentence - called tweet - to the followers, with a maximum length of 140 characters.

In this instance, the purpose of user recommendation is to identify relevant people to follow among millions of users that interact in the social network. Previous attempts include both content-based and graph-based approaches. The former focuses on metrics for measuring the topic similarity among Twitter users, the latter exploits the graph of relationships among users to infer correlations.

The main idea behind this work is that users may share similar interests but have different opinions about them. Therefore, we extend the content-based recommendation by means of the sentiments and opinions extracted from the user micro-posts in order to improve the accuracy of the suggestions. This leads us to define a novel weighting function in order to enrich content-based user profiles.

2. RELATED WORK

In spite of the growing body of research on exploiting user-generated contents in recommendation engines, there are few attempts to consider sentiment included in micro-posts during the recommendation process. Singh *et al.* [17] introduce a hybrid recommender system that improves the results of collaborative filtering by incorporating a sentiment classifier in the movie recommendation scenario. Bank and Franke [4] try to better represent public product reviews on weblogs through different text mining techniques. Faridani [9] achieves the same goal by exploiting a multivariate regression approach. As far as we are aware, there are no attempts towards sentiment user recommendation in social networks.

User recommendation approaches that ignore user opinions have been proposed by Freyne *et al.* [10] and Chen *et al.* [8] exploring different recommendations strategies. Ap-

[1] `twitter.com`

proaches for social recommendation that incorporate user opinions have been proposed in other domains, e.g., [5]. Guy *et al.* [12] propose a people recommendation engine within an enterprise social network site scenario. They aggregate several different sources to derive factors that might influence the similarity measure. Twittomender [13] lets users find pertinent profiles on Twitter exploiting different strategies, both content-based and collaborative ones. Arru *et al.* [3] propose a signal-based representation of user interests in order to draw similarities among people.

3. SENTIMENT ANALYSIS ALGORITHM

Sentiment analysis or opinion mining is formally defined as the computational study of sentiments and opinions about an entity expressed in a text. According to Liu [15], the entity is classified into five categories: *product, person, brand, event, concept*. Particularly, in this work we assume the *concept* as the sentiment analysis target entity. Sentiment analysis is a difficult task, hence - before the setup of the algorithm - some assumptions are needed. There are multiple granularity levels of sentiment analysis, as explained in [2]: feature-level, entity-level, sentence-level, document-level.

In this work we consider sentiment analysis at sentence-level. Specifically, in the Twitter domain we assume that a sentence matches the whole tweet. Moreover, we assume that each sentence contains only one opinion related to the entity.

The goal of our sentiment analysis system is to obtain an output value that represents how much positive, negative or neutral is the sentiment expressed in a tweet. For this reason, we implemented a Supervised Machine Learning algorithm based on a Naïve Bayes classifier. With a view to training our algorithm, we needed a dataset with labeled tweets. However, due to the lack of a Twitter public dataset, we decided to follow an alternative approach. Instead of manually building a labeled dataset, Bhayani *et al.* [11] propose to employ a noisy dataset of positive, negative, and neutral tweets. The labels correspond to special sequences of characters in the tweets, such as positive or negative emoticons (e.g., `:-D ;-(`), hashtags (e.g., `#iloveit`, `#ihate`) or keywords (e.g., `good`, `sad`). Even though these labels do not always correspond to the right sentiment expressed by the tweet, they allow us to collect a large amount of data for training. The Twitter API[2] have been used to retrieve a set of tweets containing the aforementioned features. The final training dataset counts 150000 tweets divided in 50000 tweets for each class. Because the experimental evaluation is conducted on events related to the 2013 Italian political elections, the TextCat language recognizer[3] is employed to limit the set to Italian tweets. In order to increase the classifier precision and reduce the presence of noise, we performed a feature selection. In particular, the terms with low values of *Salience* are discarded. The *Salience* of a term t is defined by Pak et al. [16] as follows:

$$Salience(t) = \frac{1}{N}\sum_{i=1}^{N-1}\sum_{j=i+1}^{N} 1 - \frac{min(P(t \in L_i), P(t \in L_j))}{max(P(t \in L_i), P(t \in L_j))} \tag{1}$$

where N is the number of the dataset labels, namely, $N = 3$ (i.e., positive, negative, and neutral) and $P(t \in L)$ is the likelihood that the term t belongs to the label class L. A zero value of *Salience* means that the term t appears uniformly in each dataset, thus it is a good candidate to be discarded. Finally, as for the Machine Learning algorithm, a Naïve Bayes classifier is trained on the training data, where each tweet is represented as a feature vector made up of the following groups of features:

- Bag-of-words: vectors of word unigram;
- Word polarities: using the LIWC[4] content analysis dictionary, we extracted features for positive, negative, and neutral words. Individual word polarities are inverted if the word follows a negation;
- Negations: we add the "NEG_" suffix to each word following a negation pattern (e.g., "not perfect" becomes "perfect_NEG");
- Elongated words: we represent as a feature the presence of words with one character repeated more than two times, (e.g., "looove", "yesss");
- Part-of-speech tags: they provide a rough measure of the tweet content.

4. SVO RECOMMENDATION APPROACH

4.1 User profiling

Several approaches to user recommendation are based on the definition of a similarity measure between two users u_i and u_j. Given the user u_i, the ranked list of suggested users corresponds to the set of users u_j that maximize the aforementioned measure. Content-based approaches define this measure by analyzing the user tweets. The set T of tweets $tweets(u)$ posted by the user u can be viewed as an extension of the *bag-of-word* model, where concepts are more semantically significant and less ambiguous than plain keywords. Instead of using complex semantic annotators, a concept is uniquely identified through *hashtags* contained in the tweet, namely, the metadata tags that are used in Twitter to indicate the context or the flow a tweet is associated with. Thus, we define the profile p of the user u as the set of weighted concepts:

$$p(u) = \{(c, \omega(u, c)) | c \in C_u\} \tag{2}$$

where $\omega(u, c)$ is the relevance of the concept c for the user u, and C_u is the set of concepts cited by the user u. The weighting function will be discussed in the following section.

The user profile representation is generated by monitoring the user activity, that is, all the tweets included in the observation period. Afterwards, given two users u_i and u_j, and their profiles $p(u_i)$ and $p(u_j)$, the similarity function is defined in terms of cosine similarity:

$$\begin{aligned} sim(u_i, u_j) &= sim(p(u_i), p(u_j)) = \\ &= \frac{\sum_{c \in C_{u_i} \cup C_{u_j}} \omega(u_i, c) \cdot \omega(u_j, c)}{\sqrt{\sum_{c \in C_{u_i}} \omega(u_i, c)^2} \cdot \sqrt{\sum_{c \in C_{u_j}} \omega(u_j, c)^2}} \end{aligned} \tag{3}$$

where C_{u_i} and C_{u_j} are the concepts in the profiles of users u_i and u_j, respectively.

[2]`dev.twitter.com`

[3]`www.let.rug.nl/vannoord/TextCat/`

[4]`liwc.net`

4.2 SVO Weighting Function

The idea behind this work is that taking into account user attitudes towards his own interests can yield benefits in recommending friends to follow. Specifically, we consider (i) which is the sentiment expressed by the user for a given concept, (ii) how much he is interested in that concept, and (iii) how much he expresses objective comments on it.

In our model the first contribution $S(u,c)$, namely, the *sentiment* of the user u about a concept c, is obtained as follows:

$$S(u,c) = f\left(\frac{Pos(u,c) - Neg(u,c)}{Pos(u,c) + Neg(u,c)}\right) \tag{4}$$

where $Pos(u,c)$ and $Neg(u,c)$ are the sums of the positive and negative tweets written by the user u regarding the concept c, respectively. Such values are calculated by means of our proposed Machine Learning algorithm (see Section 3) that classifies the tweets as positive, negative or neutral. A low value of $S(u,c)$ means that the user sentiments towards the concept c are negative, on the contrary a high value represents positive sentiments.

The f function is used to normalize the output value within the $[0,1]$ range:

$$f(x) = \frac{1}{1+k^{-x}} \tag{5}$$

where k = 10.

The second contribution is the *volume* $V(u,c)$, that is, how much a user u wrote about a specific concept c and is defined as follows:

$$V(u,c) = \frac{tweets(u,c)}{\sum_{i=1}^{N} tweets(u,c_i)} \tag{6}$$

where $tweets(u,c)$ is the number of tweets written by the user u about a specific concept c, and N is the total number of concepts dealt with by u.

The third contribution is the *objectivity* $O(u,c)$. With this term we denote how many tweets about a concept c do not contain sentiments or opinions and therefore may be objective. This may be important because objective tweets are typically news, so quite significant for the similarity of user profiles but less relevant for the sentiment analysis.

$O(u,c)$ is defined as follows:

$$O(u,c) = \frac{Neutral(u,c)}{Pos(u,c) + Neg(u,c) + Neutral(u,c)} \tag{7}$$

where $Pos(u,c)$, $Neg(u,c)$ and $Neutral(u,c)$ are the sums of the positive, negative, neutral tweets written by the user u relative to the concept c, respectively.

Based on such contributions, we proposed a novel weighting function, we called *sentiment-volume-objectivity (SVO)* function, that takes into account all of them. It is defined as follows:

$$SVO(u,c) = \alpha S(u,c) + \beta V(u,c) + \gamma O(u,c) \tag{8}$$

where α, β, and γ are three constants $\in [0,1]$, such that $\alpha + \beta + \gamma = 1$. The function $SVO(u,c) \in [0,1]$ is the weighting function $\omega(u,c)$ that appears in the Equations 2 and 3.

The experimental evaluations (Section 5) shows the computation of the values of the parameters α, β, and γ that maximize the performance of the recommender.

5. EXPERIMENTAL EVALUATION

5.1 Dataset

In order to evaluate the proposed model, we considered a case study rich of sentiments, such as the 2013 Italian political elections. Using the Twitter APIs we selected 31 hashtags for retrieving the Twitter streams about politician leaders and parties from Jan 25th to Feb 27th. Furthermore, because social networks are dynamic and fast-changing, we retrieved the hashtags that more often co-occur in the obtained tweets and added them to the initial hashtag set. This way, we took into account the trending topics that may be ignored in the initial query setup. The dataset counted 1085000 tweets, and over 25000 users that wrote almost one tweet. For the experimental evaluation we finally selected 1000 random users that (i) posted at least 50 tweets in the observed period, and (ii) had more than 15 friends and followers already stored into the dataset. The final dataset for the evaluation counted 805956 tweets.

5.2 Evaluation

The goal of our user recommender system is to suggest to a user someone to follow, with similar interests and opinions. In order to compare different profiling approaches and recommendation strategies, we need to understand when a user u_1 is relevant for a user u_2. In this work we suppose that u_1 is relevant for u_2 if a *following relationship* exists between them. This assumption has recently became a commonplace among social networks recommender systems [1, 14, 3] and is supported by the phenomenon of *homophily*, that is, the tendency of individuals with similar characteristics to associate with each other.

We performed a preliminary evaluation in order to assess the effectiveness of the proposed approach. For the sake of brevity, in this paper we only report the results of a comparative analysis of our approach with two traditional approaches that do not consider sentiment: (i) cosine similarity in a Vector Space Model (VSM) where vectors are weighted hashtags, and (ii) the function $S1$ proposed by Hannon*et al.* [13]. We used different metrics to express the evaluation results. *Success at Rank K* (S@K) provides the mean probability that a relevant user is located in the top K positions of the list of suggested users. *Mean Reciprocal Rank* (MRR) indicates the average position of a user in the recommended list. *Mean Average Precision* at cut-off K (MAP@K) is the average of the precision value for each of the top-K recommended users. Figure 1 shows the obtained evaluation results. As can be seen, our approach outperforms the other ones according to each evaluation metric. These findings confirm that sentiment is a valuable feature to be considered in order to improve the user recommender systems. As a marginal note, the absolute values of the achieved results are high due to the characteristics of the built dataset, where the relations among users are significantly dense. Finally, we also analyzed the user recommender performance in terms of variations of the three parameters α, β, and γ (see equation 8). In order to determine the best values of those parameters, we implemented a *mini-batch gradient descent* algorithm. The best results, according to aforementioned metrics, was achieved running the evaluation with $\alpha = 0.3$, $\beta = 0.6$, and $\gamma = 0.1$. Based on the proposed model and the used dataset, these weights appear to highlight the contribution of the *volume* and the *sentiment* in comparison with the *objectivity*.

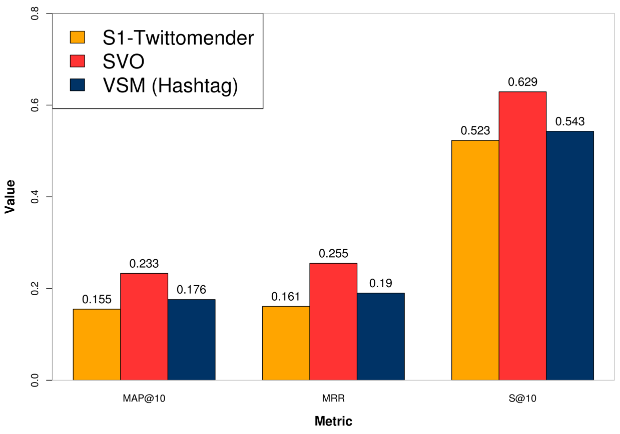

Figure 1: Comparative analysis among the proposed approach and two other state-of-the-art methods.

6. CONCLUSIONS

In this paper we have described a user recommender system for Twitter. Our work emphasizes the use of implicit sentiment analysis in order to improve the performance of the recommendation process. We have defined a novel weighting function that takes into account sentiment, volume, and objectivity related to the user interests. This technique allowed us to build more complete user profiles than traditional content-based approaches. Preliminary results show the benefits of our proposed model compared with some state-of-the-art methods.

As future work we are planning a deep sensitivity analysis to investigate whether social interactions, user preference and dataset characteristics shape parameters α, β, and γ. We will also include some improvements of the recommendation process taking into account other elements (e.g., named-entities, persons, products) and semantic representations of hashtags (e.g., [6][7]). A future study will also focus on the use of the implicit sentiment analysis within the collaborative filtering in social networks.

7. REFERENCES

[1] F. Abel, Q. Gao, G.-J. Houben, and K. Tao. Analyzing temporal dynamics in twitter profiles for personalized recommendations in the social web. In *Proceedings of ACM WebSci '11, 3rd International Conference on Web Science, Koblenz, Germany*. ACM, June 2011.

[2] A. Agarwal, B. Xie, I. Vovsha, O. Rambow, and R. Passonneau. Sentiment analysis of twitter data. In *Proceedings of the Workshop on Languages in Social Media*, LSM '11, pages 30–38, Stroudsburg, PA, USA, 2011. Association for Computational Linguistics.

[3] G. Arru, D. Feltoni Gurini, F. Gasparetti, A. Micarelli, and G. Sansonetti. Signal-based user recommendation on twitter. *Social Recommender Systems 2013*, 2013.

[4] M. Bank and J. Franke. Social networks as data source for recommendation systems. In F. Buccafurri and G. Semeraro, editors, *E-Commerce and Web Technologies*, volume 61 of *Lecture Notes in Business Information Processing*, pages 49–60. Springer Berlin Heidelberg, 2010.

[5] C. Biancalana, F. Gasparetti, A. Micarelli, and G. Sansonetti. An approach to social recommendation for context-aware mobile services. *ACM Trans. Intell. Syst. Technol.*, 4(1):10:1–10:31, Feb. 2013.

[6] C. Biancalana, F. Gasparetti, A. Micarelli, and G. Sansonetti. Social semantic query expansion. *ACM Trans. Intell. Syst. Technol.*, 4(4), 2013. Forthcoming issue.

[7] C. Biancalana and A. Micarelli. Social tagging in query expansion: A new way for personalized web search. In *CSE (4)*, pages 1060–1065. IEEE Computer Society, 2009.

[8] 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*, CHI '09, pages 201–210, New York, NY, USA, 2009. ACM.

[9] S. Faridani. Using canonical correlation analysis for generalized sentiment analysis, product recommendation and search. In *Proceedings of the fifth ACM Conference on Recommender systems*, RecSys '11, pages 355–358, New York, NY, USA, 2011. ACM.

[10] J. Freyne, M. Jacovi, I. Guy, and W. Geyer. Increasing engagement through early recommender intervention. In *Proceedings of the third ACM Conference on Recommender Systems*, RecSys '09, pages 85–92, New York, NY, USA, 2009. ACM.

[11] A. Go, R. Bhayani, and L. Huang. Twitter sentiment classification using distant supervision. *Processing*, pages 1–6, 2009.

[12] I. Guy, I. Ronen, and E. Wilcox. Do you know?: recommending people to invite into your social network. In *Proceedings of the 14th International Conference on Intelligent User Interfaces*, IUI '09, pages 77–86, New York, NY, USA, 2009. ACM.

[13] J. Hannon, K. McCarthy, and B. Smyth. Finding useful users on twitter: twittomender the followee recommender. In *Proceedings of the 33rd European Conference on Advances in Information Retrieval*, ECIR'11, pages 784–787, Berlin, Heidelberg, 2011. Springer-Verlag.

[14] H. John, B. Mike, and S. Barry. Recommending twitter users to follow using content and collaborative filtering approaches. *RecSys'10 : Proceedings of the 4th ACM Conference on Recommender Systems,*, 26-30(10):8, 09 2010.

[15] B. Liu. Sentiment analysis and subjectivity. *Handbook of Natural Language Processing,*, pages 627–666, 2010.

[16] A. Pak and P. Paroubek. Twitter as a corpus for sentiment analysis and opinion mining. volume 2010, pages 1320–1326, 2010.

[17] V. K. Singh, M. Mukherjee, and G. K. Mehta. Combining collaborative filtering and sentiment classification for improved movie recommendations. In *Proceedings of the 5th International Conference on Multi-Disciplinary Trends in Artificial Intelligence*, MIWAI'11, pages 38–50, Berlin, Heidelberg, 2011. Springer-Verlag.