

Inferring User Interests for Passive Users on Twitter by Leveraging Followee Biographies

Guangyuan Piao and John G. Breslin

Insight Centre for Data Analytics,
National University of Ireland Galway
IDA Business Park, Lower Dangan, Galway, Ireland
{guangyuan.piao}@insight-centre.org
{john.breslin}@nuigalway.ie

Abstract. User modeling based on the user-generated content of users on social networks such as Twitter has been studied widely, and has been used to provide personalized recommendations via inferred user interest profiles. Most previous studies have focused on *active users* who actively post tweets, and the corresponding inferred user interest profiles are generated by analyzing these users' tweets. However, there are also a great number of *passive users* who only consume information from Twitter but do not post any tweets. In this paper, we propose a user modeling approach using the *biographies* (i.e, self descriptions in Twitter profiles) of a user's *followees* (i.e., the accounts that they follow) to infer user interest profiles for *passive users*. We evaluate our user modeling strategy in the context of a link recommender system on Twitter. Results show that exploring the *biographies* of a user's followees improves the quality of user modeling significantly compared to two state-of-the-art approaches leveraging the *names* and *tweets* of followees.

1 Introduction

Online Social Networks (OSNs) have been growing rapidly since they first emerged in the early 2000's. A large number of users are now consuming different types of information (e.g., medical information, news) on OSNs [15] such as Twitter¹. Therefore, inferring interests for users of these OSNs can play an important role in providing them with personalized recommendations for content. Most previous studies have inferred user interest profiles from a user's posts, such as their tweets on Twitter. The research focus in these studies has been on the user modeling of *active users* who actively generate content on Twitter. However, the percentage of *passive users* in social networks is increasing² (e.g., 44% of Twitter users have never sent a tweet³). *Passive users* are not inactive accounts, but rather users that only consume information on social networks

¹ <https://twitter.com/>

² <http://www.corporate-eye.com/main/facebooks-growing-problem-passive-users/>

³ <http://guardianlv.com/2014/04/twitter-users-are-not-tweeting/>

without generating any content. In order to infer user interest profiles for passive users, some researchers have proposed linking *names* of followees (those whom a user is following) to Wikipedia⁴ entities, and then utilizing these entities to derive abstract category-based user interests [3]. For example, if a user is following famous football players such as `Cristiano_Ronaldo`, they find the Wikipedia entity for `Cristiano_Ronaldo`, and then utilize the categories of the corresponding Wikipedia entity to infer user interests. Although this approach can extract highly accurate Wikipedia entities to boost a user’s interest profile, it can only link popular Twitter accounts (e.g., the accounts of celebrities) to their corresponding Wikipedia entities. As a result, the information for a large percentage of a user’s followees is often ignored.

Another piece of information that forms an important part of followees’ profiles is their *biographies (bios)*. A *bio* on Twitter is a short personal description that appears in a user’s profile and that serves to characterize the user’s persona⁵. The length of a bio is limited to 160 characters. For example, Figure 1 shows a user named *Bob* who has filled his bio with “*Android developer. Educator.*”, which describes the user’s identity.

In this paper, we investigate the bios of followees as a source of information for boosting user interest profiles. The intuition behind this is that a user might be interested in “*Android development*” if the user is following *Bob*. Our hypothesis is that, given a large number of bios of a user’s followees, the entities mentioned in those bios can be leveraged for building quantified and qualified user interest profiles compared to using entities extracted based on the names of followees [3].

The contributions of our work are summarized as follows.

- We propose user modeling strategies leveraging the bios of followees for inferring a user’s interests by investigating two different interest propagation strategies.
- We evaluate our user modeling strategies against two state-of-the-art user modeling strategies for passive users in the context of a link recommender system on Twitter.

The organization of the rest of the paper is as follows. Section 2 gives some related work, and Section 3 describes our proposed approaches for inferring user interest profiles. In Section 4, we present the Twitter dataset for our study, and Section 5 describes the evaluation methodology of the study. Experimental results are presented in Section 6. Finally, Section 7 concludes the paper with some future work.




Fig. 1: Twitter profile.

2 Related Work

The largest area of work that is focused on inferring user interest profiles for *active users* is based on analyzing the tweets generated by them [1, 2, 9, 10, 13,

⁴ <https://www.wikipedia.org/>

⁵ <https://support.twitter.com/articles/166337>

14, 16, 17]. For example, Siehndel et al. [16] showed a prototype for generating user interest profiles based on the extracted entities from a user’s tweets, and then linking these entities to 23 top-level Wikipedia categories. Kapanipathi et al. [7] extracted Wikipedia entities from a user’s tweets, which were then used as activated nodes for applying various spreading activation functions based on a refined taxonomy of Wikipedia categories. As a result, a so-called weighted *Hierarchical Interest Graph* was generated for a given user. Instead of using Wikipedia categories, Piao et al. [14] and Orlandi et al. [11] leveraged DBpedia for propagating user interest profiles. DBpedia provides background knowledge about entities which not only includes the categories of entities, but also related entities via different properties. The authors of [14] showed that exploring some different structures of semantic information from DBpedia (i.e., categories as well as related entities) can improve the quality of user modeling in the context of a link (URL) recommender system on Twitter. Our work here is different from this line of work as we focus on inferring interests for *passive users* who do not generate tweets, but mostly just consume content from those that they follow on Twitter. In [16], the authors also suggested investigating other sources beyond tweets for user modeling. We address this research gap in our work.

Faralli et al. [5] leveraged the names of followees linked to Wikipedia entities, and then used these entities in order to infer user interest profiles for user recommendations. To the best of our knowledge, this work and the later work by [3] are the first ones exploring the use of followee profiles (in particular their names) for inferring user interest profiles, without analyzing any tweets. The authors in [5] have pointed out that leveraging followee profiles can build more stable and scalable user interest profiles than analyzing the tweets of followees. However, they also showed that only 12.7% of followees can be linked to Wikipedia entities on average. The most similar work to ours is [3]. Similar to [5], the authors in [3] first devised a method combining different heuristics for linking the followees of a user to Wikipedia entities. The linked entities were then used as activated nodes in a spreading activation function based on WiBi (Wikipedia Bitaxonomy [6]) in order to build abstracted *category-based* user interest profiles. Instead of leveraging the names of followees, we focus on the *bios* of followees for generating user interest profiles, and use the approach from [3] as one of our baseline methods (see Section 3.1).

3 User Modeling Approaches

In this section, we first describe two baseline methods (Section 3.1), and present our proposed user modeling approaches using two different propagation methods (Section 3.2). In this work, we define a user interest profile as follows.

Definition 1. *The interest profile of a user $u \in U$ is a set of weighted user interests (e.g., entities or categories of entities). The weight of each interest $i \in I$: $w(u, i)$ indicates the importance of the interest i with respect to a user u .*

$$P_u = \{(i, w(u, i)) \mid i \in I, u \in U\} \quad (1)$$

where I denotes the set of user interests, and U denotes the set of users.

3.1 Baseline Methods

SA(followees_name): Given a Twitter user u , the approach from [3] leverages the names of u 's followees for user modeling. The input of this approach is a Twitter account, and the output is a *category-based* user interest profile obtained via a spreading activation method. It has three main steps for generating user interest profiles.

1. Fetch user's followees.
2. Link these to corresponding Wikipedia entities.
3. Apply a spreading activation method for the linked entities from step 2 to generate category-based profiles based on WiBi (Wikipedia Bitaxonomy⁶).

For example, if the user account *@bob* in Figure 1 is following *@BillGates* (the Twitter account for *Bill_Gates*), this approach searches for the name *Bill_Gates* on Wikipedia in order to find the right entity for the Twitter account *@BillGates* using different heuristics. We used the author's implementation⁷ [3] to link a user's followees to Wikipedia entities. The linked Wikipedia entities are activated nodes with $w(u, i) = 1$ for the next step. This approach further applies a spreading activation function from [7] (see Algorithm 2) to propagate user interests from the extracted Wikipedia entities to Wikipedia categories, e.g., from *Bill_Gates* to *Category:Directors_of_Microsoft*. The spreading activation function is defined as follows:

$$a_t(j) \leftarrow a_{t-1}(j) + d_{subnodes} \times b_j \times a_{t-1}(i) \quad (2)$$

$$d_{subnodes} = 1/\log N_{subnodes} \quad (3)$$

$$b_j = \frac{N_{e_j}}{N_{e_{cmax}}} \quad (4)$$

where j is a node (category) being activated, and i is a sub-node of j which is activating j . $d_{subnodes}$ is a decay factor based on the number of sub-nodes (sub-entities or categories) of the current category, and b_j is an *Intersect Booster* factor introduced in [7]. b_j is calculated by Algorithm 4, where N_{e_i} is the total number of entities activating node j , and $cmax$ is the sub-category node of j which has been activated with the maximum number of entities [7]. The weight of a node is accumulated if there are several sub-nodes activating the node.

As none of the previous studies [3, 5] showed the performance of using followees' profiles (i.e., the names or bios of followees) compared to using followees' tweets, we also include a baseline method [4] using the tweets of followees for

⁶ <http://wibitaxonomy.org/>

⁷ https://bitbucket.org/beselch/interest_twitter_acmsac16

inferring user interest profiles to investigate the comparative performance of the two different approaches.

HIW(followees.tweet): This approach [4] extracts so-called *high-interest words* from each followee of a user u . The *high-interest words* consist of the top 20% of words in the ranked word list from a followee f 's tweets. The latest 200 tweets from each followee are considered for our study, which results in over 13,940,000 tweets from the followees of 48 users (see Section 5). To construct the interest profile of u , high-interest words from all followees are aggregated by excluding the words mentioned only in a single followee's tweets. Finally, the weight of each word in u 's profile is measured as $w(u, i) =$ the number of u 's followees who have i as their high-interest words.

3.2 Proposed Approaches

Figure 2 presents the overview of our user modeling process, which consists of three main steps.

1. Fetch user's followees.
2. Extract Wikipedia/DBpedia [8] entities to the bios of followees.
3. Apply one of the interest propagation methods:
 - (a) $SA(\text{followees_bio})$
 - (b) $IP(\text{followees_bio})$.

Our approach is different from the baseline method $SA(\text{followees_name})$ especially in step 2. We use the Aylien API⁸ to extract entities from the bios of a user's followees. The number of occurrences of each entity in the bios of followees is counted for measuring the importance of the entity with respect to a targeted user for inferring his or her interests.

SA(followees.bio): As one of our goals is investigating whether using the bio information of followees can improve the quality of user modeling compared to

⁸ <http://aylien.com/>




Fig. 2: Overview of our proposed approach

using the names of followees, we applied the same spreading activation algorithm (Algorithm 2) for the entities extracted from the bios of followees. Therefore, the difference between this approach and $SA(\text{followees_name})$ is the set of activated nodes for propagation. For $SA(\text{followees_bio})$, the activated nodes are extracted entities from the bios of a user’s followees with $w(u, i) = N_i$ which denotes the frequency of an interest i in their bios. Similar to $SA(\text{followees_name})$, the output of this approach is a *category-based* user interest profile.

IP(followees_bio): Differing from the propagation of user interests using the taxonomy of Wikipedia categories, this approach uses an interest propagation method from [14]. The propagation method extends user interests using *related entities* as well as corresponding *categories* from DBpedia. DBpedia is a knowledge graph providing cross-domain knowledge extracted from Wikipedia. The difference between the WiBi taxonomy and the DBpedia graph is presented in Figure 3. As we can see from Figure 3(b), the DBpedia graph provides related entities in addition to the categories of an entity. For example, as well as providing categories for the entity `Bill_Gates` via the property `dc9:subject`, DBpedia also gives related entities such as `Microsoft` via the property `dbo10:board`. Therefore, as distinct from both $SA(\text{followees_name})$ and $SA(\text{followees_bio})$, the output here is a user interest profile consisting of propagated *categories* as well as *entities*.

The authors in [14] also applied some *discounting strategies* for propagated categories, and entities via different properties. For example, a propagated category is discounted based on the log scale of the numbers of sub-pages (SP) and sub-categories (SC , see Algorithm 5). A propagated entity is discounted based on the log scale of the number of occurrences of a property in the DBpedia graph (P , see Algorithm 6), i.e., if the property appears frequently in the graph, the entities extended via this property should be discounted heavily. In addition, α

⁹ The prefix `dc` denotes `http://purl.org/dc/terms/`

¹⁰ The prefix `dbo` denotes `http://dbpedia.org/ontology/`

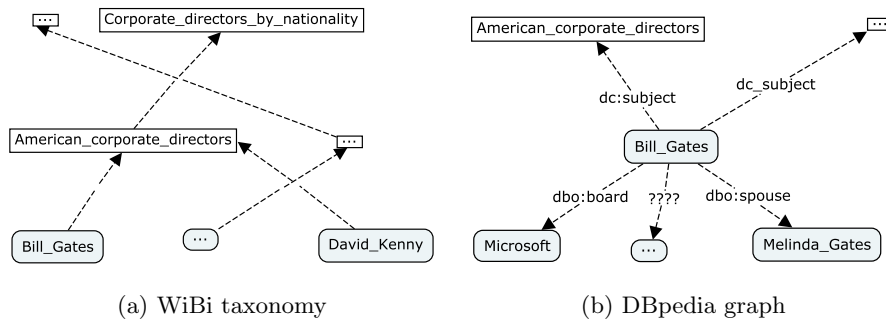


Fig. 3: Examples of WiBi taxonomy and DBpedia graph.

is a decay factor for the propagation from directly extracted entities to related categories or entities ($\alpha = 2$ as in the study [14]).

$$CategoryDiscount = \frac{1}{\alpha} \times \frac{1}{\log(SP)} \times \frac{1}{\log(SC)} \quad (5)$$

$$PropertyDiscount = \frac{1}{\alpha} \times \frac{1}{\log(P)} \quad (6)$$

For all of the aforementioned user modeling approaches, after propagating user interest profiles, we further apply IDF (Inverse Document Frequency) to the weights of user interests in order to discount user interests appearing frequently in profiles of users. Finally, the user interest profiles are normalized so that the sum of the weights of user interests is equal to one.

4 Dataset

We used a Twitter dataset from [13] for our study. The dataset consists of 480 randomly selected Twitter users, and the tweets generated by them. As the focus of our study is using the followees of Twitter users for generating user interest profiles, we further crawled information on the followees for those 480 users. It was possible to crawl followees for 461 of the original 480 users via the Twitter API¹¹ as some users did not exist anymore. As a result, the dataset consists of 461 users, and 902,544 followees of these users. Among these followees, we found that 812,483 users (around 90%) had filled out the bio field in their Twitter profiles.

Dataset for our experiment. As there can be a great number of followees even for a small number of users, we randomly selected 50 users with a corresponding set of 84,646 followees for our experiment. The descriptive statistics of the dataset are presented in Table 1. These 50 users have 77,825 distinct followees in total. 10% of these followees can be linked to Wikipedia entities using the approach from [3]. In contrast, 72,145 out of 77,825 (over 90%) followees have bios.

Table 1: Descriptive statistics of the dataset

# of users	50
# of followees	84,646
# of distinct followees	77,825
# of followees whose names can be linked to Wikipedia entities	7,785 (10%)
# of followees that have bios	72,145 (92.7%)

Comparison of extracted entities using names and bios. As the entities either linked via the names or extracted from the bios of followees play

¹¹ <https://dev.twitter.com/rest/public>

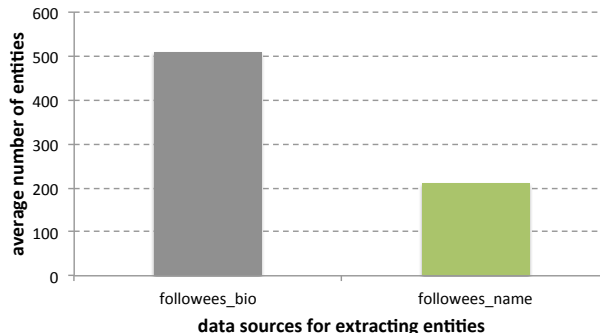


Fig. 4: Number of entities extracted via names and bios of followees.

a fundamental role in propagating user interests, we analyzed the number of entities that can be extracted using the two different sources. Figure 4 shows the difference between using the names and bios of followees in terms of the number of extracted entities. We can observe that using the bios of followees provides more than twice the number of entities when compared to using the names of followees. On average, 509 entities can be extracted for each user using the bios of followees, and 210 entities can be extracted for each user using the names of followees. This indicates that using the bios of followees can generate more quantified user interest profiles. We now move on to investigate whether the quantified user interest profiles generated by analyzing followees’ bios have a higher quality as well, compared to those generated by linked entities based on the names of followees.

5 Evaluation Methodology

We were interested in finding out if leveraging the bios of followees for a passive user improves the quality of user modeling compared to using the names of followees. To this end, we evaluate different user interest profiles generated by different user modeling strategies in the context of a link (URL) recommender system on Twitter. Given this focus of our study, we applied a lightweight content-based recommendation algorithm for generating recommendations in the same way as previous studies [2, 13, 14].

Definition 2. *Recommendation Algorithm:* given a user profile P_u and a set of candidate links $N = \{P_{i_1}, \dots, P_{i_n}\}$, which are represented via profiles using the same vector representation, the recommendation algorithm ranks the candidate items according to their cosine similarity to the user profile.

Link (item) profiles were generated by applying the same propagation strategies applied for generating user interest profiles based on the content of

a link. For example, given a link l , we first extract Wikipedia/DBpedia entities from the content of l , and then apply one of the aforementioned interest propagation strategies (see Section 3.2).

To construct a ground truth of links (URLs) for users, we assumed that links shared via a user’s tweets were links representing a user’s interests. Therefore, we further crawled the timelines of the 50 randomly selected users using the Twitter API, and extracted links shared in their tweets. In the same way as [14], we considered links that have at least four concepts to filter out non-topical ones which were automatically generated by third-party applications such as Swarm¹². 48 users were left as two of the 50 users had no topical links. On average, there were 31.46 links shared by a user. The candidate set of links consists of 1,377 distinct links shared by these 48 users. We then blinded the tweets of the 48 users, and used their followees’ information only for building user interest profiles.

Given a user interest profile and a link profile in the candidate set, the recommender system measures similarities between the two profiles, and then gives the top- N links having the highest similarity scores to the user. We focused on $N = 10$ in our experiment, i.e., the recommendation system would list 10 link recommendations to a user. We used four different evaluation metrics as used in the literature [1, 2, 11, 12, 14] for measuring the quality of recommendations using different user interest profiles as input.

- **MRR** The *MRR* (Mean Reciprocal Rank) indicates at which rank the first item *relevant* to the user occurs on average.
- **S@N** The Success at rank N (*S@N*) stands for the mean probability that a relevant item occurs within the top- N ranked.
- **R@N** The Recall at rank N (*R@N*) represents the mean probability that *relevant* items are successfully retrieved within the top- N recommendations.
- **P@N** The Precision at rank N (*P@N*) represents the mean probability that retrieved items within the top- N recommendations are *relevant* to the user.

A significance level of alpha was set to 5% for all statistical tests. We used the *bootstrapped paired t-test*¹³ to test the significance.

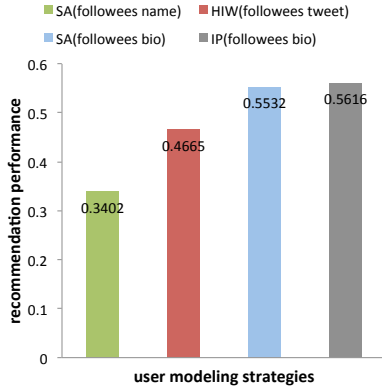
6 Results

Figure 5 presents the results of recommendations using different user modeling strategies in terms of four different evaluation metrics. Overall, *IP(followees.bio)* provides the best performance in terms of all evaluation metrics except S@10.


Comparison between using the names and bios of followees. From Figure 5, we observe that *IP(followees.bio)* as well as *SA(followees.bio)* which use the bios of followees for user modeling outperform

¹² <https://www.swarmapp.com>


¹³ http://www.sussex.ac.uk/its/pdfs/SPSS_Bootstrapping_22.pdf



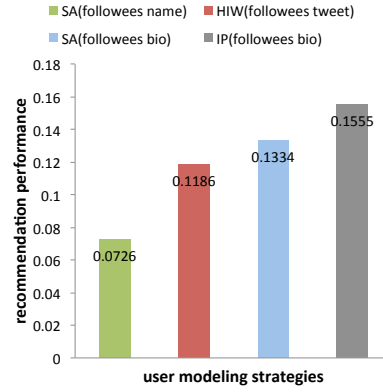
(a) MRR (Mean Reciprocal Rank)



(b) S@10 (Success rate at top-10)



(c) P@10 (Precision at top-10)



(d) R@10 (Recall at top-10)

Fig. 5: Results of the recommender system using different evaluation metrics.

SA(followees_name) which uses the names of followees. A significant improvement of *SA(followees_bio)* over *SA(followees_name)* in MRR (+63%), S@10 (+30%), P@10 (+78%), and R@10 (+84%) can be noticed ($p < 0.05$). With the same spreading activation method applied to two different sources: the names and bios of followees, the difference in terms of the four evaluation metrics clearly shows that exploring the bios of followees of passive users can infer better quality user interest profiles compared to using the names of followees.

Comparison between using the bios and tweets of followees. Figure 5 also shows the performance of the baseline method *HIW(followees_tweet)*, which analyzes followees' tweets for inferring *word-based* user interest profiles. The results show that our user modeling strategies using bios of followees out-

perform the baseline method in terms of all evaluation metrics. For instance, $IP(\textit{followers_bio})$ outperforms $HIW(\textit{followers_tweet})$ significantly in terms of S@10 as well as P@10 ($p < 0.05$). Considering $HIW(\textit{followers_tweet})$ needs to analyze over 13,940,000 tweets of followers whereas $IP(\textit{followers_bio})$ analyzes only around 77,000 bios of followers to build interest profiles for 48 users, our approach as well as $SA(\textit{followers_name})$ [5], both of which use followers' profiles (i.e., the names or bios), are more scalable in the context of OSNs such as Twitter. On the other hand, the performance of $HIW(\textit{followers_tweet})$ suggests that analyzing all the tweets of followers can lead to noisy information as an input for user modeling, which might decrease the quality of the inferred user interest profiles. For instance, a user who is following *Bob* (see Figure 1) might be interested in “*Android development*”, however, tweets posted by *Bob* would not only contain those on the topic of “*Android development*” but also other diverse topics that *Bob* might be interested in.

Comparison between using WiBi taxonomy and DBpedia graph.

Regarding the interest propagation strategies, $IP(\textit{followers_bio})$, which leverages the DBpedia graph for interest propagation, has better performance in terms of MRR, P@10 and R@10 when compared to $SA(\textit{followers_bio})$. On the other hand, $SA(\textit{followers_bio})$ has better performance in terms of S@10 than $IP(\textit{followers_bio})$. The results suggest that $IP(\textit{followers_bio})$ provides a greater number of preferred links to users who have successfully received recommendations, i.e., a higher P@10 value when S@10=1.

7 Conclusions

In this paper, we were interested in investigating whether leveraging the *bios* of followers can infer quantified as well as qualified user interest profiles. To this end, we proposed user modeling strategies leveraging the *bios* of followers for inferring user interests on Twitter. We evaluated our user modeling strategies compared to a state-of-the-art approach using the *names* of followers, and a approach using the *tweets* of followers for user modeling. The results are promising. They show that $IP(\textit{followers_bio})$, which leverages entities extracted from the bios of followers and applies an interest propagation strategy using the DBpedia graph, provides the best performance, and significantly improves upon two baseline methods in the context of a link recommender system. As a further step, we plan to study how we can combine different interest propagation strategies using the WiBi taxonomy and the DBpedia graph to improve the quality of user modeling.

Acknowledgments. This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289 (Insight Centre for Data Analytics).

References

1. Abel, F., Gao, Q., Houben, G.J., Tao, K.: Analyzing user modeling on twitter for personalized news recommendations. In: *User Modeling, Adaption and Personalization*, pp. 1–12. Springer (2011)
2. Abel, F., Hauff, C., Houben, G.J., Tao, K.: Leveraging User Modeling on the Social Web with Linked Data. In: *Web Engineering SE - 31*, pp. 378–385. Springer (2012)
3. Besel, C., Schlötterer, J., Granitzer, M.: Inferring Semantic Interest Profiles from Twitter Followeres: Does Twitter Know Better Than Your Friends? In: *Proceedings of the 31st Annual ACM Symposium on Applied Computing*. pp. 1152–1157. SAC '16, ACM, New York, NY, USA (2016)
4. Chen, J., Nairn, R., Nelson, L., Bernstein, M., Chi, E.: Short and tweet: experiments on recommending content from information streams. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. pp. 1185–1194. ACM (2010)
5. Faralli, S., Stilo, G., Velardi, P.: Recommendation of microblog users based on hierarchical interest profiles. *Social Network Analysis and Mining* 5(1), 1–23 (2015)
6. Flati, T., Vannella, D., Pasini, T., Navigli, R.: Two Is Bigger (and Better) Than One: the Wikipedia Bitaxonomy Project. In: *ACL (1)*. pp. 945–955 (2014)
7. Kapanipathi, P., Jain, P., Venkataramani, C., Sheth, A.: User Interests Identification on Twitter Using a Hierarchical Knowledge Base. In: *The Semantic Web: Trends and Challenges*. pp. 99–113. Springer (2014)
8. Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P.N., Hellmann, S., Morsey, M., van Kleef, P., Auer, S.: Dbpedia-a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web Journal* (2013)
9. Michelson, M., Macskassy, S.A.: Discovering users' topics of interest on twitter: a first look. In: *Proceedings of the fourth workshop on Analytics for noisy unstructured text data*. pp. 73–80. ACM (2010)
10. Mislove, A., Viswanath, B., Gummadi, K.P., Druschel, P.: You are who you know: inferring user profiles in online social networks. In: *Proceedings of the third ACM international conference on Web search and data mining*. pp. 251–260. ACM (2010)
11. Orlandi, F., Breslin, J., Passant, A.: Aggregated, interoperable and multi-domain user profiles for the social web. In: *Proceedings of the 8th International Conference on Semantic Systems*. pp. 41–48. ACM (2012)
12. Piao, G.: User Modeling on Twitter with WordNet Synsets and DBpedia Concepts for Personalized Recommendations. In: *The 25th ACM International Conference on Information and Knowledge Management*. ACM (2016)
13. Piao, G., Breslin, J.G.: Analyzing Aggregated Semantics-enabled User Modeling on Google+ and Twitter for Personalized Link Recommendations. In: *User Modeling, Adaptation, and Personalization*. pp. 105–109. ACM (2016)
14. Piao, G., Breslin, J.G.: Exploring Dynamics and Semantics of User Interests for User Modeling on Twitter for Link Recommendations. In: *12th International Conference on Semantic Systems*. pp. 81–88. ACM (2016)
15. Sheth, A., Kapanipathi, P.: Semantic Filtering for Social Data. *IEEE Internet Computing* 20(4), 74–78 (2016)
16. Siehdnel, P., Kawase, R.: TwikiMe!: user profiles that make sense. In: *Proceedings of the 2012th International Conference on Posters & Demonstrations Track-Volume 914*. pp. 61–64. CEUR-WS. org (2012)
17. Zarrinkalam, F., Fani, H., Bagheri, E., Kahani, M.: Inferring Implicit Topical Interests on Twitter. In: *European Conference on Information Retrieval*. pp. 479–491. Springer (2016)