Tag-aware Recommender Systems by Fusion of Collaborative Filtering Algorithms

Karen H. L. Tso-Sutter, Leandro Balby Marinho * and Lars Schmidt-Thieme
Information Systems and Machine Learning Lab (ISMLL)
University of Hildesheim
Samelsonplatz 1, 31141 Hildesheim, Germany
{tso,marinho,schmidt-thieme}@ismll.uni-hildesheim.de

ABSTRACT

Recommender Systems (RS) aim at predicting items or ratings of items that the user are interested in. Collaborative Filtering (CF) algorithms such as user- and item-based methods are the dominant techniques applied in RS algorithms. To improve recommendation quality, metadata such as content information of items has typically been used as additional knowledge. With the increasing popularity of the collaborative tagging systems, tags could be interesting and useful information to enhance RS algorithms. Unlike attributes which are "global" descriptions of items, tags are "local" descriptions of items given by the users. To the best of our knowledge, there hasn't been any prior study on tag-aware RS. In this paper, we propose a generic method that allows tags to be incorporated to standard CF algorithms, by reducing the three-dimensional correlations to three two-dimensional correlations and then applying a fusion method to re-associate these correlations. Additionally, we investigate the effect of incorporating tags information to different CF algorithms. Empirical evaluations on three CF algorithms with real-life data set demonstrate that incorporating tags to our proposed approach provides promising and significant results.

Categories and Subject Descriptors

H.3 [**Information Storage and Retrieval**]: Information Search and Retrieval - *Information Filtering*

General Terms

Algorithms, Performance, Experimentation

Keywords

Recommender Systems, Collaborative Filtering, Tags

*Brazilian National Council Scientific and Technological Research (CNPq) scholarship holder.

SAC'08 March 16-20, 2008, Fortaleza, Ceará, Brazil

1. INTRODUCTION

Collaborative tagging systems, also known as folksonomies are web-based systems that allow users to upload their resources, and to label them with arbitrary words, so-called *tags*. These systems are becoming more common among web users now-a-day. For example popular web services such as Flickr [1], del.icio.us [2] , Last.fm [3], Gmail [4], etc, provide possibility for users to tag or label an item of interest. In general, tagging is associated to the Web 2.0 and is becoming the new trend enabling people to easily add metadata to content. Hence, These additional metadata can be used to improve search mechanisms, better structure the data for browsing or provide personalized recommendations fitting the users' interests. Content information used in attribute-aware RS algorithms is typically attached to the items and is usually provided by domain experts. Therefore, an item always has the same attributes among all users. On the other hand, tags are provided by various users. Thus, tags are not only associated to the items but also to the users. Although attributes and tags are both metadata and could act as additional background knowledge to improve RS algorithms, they should be handled differently.

Despite the considerable amount of researches done in attribute-aware RS algorithms, the specific problem of integrating tags to RS algorithms is rarely explored. Most existing works on RS with tags are limited to recommending tags, i.e. assisting users for annotation purposes, while using tags as supplementary source for recommending items has never been investigated.

In this paper, we propose to integrate tags in recommender systems by first extending the user-item matrix and then applying an algorithm that fuses two popular RS algorithms such that the correlations between users, items and tags can be captured simultaneously. Our contributions are as follow: i) propose a generic method that allows tags to be incorporated to standard CF algorithms, ii) propose an adapted fusion mechanism to capture the 3-dimensional correlations between users, items and tags iii) conduct empirical evaluations on three CF algorithms with real-life data set and investigate the effect of incorporating tags information to those algorithms.

[1]http://flickr.com/
[2]http://del.icio.us/
[3]http://www.last.fm/
[4]http://mail.google.com/

2. RELATED WORK

There have already been a reasonable amount of researches in using attributes as background knowledge in RS [2, 3, 4, 8, 7, 14, 12, 17, 16, 20]. However, to the best of our knowledge, there hasn't been any research in considering tags with RS algorithms to predict items.

The existing research work on using tagging information for recommendation purposes lies basically on the recommendation of tags for assisting the user in annotation related tasks [15, 22, 5, 13, 11]. The relationship between users, items and tags could be treated by the literature as multidimensional RS, where other dimensions of data besides the traditional two-dimensional user/item approach are used. Adomavicius et al. [1] have proposed a reduction from multidimensional to 2-dimensional representations where traditional RS algorithms can be directly applied. Aspects or contexts of the additional dimensions are used to restrict or contextualize the two-dimensional data to be considered, This approach is more suitable for RS that consider time related aspect, which we do not take into account in the tag-aware RS. In our case, we use all the information available in all dimensions independent of context.

3. COLLABORATIVE FILTERING

Recommender systems (RS) predict ratings of items or suggest a list of items that is unknown to the user. They take the users, items as well as the ratings of items into account. We introduce some notations that we use throughout the rest of the paper. Let

- U be a set of users,
- I be a set of items,
- $r : U \times I \to \mathbb{R}$ be a map of ratings,
- $R \in \mathbb{R}^{U \times I}$ be the user-item matrix where each value correspond to $R_{u,i} = r(u, i)$, where $u \in U$ and $i \in I$.

Two different recommendation tasks are typically considered: (i) predicting the ratings, i.e. how much a given user will like a particular item, and (ii) predicting the items, i.e. which N items a user will rate, buy or visit next (topN). Most recommender systems derive recommendations to a user by using opinions from people who have alike tastes, called *neighborhood*, while concealing the real identity of the users neighborhood.

The prevalent method in practice is Collaborative Filtering (CF) [10]. Its idea is basically the nearest neighbor method. Given some user profiles, it predicts whether a user might be interested in a certain item, based on a section of other users or items in the database. There are in general two types of collaborative filtering: user-based and item-based. Most of the time, they share the same concept except they vary by how the neighborhood is formed.

In user-based CF [19], recommendations are generated by considering solely the ratings of users on items, by computing the pairwise similarities between users, e.g., by means of vector similarity:

$$\text{usim}^{\text{ratings}}(u, v) := \frac{\langle R_{u,.}, R_{v,.} \rangle}{||R_{u,.}||_2 ||R_{v,.}||_2} \tag{1}$$

where u, $v \in U$ are two users and $R_{u,.}$ and $R_{v,.}$ are the vectors of their ratings.

A dualistic form of user-based CF is item-based CF [9], where similarities are computed between each pair of items i, $j \in I$.

$$\text{isim}^{\text{ratings}}(i, j) := \frac{\langle R'_{.,i}, R'_{.,j} \rangle}{||R'_{.,i}||_2 ||R'_{.,j}||_2} \tag{2}$$

In user-based CF, to derive the recommendations for a target user u, usually only similarities of the k most-similar users are selected (neighborhood – N_u). When predicting a rating of a given user u for an item i, the weighted sum of the other users are computed by:

$$\hat{r}(u, i) := \bar{r_u} + \frac{\sum_{v \in N_u \cap dom r_i} w(u, v)(r(v, i) - \bar{r_v})}{\sum_{v \in N_u \cap dom r_i} w(u, v)} \tag{3}$$

where dom r is the domain of r. For $i \in I$ we define $r_i := r|_{U \times \{i\}}$ and $w(u, v)$ is the similarity between each user v in user u's neighborhood and user u him/herself.

In the case of item-based CF, the prediction would be the average of the ratings of k most-similar items N_i rated by the given user u. In similar notation as user-based CF, the prediction for a rating of a given user u for an item i is:

$$\hat{r}(u, i) := \frac{\sum_{j \in N_i \cap dom r_u} w(i, j)(r(i, j))}{\sum_{j \in N_i \cap dom r_u} w(i, j)} \tag{4}$$

Due to its simplicity and rather promising performances [10, 18], collaborative filtering has been one of the most prominent methods used in recommender systems.

3.1 Extension with Tags

As collaborative tagging is getting more widely used, this information could also be employed as background knowledge in RS. There has not been any previous investigation in integrating tags information to improve recommendation quality for such purpose. Hence, we propose a generic mechanism that allows tags to be integrated to standard CF algorithms such as user- and item-based algorithms. In the following, we will describe how the extension can be applied.

Unlike attributes which only have a two-dimensional relation $< item, attribute >$, tags hold a three-dimensional relation $< user, item, tag >$. We cope with this three dimensionalities by projecting it as three two-dimensional problem, $< user, tag >$ and $< item, tag >$ and $< user, item >$. This can be done by augmenting the standard user-item matrix horizontally and vertically with user and item tags correspondingly. *User tags*, are tags that user u, uses to tag items and are viewed as items in the user-item matrix. *Item tags*, are tags that describe an item, i, by users and play the role of users in the user-item matrix (See Fig. 1). Furthermore, instead of viewing each single tag as user or item, clustering methods can be applied to the tags such that similar tags are grouped together. In this paper, no clustering method has been applied to the tags. Let:

- a set of user tags be T_u,
- a set of item tags be T_i,
- a new set of users after item tags extension be $U_{extend} = U + T_i$,
- a new set of items after user tags extension be $I_{extend} = I + T_u$,

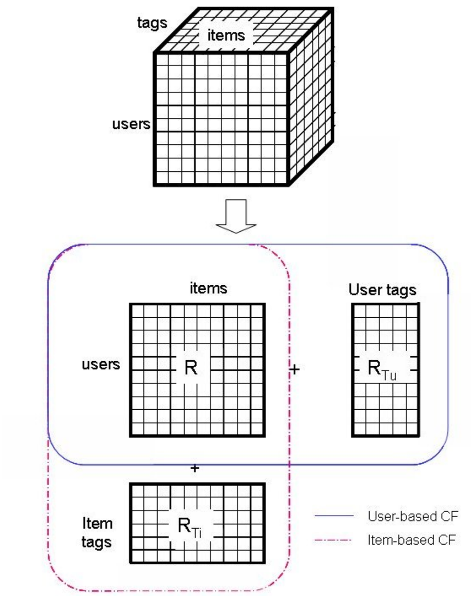

Figure 1: Extend user-item matrix by including user tags as items and item tags as users.

- $O_{u,t_u} \in \{0,1\}$ specify whether a user $u \in U$ has used the tag t_u, where $t_u \in T_u$,
- $O_{i,t_i} \in \{0,1\}$, specify whether an item $i \in I$ is described by the tag t_i, where $t_i \in T_i$,
- R_{T_u} be represented in a $U \times |T_u|$ user-tag matrix, where each value of the matrix corresponds to O_{u,t_u},
- R_{T_i} be represented in a $|T_i| \times I$ tag-item matrix, where each value of the matrix corresponds to O_{i,t_i}.

To apply user- and item-based CF after the extension with tags, both CF algorithms have to be recomputed with the newly extended user-item matrix. For user-based CF, the new user-item matrix , $R_{uextend} := R + R_{T_u}$, is represented in a $U \times I_{extend}$ matrix. In the case of item-based CF, the new user-item matrix, $R_{iextend} := R + R_{T_i}$, is represented in a $U_{extend} \times I$ matrix.

3.2 Fusing User-based and Item-based

Tag information has a slightly different nature than attribute information. In general, attributes are only attached to items and attributes of an item appear the same to all users "globally". Tags of an item, on the other hand, are descriptions of the item by one or more than one users. Thus, tags are not only attached to the item itself but also are depended on the user's preference. They are "local" descriptions of an item that might change from users to users. This suggests that a RS algorithm that is able to capture both user's and item's aspect of tags would eventually be a suitable choice.

We have selected an existing algorithm developed by Wang et al. [21], which fuses the predictions of user- and item-based CF. In general, their idea is to correspond CF to the estimation of conditional probability problem. Their paper described two types of combination: i) predicting the rating by fusing user- and item-based predictions; ii) in addition to i), it also used the similar items ratings generated by similar users. In this paper, we only consider the first type of combination as our initial experiments have shown that the second type did not provide better recommendations in the predicting item problem. The fusion of the user- and item-based predictions was done by computing the sum of the two conditional probabilities that are based on user- and item-based similarities, which are computed using standard user- and item-based CF. A parameter, λ, is introduced to adjust the significance of the two predictions.

$$P(r_{u,i}|w(u,v))\lambda + P(r_{u,i}|w(i,z))(1-\lambda).$$

In this paper, we refer the combination of user- and item-based CF as *fusion*.

3.2.1 Fusion for Predicting Item Problem

Most systems that use collaborative tagging do not contain rating information, i.e. only the occurrence, $O_{u,i} \in \{0, 1\}$, whether item $i \in I$ occurred with user $u \in U$ (e.g. u has bought/viewed item i). For example, Last.fm is a popular internet radio and music community website which allows user to tag the music. However, it does not support users to provide explicit ratings of the music. A predicting item problem would be more suitable for tag-aware RS. Hence, we focus on the predicting item problem in this paper. Yet, the fusion method by Wang et al. [21] does not consider tags. Also, their algorithm is only suitable for predicting rating problem and not the predicting item problem. Thus, we propose a fusion algorithm that tackles the predicting item problem and also takes tags into account.

Again, fusion for the predicting item problem is done by combining the predictions of user and item-based; however, this time, the predictions are computed differently. For the predicting item problem in user-based CF, recommendations are a list of items that is ranked by decreasing frequency of occurrence in the ratings of his/her neighbors.

$$p^{\text{ucf}}(O_{u,i}=1) := \frac{|\{v \in N_u \,|\, O_{v,i}=1\}|}{|N_u|} \tag{5}$$

For item-based CF, the *topN* recommendation suggested by [9] is to compute a list of items that is ranked by decreasing sum of the similarities of neighboring items, N_i, which have been rated by user u.

$$p^{\text{icf}}(O_{u,i}=1) := \sum_{j \in N_i \cap O_{u,j}=1} w(i,j) \tag{6}$$

As the values of the prediction lists computed by user- and item-based have different units (user-based being the frequency of items and item-based the similarity of items). To address the different meanings of the values, we normalized the prediction lists to unity. Note that other ways of normalization techniques can be explored but for simplicity we have chosen to normalize to unity.

The *topN* combined prediction list is thus:

$$\begin{aligned} p^{\text{iucf}}(O_{u,i}=1) := {} & \lambda \frac{p^{ucf}(O_{u,i}=1)}{\sum_i p^{ucf}(O_{u,i}=1)} \\ & +(1-\lambda)\frac{p^{icf}(O_{u,i}=1)}{\sum_i p^{icf}(O_{u,i}=1)} \end{aligned} \tag{7}$$

$$\arg\max_{i}^{N} p^{\text{iucf}}(O_{u,i} = 1) \tag{8}$$

Fusion with Tags
As user- and item-based CF are the basic components of the fusion method, when tag information is available, fusion can simply be extended with tags by applying the tag extension method described in Section 3.1.

We deem that this tag-fusion algorithm is a suitable tag-aware RS algorithm because i) user tags and item tags provide extra indications of the user's and item's preferences, ii) our adapted fusion approach then brings about both user and item aspects of the tags concurrently. In fact, our empirical analysis has shown that our tag-fusion RS approach provides promising results.

4. EXPERIMENTS

4.1 Data Set

The data set we have used in our experiments is the Last.fm data. Last.fm is an internet radio and music community website and is one of the world's largest social music platforms. It collects the profile of each user's musical preference by keeping track of the songs that the user listens to. The reason we have chosen this data set is because it is rich in users preference information – songs the users listen to, as well as tags information – user-end tagging of artists, albums, and tracks of the music. The data was gathered in July 2006 by crawling the Last.fm site. Here the items are artist names, which are already normalized by the system. Many recommendation algorithms suffer from sparse data or the "long tail" of items which were used by only few users. Hence, to eliminate some of the noise and improve the chances of good results for all algorithms we will adopt the pruning procedure described in [11] and restrict the evaluation to the "dense" part of the tripartite graph. In this case we considered only the users, resources and tags that appear at least 10 times in the tag assignments, i.e. (user, resource, tag) triples. This gives 1853 items (artists), 2917 users and 2045 tags.

4.2 Evaluation Protocol and Setting

We have used a leave-one-out protocol to evaluate the obtained recommendations. Hence, obtaining the test set by randomly select one listened music from every user. The rest of the data is used as the training set, where the model is trained on this data set. This protocol has been called *All-But1* in [6]. We have performed 10-fold cross validation. In addition, we have further split the training data to validation data to optimize the parameters λ and k, the neighborhood size. We have varied the lambda from 0 to 1 by an interval of 0.1 and the neighborhood from 10-150 by an interval of 10. Using the validation data, we have found the best λ to be 0.4 and k to be 20. We have then retrained the model using the training data with the optimized lambda.

Our paper focuses on the item prediction problem, which is to predict a fixed number of *topN* recommendations and not the ratings. Suitable evaluation metrics are Precision, Recall and F1. Similar to Sarwar et al. [19], our evaluations consider any item in the recommendation set that matches any item in the testing set as a "hit". The number of *topN* items to be predicted is 10.

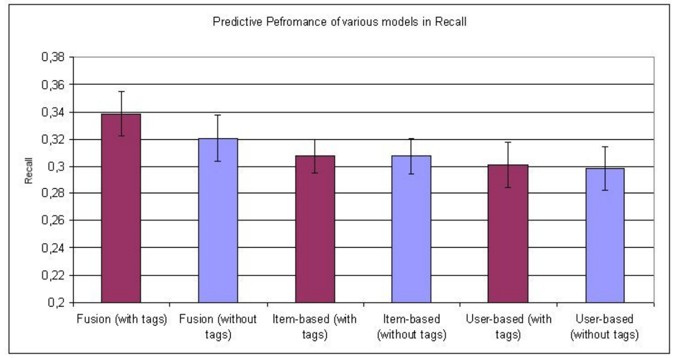

Figure 2: Results of comparing baseline models and the unification model using the optimized lambda and neighborhood.

Methods in Comparison
We have implemented two baseline models: user-based CF [19] and item-based CF [9] as well as the fusion of user- and item-based CF described in Section 3.2.1. In addition, to evaluate the impact of the presence of tags in CF algorithms, we have compared the evaluations with the incorporation of tags and without tags to each algorithm. Hence, each algorithm is first evaluated with the standard approach and then evaluated with the extension of tags using the method described in Section 3.1. Note that all algorithms are implemented for the predicting item problem and all CF algorithms are computed by the vector similarity measures.

4.3 Experiment Results

Figure 2 summarizes the results, including a comparison with the baseline models: standard user- and item-based CF as well as the fusion model, each *with* and *without* the presence of tags. It can be seen that the fusion method, both with and without tags, significantly outperform the standard CF models. Furthermore, the results show that after the introduction of tags to the fusion method, there is a significant increase in the performance. It is interesting to see that incorporating tags to the baseline models does not improve the recommendation quality at all, in contrast to the promising results of including tags in the fusion method. As mentioned in previous section, tags hold the 3-dimensional relationship between users, items and tags. The use of tags has shown to increase the interconnectivity amongst users and items. Applying user/item tags alone does not exploit the characteristic of tags correctly. Hence, attaching tags to standard CF algorithms, such as user-based CF, does not improve the performance at all, the tags are then only seen as noise. This reflects that by simply extending the standard CF algorithms with tags, it fails to denote the 3-dimensional correlations between user, item and tag, whereas the proposed fusion method has shown to be able to capture this relationship.

5. CONCLUSIONS

We have conducted a novel study in the utilization of tags as supplementary source to predict item recommendations. Here, we have presented a generic method to include tags to standard CF algorithms such as user- and item-based CF. In addition, we have found an approach that deals with the 3-dimensional correlation between the users, items and tags by first applying our tag extension mechanism and then an fusion method which we have adapted from a predicting rating problem to predicting item problem. Our empirical analysis has shown that the proposed adapted fusion method outperforms standard baseline models, especially with the incorporation of tags. Moreover, our findings have suggested that our adapted fusion method has successfully captured the relationships between users, items and tags. Although our investigation has provided promising results, we believe that our contribution is an initial step in the study of tag-aware RS, additional research in this field is still to be explored.

6. REFERENCES

[1] G. Adomavicius, R. Sankaranarayanan, S. Sen, and A. Tuzhilin. Incorporating contextual information in recommender systems using a multidimensional approach. *ACM Trans. Inf. Syst.*, 23(1):103–145, 2005.

[2] M. Balabanovic and Y. Shoham. Fab: Content-based, collaborative recommendation. *Commun. ACM*, 40:66–72, 1997.

[3] J. Basilico and T. Hofmann. Unifying collaborative and content-based filtering. In *Proceedings of the twenty-first international conference on Machine learning*, pages 65–72, New York, NY, USA, 2004. ACM Press.

[4] C. Basu, H. Hirsh, and W. Cohen. Recommendation as classification: using social and content-based information in recommendation. In *AAAI '98/IAAI '98: Proceedings of the fifteenth national/tenth conference on Artificial intelligence/Innovative applications of artificial intelligence*, pages 714–720, Menlo Park, CA, USA, 1998. American Association for Artificial Intelligence.

[5] D. Benz, K. Tso, and L. Schmidt-Thieme. Automatic bookmark classification: A collaborative approach. In *Proceedings of the Second Workshop on Innovations in Web Infrastructure (IWI 2006), Edinburgh, Scotland*, 2006.

[6] J. S. Breese, D. Heckerman, and C. Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (UAI-98). G. F. Cooper, and S. Moral, Eds. Morgan-Kaufmann, San Francisco, Calif*, pages 43–52, 1998.

[7] R. Burke. Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction*, 12:331–370, 2002.

[8] M. Claypool, A. Gokhale, T. Miranda, P. Murnikov, D. Netes, and M. Sartin. Combining content-based and collaborative filters in an online newspaper. In *Proceedings of ACM SIGIR Workshop on Recommender Systems*, August 1999.

[9] M. Deshpande and G. Karypis. Item-based top-n recommendation algorithms. *ACM Transactions on Information Systems. Springer-Verlag*, 22/1, 2004.

[10] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry. Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, 35(12):61–70, 1992.

[11] R. Jaeschke, L. Marinho, A. Hotho, L. Schmidt-Thieme, and G. Stumme. Tag recommendations in folksonomies. In *Proceedings of the 11th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD), Warsaw, Poland (to appear)*, 2007.

[12] Q. Lin and B. M. Kim. An approach for combining content-based and collaborative filters. In *Proceedings of the Sixth international workshop on Information retrieval with Asian languages (ACL-2003)*, pages 17–24, 2003.

[13] L. B. Marinho and L. Schmidt-Thieme. Collaborative tag recommendations. In *Proceedings of 31st Annual Conference of the Gesellschaft fr Klassifikation (GfKl), Freiburg (to appear)*. Springer, 2007.

[14] P. Melville, R. Mooney, and R. Nagarajan. Content-boosted collaborative filtering. In *Proceedings of Eighteenth National Conference on Artificial Intelligence (AAAI-2002)*, pages 187–192, 2001.

[15] G. Mishne. Autotag: a collaborative approach to automated tag assignment for weblog posts. In *WWW '06: Proceedings of the 15th international conference on World Wide Web*, pages 953–954, New York, NY, USA, 2006. ACM Press.

[16] M. J. Pazzani. A framework for collaborative, content-based and demographic filtering. *Artif. Intell. Rev.*, 13(5-6):393–408, 1999.

[17] A. Popescul, L. Ungar, D. Pennock, and S. Lawrence. Probabilistic models for unified collaborative and content-based recommendation in sparse-data environments. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 437–444, 2001.

[18] P. Resnick, N. Iacovou, M. Suchak, P. Bergstorm, and J. Riedl. GroupLens: an open architecture for collaborative filtering of netnews. In *Proceedings of the ACM1994 Conference on Computer Supported Cooperative Work*, pages 175–186, Chapel Hill, North Carolina, 1994. ACM.

[19] B. M. Sarwar, G. Karypis, J. A. Konstan, and J. Riedl. Analysis of recommendation algorithms for e-commerce. In *Proceedings of the Second ACM Conference on Electronic Commerce (EC'00)*, pages 285–295, 2000.

[20] K. Tso and L. Schmidt-Thieme. Evaluation of attribute-aware recommender system algorithms on data with varying characteristics. In *Proceedings of the 10th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD 2006)*, pages 831–840. Springer, 2006.

[21] J. Wang, A. P. de Vries, and M. J. T. Reinders. Unifying user-based and item-based collaborative filtering approaches by similarity fusion. In *SIGIR '06: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 501–508, New York, NY, USA, 2006. ACM Press.

[22] Z. Xu, Y. Fu, J. Mao, and D. Su. Towards the semantic web: Collaborative tag suggestions. In *Proceedings of the Collaborative Web Tagging Workshop at the WWW 2006*, Edinburgh, Scotland, 2006.