

Tag Recommendation for Folksonomies Oriented towards Individual Users

Marek Lipczak

Faculty of Computer Science, Dalhousie University, Halifax, Canada, B3H 1W5
lipczak@cs.dal.ca

Abstract. Tagging has become a standard way of organizing information on the Web, particularly in folksonomies – data repositories freely created by communities of users. A few tags attached to each resource create a bridge between heterogeneous data and users accustomed to keyword-based search and browsing. To establish this connection, tagging requires users to manually define tags for each resource they enter to the system. This potentially time-consuming step can be eased by tag recommender systems, which propose terms that users may choose to use as tags. This paper suggests and evaluates potential sources of recommended tags, focusing on folksonomies oriented towards individual users. These suggestions are used to propose a three-step tag recommendation system. Basic tags are extracted from the resource title. In the next step, the set of potential recommendations is extended by related tags proposed by a lexicon based on co-occurrences of tags within resource’s posts. Finally, tags are filtered by the user’s personomy – a set of tags previously used by the user.

1 Introduction

Folksonomy services allow users to store and share various types of Internet resources. The content of folksonomies is completely defined by communities of their users. Large number of creators and resources push the folksonomies from the traditional hierarchical data structure design based on directories created by system editors (e.g., Open Directory Project¹) to tag-based taxonomies defined jointly by service users (e.g., BibSonomy², del.icio.us³, Flickr⁴, Technorati⁵). While adding a resource to the system, users are asked to define a set of tags – keywords which describe it and relate it to other resources gathered in the system. To ease this process, some folksonomy services recommend a set of potentially matching tags. Proposing a tag recommendation system was a task of ECML PKDD discovery challenge 2008⁶. This paper presents a tag recommendation system submitted to the challenge.

¹ <http://www.dmoz.org/about.html>

² <http://bibsonomy.org/help/about/>

³ <http://del.icio.us/about/>

⁴ <http://flickr.com/about/>

⁵ <http://technorati.com/about/>

⁶ <http://www.kde.cs.uni-kassel.de/ws/rsdc08/>

The formal definition of folksonomy can be found in [6]. A folksonomy is a collection of resources entered by users in posts. Each *post* consists of a resource and a set of tags attached to it by a user. Generally, the resource is specific to the user who added it to the system. However, for some types of resources (e.g., bookmarks) identical resources can be added to the system by different users. In the latter case, by the *set of resource tags* we denote all tags attached to a given resource by various users.

Folksonomies can be classified into two types based on the objective of the tagging process. The first type, represented by BibSonomy and del.icio.us, treats resources (e.g., personal bookmarks) as an individual property of a user. Here, the aim of tags is to create a repository tailored to individual user interests. In this paper, this type is referred to as *folksonomies oriented towards individual users*. The second type of folksonomies, represented by Flickr and Technorati, is a shared repository of public resources (e.g., blog entries). In this case tags are added keeping in mind a broad audience that in the future would like to search for the resource. In this paper, this type is referred to as *folksonomies oriented towards broad audience*. As the reason of tagging a resource is fundamentally different, we may expect that a tag recommendation system that suits one folksonomy type would be inappropriate for the other. This paper focuses on the first type, proposing a tag recommender for individual users.

2 Related work

The attention of researchers is mostly directed to tag recommendation systems for broad audience folksonomies. TagAssist [12] is a system designed to recommend tags of blog posts. The recommendation is built on tags previously attached to similar resources. Earlier, meaning disambiguation is performed based on co-occurrence of tags in the complete repository. Co-occurrence of tags was also used by Sigurbjörnsson and van Zwol [11] to propose tags that complement user-defined tags of photographs in Flickr.

The problem of tag recommendation in folksonomies oriented towards individual users was addressed by Jäschke et al. [7]. They compared a number of recommendation techniques including collaborative filtering, PageRank, and its modification suited for folksonomies – FolkRank. The evaluation showed that the FolkRank based recommender outperforms other approaches; however, the tests were performed on a dense core of folksonomy, thus might be not representative.

Most of the tag recommendation systems are based on the tags that are already present in the system. An exception from this rule is the system presented by Lee and Chun [9]. The system recommends tags retrieved from the content of a blog, using artificial neural network. The network is trained based on statistical information about word frequencies and lexical information about word semantics extracted from WordNet.

Schmitz et al. [10] proposed association rule mining as a technique that might be useful in the tag recommendation process. The intuition behind this concept was also used in the system presented by this paper.

3 Examined dataset

All presented experiments and the evaluation of proposed tag recommendation system were performed on a snapshot of BibSonomy [5] containing 2,570 users, 242,175 resources and 274,139 posts (after preprocessing). The snapshot was provided by the organizers of the ECML PKDD discovery challenge 2008. The preprocessing phase included removing useless tags (e.g., “system:unfiled”), changing all letters to lower case and removing non-alphabetical and non-numerical characters from tags.

The statistical characteristics of folksonomies have been an object of many research publications [2, 3, 8, 11]. In the following sections I present experiments particularly important from the perspective of the tag recommendation task.

3.1 General characteristics

The frequency distribution of tags from the Bibsonomy snapshot shows that mid- and low-frequency tags follow Zipf’s distribution (Fig. 1). Zipf’s distribution does not hold for high-frequency tags. The frequency distribution of tags from Flickr, which represents folksonomies oriented towards broad audience shows important differences [11]. Flickr’s low-frequency tags does not follow Zipf’s distribution. A possible explanation of this fact is a smaller number of user specific tags in comparison to folksonomies oriented towards individual users. In addition, Flickr’s high-frequency tags follows Zipf’s distribution and are too general to be used as recommendation. The list of the most frequent tags from Bibsonomy (“software”, “web20”, “tools”, “web”, “blog”) shows that tag recommenders for folksonomies oriented towards individual users should not ignore high-frequency terms.

The difference between two folksonomy types may have impact on the efficiency of applied tag recommendation methods. A commonly used collaborative filtering approach is based on the intuition that the best recommendation consists of tags attached to the resource by people similar to the user. This approach proved its quality in many recommendation systems; however, the intuition behind it can be deceiving. Folksonomies like BibSonomy or del.icio.us are mainly designed as a collection of repositories of individual users. By adding posts, each user defines his/her own set of used tags – personomy [6], which describes the resources from a user’s point of view. As a result, users addressing similar resources do not have to use similar tags, and similar personomies do not have to be associated with similarity in tagged resources. In fact, there is no such correlation in the processed BibSonomy snapshot. The cosine similarity between users calculated based on tags seems to be uncorrelated with that calculated based on resources (Fig. 2). In this situation recommending tags assigned to a resource by similar users (collaborative filtering) should give similar results as recommending the tags frequently attached to the resource by any user. This conclusion seems to be confirmed by the experiment presented by Jäschke et al. [7]. Minding the limitations of the collaborative approach I decided to focus on a tag space that is directly related to a post.




Fig. 1. The overall frequency distribution of tags (after preprocessing and removing posts classified as imported).




Fig. 2. Cosine similarity between each pair of users calculated based on tags (tf-idf weights) and resources (binary weights). The two values seem to be independent.

3.2 Characteristics based on individual posts

Considering only the direct surrounding of the post, the potential tag recommendations can be obtained from the resource itself, the set of tags attached to the resource in previous posts, or the set of tags that were already used by the user (user’s personomy). Exploiting tags from the resource depends on the folksonomy character. In BibSonomy the resource can be a bibtex entry or a webpage bookmark. The first contains bibliographic information about a research publication including its title and abstract. The second contains web-page title and URL. Preliminary experiments showed that using title words as tags outperforms the results of abstracts and URLs. The latter two contain lesser amount of correct tags. The title is the only element that joins both resource types and it is common in other folksonomies, which are its additional advantages. I decided to use the title as the representation of resource.

To evaluate the three potential sources of tag recommendations, namely words from the resource title, resource tags and user’s personomy, I checked for each post if its tags can be found in any of these sources associated with all other posts in the folksonomy. The quality of sources was measured by precision (i.e., number of correct tags retrieved divided by the total number of retrieved tags) and recall (i.e., number of correct tags retrieved divided by the total number of correct tags). These are standard information retrieval metrics [4]. The value of recall was averaged over all tested posts. The averaged recall informs us how many correct tags can be found in a source. The value of precision was averaged only over posts, for which the source returned any tags. Precision averaged this way is the ratio of correct tags among all tags retrieved. In addition, I present the total number of potential tags obtained from the sources, and the number of correct tags among them (Fig. 3).

User’s personomy is the richest source of correct tag recommendations. For the tested BibSonomy snapshot it gave access to 90% of tags from test posts. On

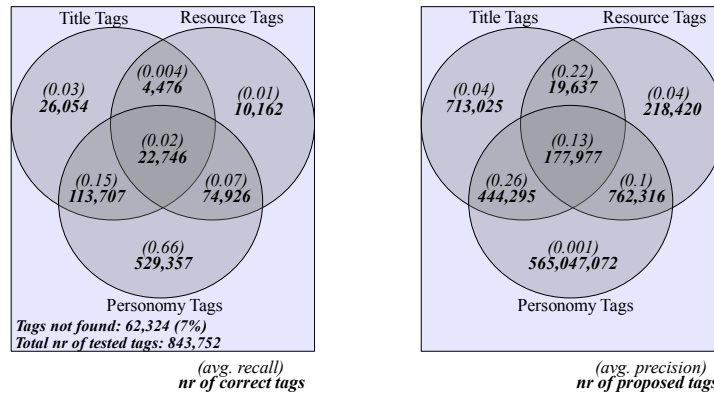


Fig. 3. Venn diagrams presenting average recall, plus the number of correct tags found in three potential sources of tags (left) and average precision, plus the total number of tags retrieved from these sources (right).

the other hand, correct tags from personomy are accompanied by a large number of incorrect tags (precision around 0.001). Compared to tags retrieved from personomy, the recommendation based on resource title is much more precise; however, the number of correct tags found this way is lower. In addition, most of these tags can be also found in the user’s personomy. Finally, both recall and precision values show that resource tags are not a good source of potential tag recommendations. The character of each tag recommendation source and their potential usability in tag recommendation system are discussed in the following sections.

Resource title Resource title appears to be the most robust source of tag recommendations. Among all posts in processed BibSonomy snapshot only 51 resource titles were unable to produce any tags (no letters or numbers in the title). In addition, among all discussed sources the title seems to be the most strongly related to the resource. The drawback of this source is low recall which makes the title inappropriate as a stand-alone tag recommender. The title is a simplified natural language sentence, which should be cleaned of words with no informative value (e.g., stopwords).

Resource tags Tags assigned to the resource by other folksonomy users are not a good source of tag recommendations. One of the reasons is the sparsity of data; 92% of resources were added to the system only once. This fact significantly limits the possible recall of this source of tags. The other issue is the personal character of posts (discussed in section 3.1), which hurts the precision of retrieved tags. The variety of tags attached by users creates, however, another application of resource tag sets. Mining relations between tags attached to the same resource can result in a simplified semantic lexicon. The lexicon would not give us the

information about the character of relation, but given a tag, the lexicon can point out related tags which are also potential recommendations. The lexicon consists of general relations between tags and can be used independently of the resources. This fact reduces the negative impact of data sparsity. In addition, it is suited for a particular folksonomy and it can capture specific relations between its tags.

Personomy tags Building his/her personomy the user is interested in representing his/her interests using a limited number of tags. The same tag will be attached to resources fitting a particular interest, for example, all articles related to user's master thesis will be tagged by the same keyword. In addition, users are likely to stick with one lexical form of a word or expression, for example using constantly singular or plural form of a noun (e.g., "*publication*" or "*publications*"). These are the reasons why we are likely to find a lot of good recommendations among user's tags. The problem is that the choice of the lexical form or the word that describes the interest is completely up to the user. For the given example of resources related to master thesis the tag may be "*masterthesis*", "*msc*", "*thesis*", "*work*", or any other that according to user's opinion conveys the information.

To describe the resource more accurately users pick additional tags, specific not only to the user, but also to the resources. This is likely the cause of a large number of low-frequency tags (see section 3.1) and complicates the process of retrieving potential recommendations from personomy.

4 Tag recommendation system

The tag recommendation system (Algorithm 1), described in this section, is based on observations from the presented statistical experiments. The system is built of three steps. The first step produces tags from resource title words and assigns a score that represents their usefulness for previously tagged resources. The second step uses the resource tag based lexicon to propose tags related to tags taken from the title. The third step checks the tags proposed by the lexicon against user's personomy. The tags recommended to the user are a union of most promising tags produced in step one and three. The following sections give the detailed description of each step.

Extraction of title based tags The resource title is divided into words, which are then cleaned of non-alphabetical and non-numerical characters. The system assigns a score to each word, which represents the probability of being chosen as a tag – number of times being chosen as a tag divided by the number of occurrences. If the word occurred in the titles of previously entered resources less than 100 times its probability of being a correct tag is set to 0.1 which is an empirically estimated value for low-frequency tags. The probability score is introduced to reduce the impact of stopwords. It is important to notice that the standard stopwords list, which is often used in information retrieval systems, is

Algorithm 1: Tag recommendation system

Data: a resource p_{res} and user u

Result: a set of recommended tags $T_{Recommendation}$, a tag consist of a keyword (w) and recommendation score (**score**)

```
begin
  /*Step 1 - Extraction of title based tags*/
   $W_{Title} \leftarrow extractTitleWords(p_{res})$ 
   $T_{Title} \leftarrow \emptyset$ 
  foreach  $w \in W_{Title}$  do
     $T_{Title}$  add  $makeTag(w, getPriorUsefulness(w))$ 
  /*Step 2 - Retrieval of tags related to title*/
  foreach  $t \in T_{Title}$  do
     $T_{t\ Related} \leftarrow \emptyset$ 
     $T_{t\ RelTags} \leftarrow \emptyset$  // related tags from Tag-to-Tag lexicon
     $T_{t\ RelTitle} \leftarrow \emptyset$  // related tags from Title-to-Tag lexicon
    foreach  $r \in getRelated(l_{TagToTag}, t)$  do
       $T_{t\ RelTags}$  add  $makeTag(r.w, t.score * getRelScore(l_{TagToTag}, t, r))$ 
    foreach  $r \in getRelated(l_{TitleToTag}, t)$  do
       $T_{t\ RelTitle}$  add  $makeTag(r.w, t.score * getRelScore(l_{TitleToTag}, t, r))$ 
     $T_{t\ RelTags} \leftarrow limitSize(T_{t\ RelTags}, 20)$ 
     $T_{t\ RelTitle} \leftarrow limitSize(T_{t\ RelTitle}, 20)$ 
     $T_{t\ Related} \leftarrow unionProb(T_{t\ RelTags}, T_{t\ RelTitle})$ 
   $T_{Related} \leftarrow unionProb(T_{t_1\ Related}, \dots, T_{t_n\ Related})$ 
  /*Step 3 - Personomy based filtering*/
   $P \leftarrow getPersonomy(u)$ 
  foreach  $t \in T_{Related}$  do
     $T_{t\ RelPersonomy} \leftarrow \emptyset$  // tags retrieved from user's personomy
    if  $t \in P$  then
      foreach  $r \in P$  do
         $T_{t\ RelPersonomy}$  add  $makeTag(r.w, t.score * getRelScore(P, t, r))$ 
   $T_{RelPersonomy} \leftarrow unionProb(T_{t_1\ RelPersonomy}, \dots, T_{t_n\ RelPersonomy})$ 
   $T_{RelPersonomy} \leftarrow normalizeScores(limitSize(T_{RelPersonomy}, 10))$ 
   $T_{Title} \leftarrow normalizeScores(T_{Title})$ 
   $T_{Recommendation} \leftarrow limitSize(unionProb(T_{Title}, T_{RelPersonomy}), 10)$ 
end
```

not sufficient here, because we have to deal with titles in various languages and stopwords specific for the folksonomy (e.g., word “page” is frequent in web-page titles, but it is rarely used as a tag).

Retrieval of tags related to title The most important element of this step is the definition of the lexicon. It can be built based on two types of relations. As introduced in section 3.2, the lexicon can be built based on tags attached to the same resource, which are considered as related. The calculation of the factor that represents the relation strength can be solved based on various approaches

(e.g., association rule mining). In the presented system the score for a tag t_1 is the number of its co-occurrences with another tag t_2 among all resources, divided by the total number of occurrences of tag t_1 . The score is analogous to the confidence score (Eq. 1) in association rule mining [1].

$$confidence(t_1, t_2) = \frac{support(\{t_1 \cap t_2\})}{support(\{t_1\})} \quad (1)$$

Considering title words as the source of tags we can think of the second type of the lexicon representing relations between the words extracted from resource title and resource tags. The method of construction is analogical to the previous lexicon, the only difference is that tag t_1 is drawn from the title not the resource tags.

Both lexicons present different perspective of tag relations and give slightly different results (Table 1). The latter approach seems to be more adequate to the input tags; however, it is biased for general words that are often used in the title. For this type of words the related tags given by the second lexicon are simply the most frequently used tags (Table 2). To avoid the need of disambiguation between words more appropriate for either of the lexicons I decided to join the list of related tags produced by both of them (limited to twenty tags). The scores of tags that were present in both lists are summed as they were probabilities of two independent events.

This step is performed independently for each tag extracted from the title. Based on the lexicon the list of related tags with scores defining the strength of relation is retrieved. Finally, the lists are joined. Scores of multiple occurrences of identical tags are summed as they were independent probabilistic events, where the probability is defined by the relation score. Tags related to a word that is not likely to become a tag (e.g., “page”) are also not good candidates for recommendation. These are very general terms which are hard to connect with any concept. This is the reason why before joining the relation score is multiplied by the title tag score computed in the previous step.

Tag-to-Tag lex. occurrence: 317		Title-to-Tag lex. occurrence: 204		Tag-to-Tag lex. occurrence: 53		Title-to-Tag lex. occurrence: 2439	
Tag	Score	Tag	Score	Tag	Score	Tag	Score
1. semantics	1.000	semantics	0.392	1. home	1.000	software	0.081
2. semanticweb	0.306	semanticweb	0.348	2. page	0.113	tools	0.073
3. ontology	0.177	semantic	0.313	3. software	0.094	computing	0.064
4. semantic	0.167	folksonomy	0.215	4. server	0.075	java	0.059
5. semweb	0.158	tagging	0.196	5. photos	0.056	opensource	0.051

Table 1. Top 5 tags related to “semantics” according to two types of lexicon. **Table 2.** Top 5 tags related to “home” according to two types of lexicon.

Personomy based filtering The set of tags retrieved in the second step is likely to consist of a lot of correct recommendations. However, low precision caused by the size of the set, makes its usefulness low. The last processing step is used to filter the tags that are most likely to be chosen by a user. Checking the tags against the user’s personomy allows the system to choose lexical forms preferred by user (e.g., “semantics” instead of “semantic”). In addition, the personomy gives access to user specific tags (e.g., “masterthesis”). The retrieval of related tags is done analogously to the lexicon based approach used in the second step. The strength of relation is calculated based on Eq. 1; however, now the set of resources is limited to user’s own posts. It is important to notice that this approach gives access not only to tags that are explicitly found in the personomy, but also to tags that co-occurred with them in user’s posts. Subsequently, the scores are multiplied by the relation score of the base tag, which was calculated in the second step. Again the scores are calculated for each base tag separately and then the lists of results are joined, summing scores of multiple occurrences of the same tag in probabilistic way. The list of tags proposed as a recommendation is limited to the ten tags with the highest score.

As mentioned in section 3.2, the objective of some tags is to describe the resource, not to relate it to user’s interests. To give the user access to recommendation of such tags, the system recommends also the tags retrieved from the title in the first step. As scores defined in first and third step are not comparable, I decided to normalize the scores in both lists, to make the sum of scores in each list equal to one. After normalization the lists are joined, again using the probabilistic sum and limiting the final list to ten tags.

5 Evaluation

This section presents the results of the off-line system evaluation based on the available BibSonomy snapshot. The used evaluation approach assumed that all and only relevant tags were given by the user. Although this method simplifies the problem it is robust and objective. The used quality metrics were recall and precision, commonly used in recommender system evaluations [4].

Methodology The commonly used evaluation approach is to keep strict division between training and testing set. This approach was used by the organizers of the ECML PKDD discovery challenge 2008. It allowed the organizers to keep the list of correct tags in secret during the contest. However, assuming that a user provides all and only relevant tags in a post, tag recommendation becomes a specific problem in which the complete feedback about the quality of recommendation is entered to the system with each post. In such case, we should consider incremental way of evaluation in which each tested post trains the system with tags provided by the user. The paper presents both evaluation approaches. The first experiment followed strictly the approach proposed by the organizers of the ECML PKDD discovery challenge 2008 – 59,542 newest posts were used as test set. In the second experiment, in addition to incremental training, I decided

to reduce the impact of posts imported from an external repository (e.g., web browser), by not testing the system on groups of user’s posts with the same timestamp. This limited the number of test posts to 7,133. Imported posts have their tags assigned automatically. In real use a tag recommender is not used for the imported tags, therefore it should not be tested by them.

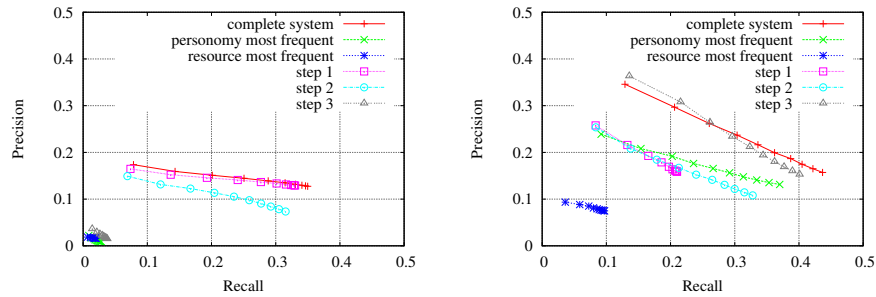
To give more insights about the system its final recommendation is presented together with tags produced in each of its three steps. The first step, that simply proposes words from the title as tags, can be considered as a baseline system. Additional baseline systems presented are the recommenders that proposes most frequent tags from user’s personomy and resource tags. As each approach returns ranked list of tags it is possible to freely limit the number of recommended tags. The plots (Fig. 4) present consecutive results for the top n tags, where $1 \leq n \leq 10$.

Results The first experiment shows low quality of personomy based recommendation, represented by the third step of the system and the baseline system which proposes the most frequent user’s tags (Fig. 4(a)). This unexpected situation is caused by the evaluation approach used in this experiment (strict division between training and testing set). Among 59,542 tested posts only 16,169 (27%) were entered by users who have their previous posts in the training set. For the rest of tested posts the personomy based recommenders could not propose any tags. Clearly such large percentage of “first-time” users is not possible in reality, thus this evaluation approach seems to underestimate the score of personomy based recommenders. The results are also strongly biased by the choice of test posts. Especially not representative is a single user that is responsible for 65% of all test posts. His/Her posts are likely to be imported from an external repository. The tags in these posts look like being mechanically extracted from the article content, which supports the recall result of the title based recommender (first step of the system). The overall result of the system is therefore completely determined by the tags proposed in the first step, that was not meant to be the main element of the system.

The second evaluation approach, in which tested posts were used to train the system, solves the problem of extraordinarily large number of “first-time” users (3% of test posts). For this evaluation method personomy based recommender (the third step) outperforms title based solutions (Fig. 4(b)). Low and slowly decreasing (with increasing number of recommended tags) precision of the baseline approach shows that most frequent tags from personomy are not necessary a good recommendation. These results confirm previous experiments, which showed that personomy is the richest, but noisiest source of tags. The title also confirmed its usefulness as a source of tags. At some point increasing the number of recommended tags does not improve precision and recall of the first step – the number of words in the title hardly ever reaches 10. The results of the second step are consistent with the first step for the top tags. Tags tend to have high self-relation score which makes title base tags likely to be high in the ranking produced by the second step. The results of the last baseline system,

the most frequent tags from the resource, confirmed that data sparsity greatly reduces the usefulness of recommenders based on resource tags.

Considering only the top tags, the third step of the system seems to be more precise than the overall system. It shows that the ranking method used to join the results of first and third step should be improved. The advantage of joining the results of these two steps is visible in the total recall of the overall system (0.44 compared to 0.4 for the third step).



(a) Evaluation approach proposed by the organizers of the ECML PKDD discovery challenge 2008. (b) Training set is iteratively enriched by tested posts. Posts classified as imported are not tested.

Fig. 4. Recall and precision of tag recommendation system compared to results of its three steps, and two baseline systems: most frequent personomy and resource tags.

6 Conclusions and future work

The key conclusion of the presented experiments is that, when recommending tags in folksonomies oriented towards individual users we should not rely only on tags previously attached to given resource. Sparsity of data and individuality of users greatly reduce their usefulness. Looking for potential tags we should focus on the direct surrounding of the post, as in this type of folksonomies a collaborative filtering approach may be deceiving. The presented tag recommendation system tries to follow these directions starting with tags from the most robust source – resource title, and expanding them by the richest source – user’s personomy.

The introduced three steps of tag recommender system can be used as a basis of more sophisticated approaches. The element that potentially is the most promising area of improvements is the folksonomy-based lexicon. In my future research I plan to experiment with mining the relations between tags using techniques from Data Mining (e.g., association rule mining) and Information Retrieval (e.g., PageRank algorithm). Specific characteristics of the dataset used

in system evaluation decrease its reliability. To confirm the results I plan to repeat the experiments on different folksonomies (e.g., del.icio.us). However, large variance of results of two evaluation approaches points to the need for a unified evaluation method that is representative of the real applications of tag recommenders, before new experiments are made.

References

1. Rakesh Agrawal, Tomasz Imieliński, and Arun Swami. Mining association rules between sets of items in large databases. *SIGMOD Rec.*, 22(2):207–216, 1993.
2. Ralitsa Angelova, Marek Lipczak, Evangelos Milios, and Pawel Pralat. Characterizing a social bookmarking and tagging network. In *Proc. of the ECAI 2008 Workshop on Mining Social Data (MSoDa)*, pages 21–26, 2008.
3. Ciro Cattuto, Christoph Schmitz, Andrea Baldassarri, Vito D. P. Servedio, Vittorio Loreto, Andreas Hotho, Miranda Grahl, and Gerd Stumme. Network properties of folksonomies. *AI Commun.*, 20(4):245–262, 2007.
4. Jonathan L. Herlocker, Joseph A. Konstan, Loren G. Terveen, and John T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst.*, 22(1):5–53, 2004.
5. Andreas Hotho, Robert Jäschke, Christoph Schmitz, and Gerd Stumme. BibSonomy: A social bookmark and publication sharing system. In *Proc. the First Conceptual Structures Tool Interoperability Workshop at the 14th Int. Conf. on Conceptual Structures*, pages 87–102, Aalborg, 2006. Aalborg Universitetsforlag.
6. Andreas Hotho, Robert Jäschke, Christoph Schmitz, and Gerd Stumme. Trend detection in folksonomies. In *Proc. First International Conference on Semantics And Digital Media Technology (SAMT)*, volume 4306 of *LNCS*, pages 56–70, Heidelberg, dec 2006. Springer.
7. Robert Jäschke, Leandro Balby Marinho, Andreas Hotho, Lars Schmidt-Thieme, and Gerd Stumme. Tag recommendations in folksonomies. In *Knowledge Discovery in Databases: PKDD 2007, 11th European Conference on Principles and Practice of Knowledge Discovery in Databases, Warsaw, Poland, September 17-21, 2007, Proceedings*, volume 4702 of *LNCS*, pages 506–514. Springer, 2007.
8. Beate Krause, Robert Jäschke, Andreas Hotho, and Gerd Stumme. Logsonomy - social information retrieval with logdata. In *HT '08: Proc. the 19th ACM conf. on Hypertext and hypermedia*, pages 157–166, New York, NY, USA, 2008. ACM.
9. Sigma On Kee Lee and Andy Hon Wai Chun. Automatic tag recommendation for the web 2.0 blogosphere using collaborative tagging and hybrid ANN semantic structures. In *ACOS'07: Proc. the 6th Conf. on WSEAS Int. Conf. on Applied Computer Science*, pages 88–93, Stevens Point, Wisconsin, USA, 2007. WSEAS.
10. Christoph Schmitz, Andreas Hotho, Robert Jäschke, and Gerd Stumme. Mining association rules in folksonomies. In *Data Science and Classification (Proc. IFCS 2006 Conference)*, pages 261–270, Ljubljana, July 2006. Springer.
11. Börkur Sigurbjörnsson and Roelof van Zwol. Flickr tag recommendation based on collective knowledge. In *WWW '08: Proc. the 17th international conference on World Wide Web*, pages 327–336, New York, NY, USA, 2008. ACM.
12. S.C. Sood, K.J. Hammond, S.H. Owsley, and L. Birnbaum. TagAssist: Automatic tag suggestion for blog posts. In *Proc. the International Conference on Weblogs and Social Media (ICWSM 2007)*, 2007.