Interweaving Trend and User Modeling for Personalized News Recommendation

Qi Gao, Fabian Abel, Geert-Jan Houben, Ke Tao
Web Information Systems Group
Delft University of Technology
Delft, the Netherlands
Email:{q.gao, f.abel, g.j.p.m.houben, k.tao}@tudelft.nl

***Abstract*—In this paper, we study user modeling on Twitter and investigate the interplay between personal interests and public trends. To generate semantically meaningful profiles, we present a framework that allows us to enrich the semantics of individual Twitter messages and features user modeling as well as trend modeling strategies. These profiles can be re-used in other applications for (trend-aware) personalization. Given a large Twitter dataset, we analyze the characteristics of user and trend profiles and evaluate the quality of the profiles in the context of a personalized news recommendation system and show that personal interests are more important for the recommendation process than public trends and that by combining both types of profiles we can further improve recommendation quality.**

***Keywords*-twitter; user modeling; trend modeling; personalized news recommendation; social web**

I. INTRODUCTION

Today, people publish millions of messages per day on Twitter, which is the most popular microblogging service on the Web. Recent research shows that the majority of Twitter messages (tweets) are related to news [1] and that the trending topics propagate quickly through the Twitter social network which allows for applications such as early warning systems [2]. So far, most research initiatives focused on the analysis of structural properties of the Twitter network. In this paper, we investigate how individual Twitter activities can be exploited to infer personal interests and generate semantic user profiles that can be re-used also by other applications than Twitter.

In contrast, to other Social Web services like Last.fm, which allows for the deduction of users' musical taste, or Flickr, which primarily provides information to infer users' interests in locations or events, tweets are not restricted to a certain domain. Instead, Twitter users can discuss about any topic they are interested in or concerned with which makes it worthwhile to explore for user modeling. Furthermore, the real-time nature of information people publish on Twitter poses new challenges and possibilities for user modeling. In this paper, we research whether information from Twitter can be exploited to generate profiles that reflect the interests of a user in current news topics. We present a Twitter-based trend and user modeling framework and evaluate different strategies for generating trend-aware user profiles in the context of personalized news recommendations.

Generating semantically meaningful profiles based on tweets is a non-trivial problem as tweets are limited to 140 characters. Rowe et al. propose the use of contextual information to enrich the semantics of tweets [3]. In previous work, we introduced strategies that link tweets with external Web resources and proposed methods for enriching the semantics of tweets with meaningful concepts extracted from the linked Web resources [4]. In this paper, we exploit these enrichment strategies for user and trend modeling on Twitter. We analyze the temporal characteristics of user and trend profiles and evaluate trend and user modeling strategies for recommending news articles to users.

II. USER AND TREND MODELING ON TWITTER

To obtain semantically meaningful concepts for representing the users' interests, we propose a Twitter-based User Modeling Framework that allows for the generation of different types of semantic user profiles. The generic model of user profiles generated by our user modeling framework is specified in Definition 1.

Definition 1 (User Profile): The profile of a user $u \in U$ is a set of weighted concepts where with respect to the given user u for a concept $c \in C$ its weight $w(u, c)$ is computed by a certain weighting function w.

$$P(u) = \{(c, w(u, c)) | c \in C, u \in U\} \tag{1}$$

Here, C and U denote the set of concepts and users respectively. With $\vec{p}(u)$ we refer to $P(u)$ in its vector space model representation, where the value of the i-th element refers to $w(u, c_i)$ and the values of all elements are normalized to make the sum of values to 1. In previous work, we introduce three types of profiles that differ in the type of concepts C that occur in a profile: hashtag-based, topic-based, and entity-based profiles and various weighting functions [4].

While user profiles represent personal interests of a specific user, trend profiles describe the trending interests of the entire user community. In line with the three types of user profiles, trends may also refer to hashtags, topics or entities. The generic model of profiles that represent public trends for a given time interval can be defined analogously to the user profile model (cf. Definition 1):

Definition 2 (Trend Profile): The *trend profile* $T(I_j)$ for a given time interval I_j is a set of weighted concepts where for a concept $c \in C$ its weight $w(I_j, c)$ is computed by a certain function w.

$$T(I_j) = \{(c, w(I_j, c)) | c \in C\} \quad (2)$$

Here, C denotes the set of candidate concepts from which the trends can be extracted in the given time interval I_j. With $\vec{t}(I_j)$ we refer to $T(I_j)$ in its vector space model representation, where the value of the i-th element refers to $w(I_j, c_i)$ and the values of all elements are normalized to make the sum of values to 1.

The weighting function $w(I_j, c)$ is applied to measure the importance and popularity of a concept in a specific period of time. In particular, we make use of term frequency and inverse document frequency and also introduce time-sensitive variations of these measures.

TF: For a given time interval I_j, the term frequency TF of a concept c is the fraction of concept references that refer to c.

$$w_{TF}(I_j, c) = \frac{n_{c,j}}{\sum_{c \in C} n_{c,j}} \quad (3)$$

where $n_{c,j}$ denotes the number of (enriched) tweets that refer to concept c during time interval I_j.

TF×IDF: The inverse document frequency (IDF) can be applied to value the specificity of a concept c within a given period of time I_j.

$$w_{TF \times IDF}(I_j, c) = w_{TF}(I_j, c) \cdot log(\frac{|I|}{1 + |\{I_i : n_{c,i} > 0\}|}) \quad (4)$$

where $w_{TF}(I_j, c)$ is the term frequency of concept c in the time interval I_j, $|I|$ denotes the number of separated time intervals and $|\{I_i : n_{c,i} > 0\}|$ is the number of time intervals in which the concept c was referenced at least once.

Our analysis of the Twitter activities (see Section III) reveals that there are concepts (entities) that show constant popularity and are not specific for certain time periods. In contrast, there exist other new occurring entities quickly gain popularity in a short period of time and fade out after some time. To measure the temporal dynamics of a concept c on a more continuous spectrum, we calculate the standard deviation of the timestamps of (semantically enriched) tweets that refer to c.

$$\sigma(c) = \sqrt{\frac{\sum_{k=1}^{N}(ts_k - \overline{ts})^2}{N-1}} \quad (5)$$

Here, ts_k is the timestamp of the k-th tweet that refers to concept c, $\overline{ts}$ is the average timestamp of tweets that relate to c and N is the overall number of tweets that refer to c. Similar to an interpretation that characterize the temporal statbility of hashtags in [5], here the lower the value of $\sigma(c)$ the shorter is the time period in which concept c is referenced by tweets. If a concept c is just mentioned once then the standard deviation will be zero ($\sigma(c) = 0$). In contrast, the higher the standard deviation, the more constantly is a concept referenced from tweets. Given this notion of standard deviation, we modify the conventional TF and $TF \times IDF$ weighting functions and introduce two new time-sensitive weighting schemes.

Time-sensitive TF: For a given time interval I_j, the time-sensitive term frequency TF of a concept c is defined as follows:

$$w_{t-TF}(I_j, c) = w_{TF}(I_j, c) \cdot (1 - \hat{\sigma}(c)) \quad (6)$$

where $w_{TF}(I_j, c)$ is the term frequency of concept c for time interval I_j and the standard deviation $\sigma(c)$ is normalized by the maximum value, denoted as $\hat{\sigma}(c)$.

Time-sensitive TF×IDF: Similarly, the time-sensitive $TF \times IDF$ is specified as follows:

$$w_{t-TF \times IDF}(I_j, c) = w_{TF \times IDF}(I_j, c) \cdot (1 - \hat{\sigma}(c)) \quad (7)$$

where $w_{TF \times IDF}(I_j, c)$ denotes the weight fro conventional $TF \times IDF$ function and $\hat{\sigma}(c)$ denotes normalized standard deviation value for concept c.

Hence, $\hat{\sigma}(c) \in [0..1]$ and in particular the factor $(1 - \hat{\sigma}(c))$ is used to de-emphasize TF and $TF \times IDF$. The higher the normalized standard deviation $\hat{\sigma}(c)$, the lower the weight $w_{t-TF}(I_j, c)$ and $w_{t-TF \times IDF}(I_j, c)$ respectively. We use $T_{w@k}(I_j)$ to denote the trend profile that selects the top k weighted concepts $c \in C$ when applying a certain weighting function w.

Furthermore, to generate a profile that better estimates user interests in the context of public trends for supporting trend-aware personalization, we apply a classical mixture method to combine a given user profile with a trend profile of a certain time period I_j.

$$\vec{m}(I_j, u) = d * \vec{p}(u) + (1 - d) * \vec{t}(I_j) \quad (8)$$

III. Temporal Analysis of User and Trend Profiles

In this section we analyze the temporal dynamics of user and trend profiles based on a large Twitter dataset that we collected in [4].

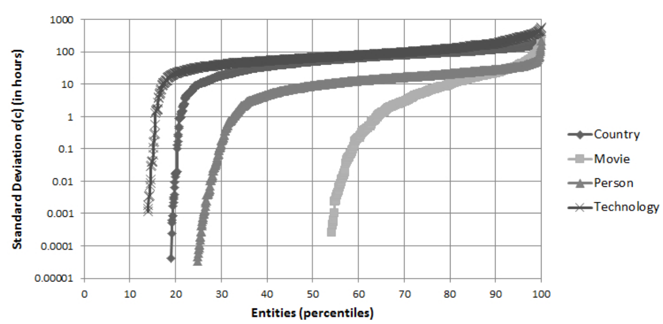

Figure 1. Standard deviation of different types of entities (cf. Equation 5).

First, we observe that different types of entities have a different life span within entity-based profiles. Figure 1 shows the standard deviation (see Equation 5) of different types of entities. The higher the standard deviation of a certain entity is, the more consistently does the entity occur in Twitter messages posted by the users. For example, we see that entities of type country or technology have, on average, a higher standard deviation than movies or persons, i.e.

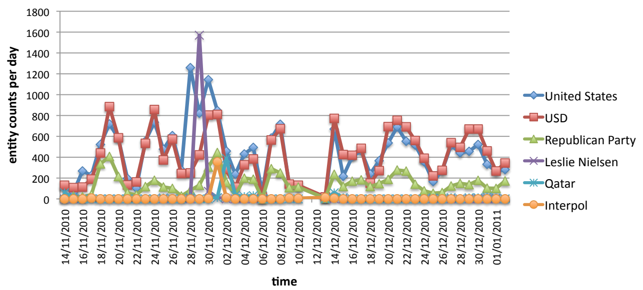

Figure 2. Top entities for a given week.

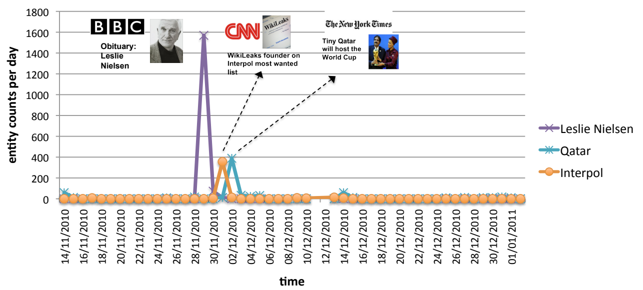

Figure 3. Top trends for a given week.

country or technology entities occur more constantly within profiles than movie and person entity. In fact, for more than 50% of the movies standard deviation is 0, which means that the corresponding entities are mentioned just once by a user. In summary, we observe that some fractions of the entity-based profiles are rather constant while others change dynamically over time.

Second, we investigate the temporal characteristics of trends in Twitter and analyze the effectiveness of trend modeling strategies. In Figure 2 we plot the occurrence frequency of popular entities over time. Some entities such as "United States" are continuously among the most frequently mentioned entities. We assume that interest in those entities will anyhow be captured by the long-term user profiles. Hence, when generating trend profiles, we are rather interested in trending entities that have a peak at a certain point of time. For example, Figure 2 shows a peak for "Leslie Nielson" who was usually not mentioned frequently but suddenly became the most popular entity at the end of November 2010. While this trend is clearly visible, other trends are overloaded by entities that are constantly popular. Figure 3 shows the entities from Figure 2 which are identified as trending entities according to the time-sensitive $TF \times IDF$ weighting function. By using the time-sensitive weighting function in particular we are able to filter out those popular entities and can identify trending entities which are of particular importance for a specific period in time. For example, We observe that there are further peaks for "Qatar" and "Interpol" at the beginning of December 2010 which are overloaded by "United States" and "USD" in Figure 2.

IV. Evaluation of Trend and User Modeling for Recommending News Articles

In this section, we apply the Twitter-based trend and user modeling framework proposed in Section II to provide personalized news recommendations. Our goal is to compare the quality achieved by the same recommendation algorithm when inputting different types of profiles. Therefore, we apply a lightweight content-based algorithm that recommends news according to their cosine similarity with a given profile generated by the trend and user modeling framework.

We considered the last week of our observation period in our dataset as the time interval for computing recommendations. The ground truth of news articles, which we consider as *relevant* for a specific user u, is obtained via the Twitter messages (including re-tweets) posted by u in this week that explicitly link to a news article published by BBC, CNN or New York Times. We run our experiments for these 577 users, for whom we identified at least five relevant news articles during our recommendation period. For each of these users, we compare the three alternative modeling strategies for generating an input profile to be fed into the news recommendation algorithm: the user profile $P(u)$, the trend profile $T(I_r)$, and the combined profile $M(I_r, u)$. The quality of the recommendations is measured by means of *MRR* (Mean Reciprocal Rank), which indicates at which rank the first item relevant to the user occurs on average, and *S@k* (Success at rank k), which stands for the mean probability that a relevant item occurs within the top k of the ranking.

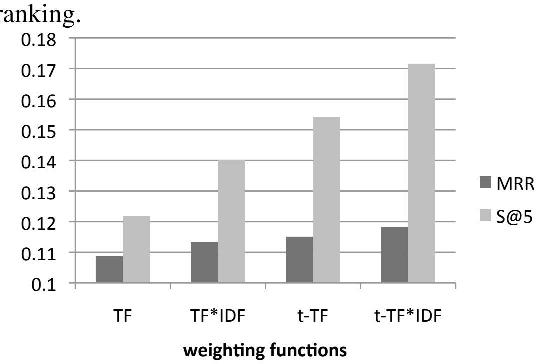

Figure 4. Comparison of different weighting functions (for trend profile $T_{w@500}(I_r)$)

The results of the news recommendation experiments are demonstrated in Figure 4, 5 and 6. First, we investigate which weighting function is best for generating trend profiles in the context of news recommendations. Our hypothesis is that our time-sensitive methods that adjust conventional TF and $TF{\times}IDF$ weighting functions by means of standard deviation allow us to better emphasize the emerging and popular concepts in a specific period of time. To validate our hypothesis, we generate trend profiles by applying different weighting functions and compare *MRR* and *S@5* measures for the recommendations. Figure 4 reveals that the time-sensitive weighting functions improve the quality of news recommendations clearly. With the time-sensitive $TF{\times}IDF$ weighting function we reach the best recommendation performance and improve over the TF baseline by 9.3% and 40.1% with respect to MRR and S@5 respectively. These

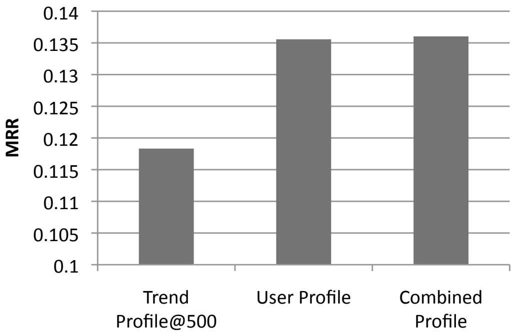

Figure 5. Comparison of different type of profiles ($T_{t-TF\times IDF@500}(I_r)$, $P(u)$ and $M(I_r, u)$)

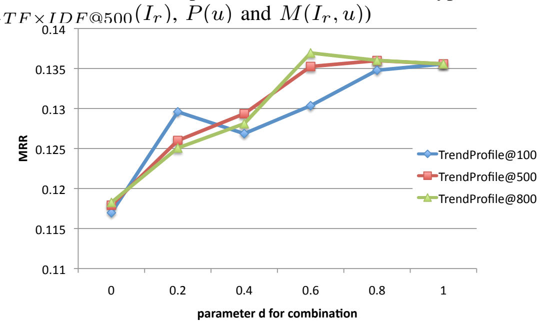

Figure 6. Comparison of different strategies for combining user and trend profiles

results thus confirm our hypothesis and show that time-sensitive weighting functions and time-sensitive $TF\times IDF$ in particular are best for generating trend profiles in the context of recommending news articles.

Second, we compare the performance of the three types of profiles that are generated by our trend and user modeling framework: user profiles, trend profiles and the mixture of both of these profiles. By interweaving the personal interests into global trends that we detect in the Twitter community, we expect to better estimate the current and future interests for supporting trend-aware personalization and news article recommendations in particular. Figure 5 compares the performance of the strategies with respect to MRR when using $TF\times IDF$ as weighting function for modeling the top 500 trends ($T_{t-TF\times IDF@500}(I_r)$). We observe that user profiles allow for better recommendation quality. In fact, personal user interest profiles $P(u)$ improve the performance of public trend profiles $T_{t-TF\times IDF@500}(I_r)$ by 15.3% regarding mean reciprocal rank of the first relevant item (MRR). Furthermore, by combining personal interests with public trends, the recommendation quality can be improved slightly.

Finally, to further investigate how the combination of user and trend profiles impacts personalized news recommendations, we evaluate the mixture strategy of combining trend and user profiles for different configurations. Figure 6 shows the results for varying parameter d when combining $P(u)$ and $T(I_r)$ (cf. Equation 8) and moreover for a varying the number of top k concepts selected for generating the trend profiles. The results reveal that personal user interest profiles seem to be more important than public trends as the recommendation quality with respect to MRR increases when the influence of $P(u)$ is increased. When combining $P(u)$ with $T_{t-TF\times IDF@800}(I_r)$ we achieve a global maximum in performance for $d = 0.6$ which also clearly improves over the strategy that is merely based on personal user interests ($d = 1$, $P(u)$). We conclude that user profiles are more important in the context of personalized news recommendations. However, by combining trend and user profiles we achieve the best recommendation performance.

V. Conclusion

In this paper, we developed a Twitter-based trend and user modeling framework. Our framework features functionality for enriching the semantics of tweets and therefore allows for the generation of semantically meaningful profiles which represent both personal interests and public trends. Those profiles can be re-used outside of Twitter by other applications that aim for trend-aware personalization. We have evaluated the trend and user modeling strategies in the context of news article recommendations. Given the trend and user modeling strategies, we showed that personal user interest profiles are more important for the news article recommendation process than public trends. By interweaving trend and user profiles we succeeded in further improving the recommendation quality.

Acknowledgements: The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no ICT 257831 (ImREAL project).

References

[1] 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)*. New York, NY, USA: ACM, 2010, pp. 591–600.

[2] T. Sakaki, M. Okazaki, and Y. Matsuo, "Earthquake shakes twitter users: real-time event detection by social sensors," in *Proceedings of the 19th international conference on World wide web*, ser. WWW '10. New York, NY, USA: ACM, 2010, pp. 851–860.

[3] M. Rowe, M. Stankovic, and P. Laublet, "Mapping Tweets to Conference Talks: A Goldmine for Semantics," in *Workshop on Social Data on the Web (SDoW '10), colocated with ISWC '10, Shanghai, China*, A. Passant, , J. Breslin, S. Fernandez, and U. Bojars, Eds., vol. 664. CEUR-WS.org, 2010.

[4] F. Abel, Q. Gao, G.-J. Houben, and K. Tao, "Analyzing User Modeling on Twitter for Personalized News Recommendations," in *International Conference on User Modeling, Adaptation and Personalization (UMAP), Girona, Spain*. Springer, July 2011.

[5] J. Huang, K. M. Thornton, and E. N. Efthimiadis, "Conversational Tagging in Twitter," in *Proceedings of the 21st ACM conference on Hypertext and hypermedia (HT '10)*, M. H. Chignell and E. Toms, Eds. New York, NY, USA: ACM, 2010, pp. 173–178.