

Tag Recommendation using Probabilistic Topic Models

Ralf Krestel and Peter Fankhauser

L3S Research Center
Leibniz Universität
Hannover, Germany

Abstract. Tagging systems have become major infrastructures on the Web. They allow users to create tags that annotate and categorize content and share them with other users, very helpful in particular for searching multimedia content. However, as tagging is not constrained by a controlled vocabulary and annotation guidelines, tags tend to be noisy and sparse. Especially new resources annotated by only a few users have often rather idiosyncratic tags that do not reflect a common perspective useful for search. In this paper we introduce an approach based on Latent Dirichlet Allocation (LDA) for recommending tags of resources. Resources annotated by many users and thus equipped with a fairly stable and complete tag set are used to elicit latent topics represented as a mixture of description tokens and tags. Based on this, new resources are mapped to latent topics based on their content in order to recommend the most likely tags from the latent topics. We evaluate recall and precision for the bibsonomy benchmark provided within the ECML PKDD Discovery Challenge 2009.

1 Introduction

Tagging systems [1] like Flickr, Last.fm, Delicious or Bibsonomy have become major infrastructures on the Web. These systems allow users to create and manage tags to annotate and categorize content. In *social* tagging systems like Delicious the user can not only annotate his own content but also content of others. The service offered by these systems is twofold: They allow users to publish content and to search for content, thus *tagging* also serves two purposes for the user:

1. Tags help to organize and manage own content, and
2. Find relevant content shared by other users.

Tag recommendation can focus on one of the two aspects. Personalized tag recommendation helps individual users to annotate their content in order to manage and retrieve their own resources. Collective tag recommendation aims at making resources more visible to other users by recommending tags that facilitate browsing and search.

However, since tags are not restricted to a certain vocabulary, users can pick any tags they like to describe resources. Thus, these tags can be inconsistent and idiosyncratic, both due to users’ personal terminology as well as due to the different purposes tags fulfill [2]. This reduces the usefulness of tags in particular for resources annotated by only a few users (aka cold start problem in tagging), whereas for popular resources collaborative tagging typically saturates at some point, i.e., the rate of new descriptive tags quickly decreases with the number of users annotating a resource [3].

The main goal of the approach presented in this paper is to overcome the cold start problem for tagging new resources. To this end, we use Latent Dirichlet Allocation (LDA) to elicit latent topics from resources with a fairly stable and complete tag set. The latent topics are represented as a mixture of description tokens like URL, title, and other metadata, and tags, which typically co-occur. Based on this, new resources are mapped to latent topics based on their description in order to recommend the most likely tags from the latent topics.

The remainder of this paper is organized as follows. In Section 2, we define the problem of tag recommendation more formally, and introduce the approach based on LDA. In Section 3 we present our evaluation results. In Section 4 we discuss related work, and in Section 5 we summarize and outline possible future research directions.

2 Tag Recommendation

2.1 Problem Definition

Given a set of resources R , tags T , and users U , the ternary relation $X \subseteq R \times T \times U$ represents the user specific assignment of tags to resources. T consists of two disjoint sets T_{tag} and T_{desc} . T_{tag} contains all user assigned tags, T_{desc} contains the vocabulary of content and meta information, such as abstract or resource description, which is represented as tag assignment by a special “user”. A post $b(r_i, u_j)$ for resource $r_i \in R$ and a user $u_j \in U$ comprises all tags assigned by u_j to r_i : $b(r_i, u_j) = \pi_t \sigma_{r_i, u_j} X$ ¹. The goal of collective tag recommendation is to suggest tags to a user u_j for a resource r_i based on tag assignments to other resources by other users collected in $Y = \sigma_{r \neq r_i \vee u \neq u_j} \pi_{r, t} X \subseteq R \times T$.

2.2 Latent Dirichlet Allocation

The general idea of Latent Dirichlet Allocation (LDA) is based on the hypothesis that a person writing a document has certain topics in mind. To write about a topic then means to pick a word with a certain probability from the pool of words of that topic. A whole document can then be represented as a mixture of different topics. When the author of a document is one person, these topics reflect the person’s view of a document and her particular vocabulary. In the context of tagging systems where multiple users are annotating resources, the

¹ projection π and selection σ operate on multisets without removing duplicate tuples

resulting topics reflect a collaborative shared view of the document and the tags of the topics reflect a common vocabulary to describe the document.

More generally, LDA helps to explain the similarity of data by grouping features of this data into unobserved sets. A mixture of these sets then constitutes the observable data. The method was first introduced by Blei, et. al. [4] and applied to solve various tasks including topic identification [5], entity resolution [6], and Web spam classification [7].

The modeling process of LDA can be described as finding a mixture of topics for each resource, i.e., $P(z | d)$, with each topic described by terms following another probability distribution, i.e., $P(t | z)$. This can be formalized as

$$P(t_i | d) = \sum_{j=1}^N P(t_i | z_i = j) P(z_i = j | d), \quad (1)$$

where $P(t_i)$ is the probability of the i th term for a given document and z_i is the latent topic. $P(t_i | z_i = j)$ is the probability of t_i within topic j . $P(z_i = j)$ is the probability of picking a term from topic j in the document. These probability distributions are specified by LDA using Dirichlet distributions. The number of latent topics N has to be defined in advance and allows to adjust the degree of specialization of the latent topics. The algorithm has to estimate the parameters of an LDA model from an unlabeled corpus of documents given the two Dirichlet priors and a fixed number of topics. Gibbs sampling [5] is one possible approach to this end: It iterates multiple times over each tag t , and samples a new topic j for the tag based on the probability $P(z_i = j | t, z_{-i})$, where z_{-i} represents all topic-word and document-topic assignments except the current assignment z_i for tag t , until the LDA model parameters converge.

Application to Tagging Systems LDA assigns to each document latent topics together with a probability value that each topic contributes to the overall document. For tagging systems the documents are resources $r \in R$, and each resource in addition to its description from T_{desc} is described by tags $t \in T_{tag}$ assigned by users $u \in U$. Instead of documents composed of terms, we have resources composed of tags. To build an LDA model we need resources and associated tags previously assigned by users. For each resource r we need some posts $b(r, u_i)$ assigned by users $u_i, i \in \{1 \dots n\}$. Note that for each resource, at least the tag assignments from its description is available. Then we can represent each resource in the system not with its actual tags but with the tags from topics discovered by LDA.

For a new resource r_{new} with few or no posts, we can expand the latent topic representation of this resource with the top tags of each latent topic. To accommodate the fact of some tags being added by multiple users whereas others are only added by one or two users we can use the probabilities that LDA assigns. As formalized in Equation 1 this is a two level process. Probabilities are assigned not only to the latent topics for a single resource but also to each tag within a latent topic to indicate the probability of this tag being part of that particular

Table 1. Top terms composing the latent topic “images” and “tutorial”

Tag	Count	Prob.	Tag	Count	Prob.
images(tag)	243	0.064	tutorial(tag)	640	0.185
photo(tag)	218	0.057	howto(tag)	484	0.140
photography(tag)	205	0.054	tutorial(desc)	204	0.059
image(tag)	188	0.049	tutorials(tag)	184	0.053
photos(tag)	164	0.043	tutorials(desc)	173	0.050
photo(desc)	138	0.036	tips(tag)	126	0.037
images(desc)	106	0.028	reference(tag)	118	0.034
photos(desc)	98	0.026	guide(tag)	79	0.023
flickr(tag)	93	0.024	lessons(tag)	50	0.014
pictures(desc)	61	0.016	tips(desc)	48	0.014
graphics(tag)	49	0.013	wschools(desc)	45	0.013
media(tag)	48	0.013	tutoriel(tag)	33	0.010
art(tag)	48	0.013	comment(tag)	29	0.008

topic. We represent each resource r_i as the probabilities $P(z_j|r_i)$ for each latent topic $z_j \in Z$. Every topic z_j is represented as the probabilities $P(t_n|z_j)$ for each tag $t_n \in T$. By combining these two probabilities for each tag for r_{new} , we get a probability value for each tag that can be interpreted similarly as the tag frequency of a resource. Setting a threshold allows to adjust the number of recommended tags and emphasis can be shifted from recall to precision.

Imagine a resource with the following tags: “photo”, “photography”, and “howto”. Table 1 shows the top terms for two topics related with the assigned tags. The latent topics comprise a broad notion of (digital) photography and the various aspects of tutorial material. Given these topics we can easily extend the current tag set or recommend new tags to users by looking at the latent topics. If LDA assumes that our resource in question belongs to 66% to the “photo”-topic and to 33% to the “howto”-topic, these probabilities are multiplied with the individual topic/tag probabilities, and the top five tags recommended are “tutorial”, “howto”, “images”, “photo”, and “photography”.

3 Evaluation

We used the data provided by the ECML PKDD Discovery Challenge 2009 to evaluate our approach and fine-tune our parameters. For assessing precision, recall, and f-measure we used the supplied evaluation script.

3.1 Dataset

Our dataset consists of the provided training data for the Discovery Challenge. All experiments were performed on the post-core at Level 2, where all tags, users, and resources occur at least in two posts. To measure the performance of our system, we split the training data into a 90% training set and a 10% test set based on posts (called content IDs in the dataset).

Table 2. Fields parsed to represent a resource

Bibtex		Bookmark
Author	Title	URL
Editor	Description	Description
Booktitle	Journal	Extended
Abstract		

Table 3. Actual tags and recommended tags with computed probability for URL <http://jo.irisson.free.fr/bstdatabase/>

Real Tag	LDA Tag	LDA Prob.
latex	bibtex(tag)	0.017
bibtex	latex(tag)	0.017
bibliography	bibtex(desc)	0.014
database	latex(desc)	0.008
engine	theory(desc)	0.005
style	citeulike(desc)	0.005
tex	bibliography(tag)	0.004
reference	database(tag)	0.003
academic	styles(desc)	0.003

For each resource, as defined by the hash values, we build up a textual representation. This representation contains all the tags that were assigned by users in the training set to a particular resource. In addition, we add terms extracted from the description of the resource. More precisely, we tokenized different fields describing a bookmark or bibtex entry. An overview of the fields can be seen in Table 2. Afterwards, we removed stopwords and punctuation marks. Using also the description ensures that we have some terms related to a resource even if no other user before tagged it.

3.2 Results

The tag recommendation algorithm is implemented in Java. We used Mallet [8], which provides an efficient SparseLDA implementation [9], to perform the Latent Dirichlet Allocation with Gibbs sampling. The LDA algorithm takes three input parameters: the number of terms to represent a latent topic, the number of latent topics to represent a document, and the overall number of latent topics to be identified in the given corpus.

Table 3 shows the actual tag distribution for a randomly selected resource (<http://jo.irisson.free.fr/bstdatabase/>), the top tags recommended by LDA with aggregated probabilities, and all the tags provided by a sample user. As the actual tags indicate, the url is a database/latex related site. The tags recommended by LDA come from six latent topics, comprising latex, databases, academia, references, bibliography, and style. These tags characterize the resource quite well.

Table 4. F-measure for different number of recommended tags and different number of LDA topics compared with recommending the most frequent tags (mf)

No. Tags	# LDA topics										mf
	50	100	200	400	600	800	1000	2500	5000	10000	
1	0.170	0.191	0.214	0.229	0.229	0.230	0.229	0.238	0.240	0.235	0.270
2	0.200	0.225	0.248	0.266	0.271	0.271	0.274	0.289	0.288	0.283	0.335
3	0.209	0.233	0.257	0.277	0.282	0.285	0.287	0.302	0.303	0.300	0.362
4	0.209	0.237	0.257	0.279	0.287	0.289	0.292	0.305	0.307	0.303	0.379
5	0.209	0.238	0.258	0.280	0.286	0.291	0.293	0.307	0.307	0.304	0.388

Table 4 compares the f-measure reached for various numbers of latent topics and the baseline which simply recommends the top most frequent tags for each resource (mf)². As can be seen, the best f-measure for LDA is reached between 2500 and 5000 latent topics, but it does not reach the baseline by far. The main reason for this seems to be that the average number of tags per resource is just 10.3 (7.4 distinct tags). This is significantly smaller than the number of (distinct) tokens in a full-text abstract or document, to which LDA has been applied traditionally. Moreover, there are only about 2.8 posts per resource. Thus, there is on the one hand too little co-occurrence evidence for eliciting latent topics, on the other hand there is too little overlap between users on a resource to effectively predict tags via the latent topics of a resource for a new post.

However, to deal with resources that have only few tags associated it makes sense to combine tag recommendations based on most frequent tags with tag recommendations based on latent topics. With $freq(t, r)$ the frequency of tag t annotated for resource r , one estimate of the probability of tag t given resource r is as follows:

$$P_1(t | r) = \frac{freq(t, r)}{\sum_{t_i \in r} freq(t_i, r)} \quad (2)$$

This estimate can be combined with the estimate $P_2(t | r)$ via latent topics in Equation 1 by means of a mixture:

$$P(t | r) = \lambda P_1(t | r) + (1 - \lambda) P_2(t | r). \quad (3)$$

Table 5 shows that this combination achieves consistently better recall and precision than the individual approaches. The largest gain is achieved for the first recommended tag. Similar accuracies are achieved when varying the mixture parameter λ between 0.3 and 0.9, and for a number of latent topics ≥ 1000 .

² Unless stated explicitly otherwise, we recommend at least one tag and at most the number of tags annotated to a resource

Table 5. Evaluation results for tag recommendation based on most frequent tags, based on 5000 latent topics, and their combination with $\lambda = 0.5$.

No. Tags	Most Frequent Tags			Latent Topics			Combination		
	Recall	Prec.	F-Meas.	Recall	Prec.	F-Meas.	Recall	Prec.	F-Meas.
1	0.190	0.467	0.270	0.165	0.437	0.240	0.214	0.537	0.306
2	0.274	0.430	0.335	0.232	0.380	0.288	0.302	0.479	0.370
3	0.329	0.403	0.362	0.271	0.343	0.302	0.357	0.441	0.394
4	0.370	0.388	0.379	0.298	0.316	0.307	0.393	0.415	0.404
5	0.400	0.377	0.388	0.316	0.299	0.307	0.421	0.398	0.409

Table 6. Evaluation results DC09 challenge Task 1 based on most frequent tags, based on 5000 latent topics, and their combination with $\lambda = 0.5$.

No. Tags	Most Frequent Tags			Latent Topics			Combination		
	Recall	Prec.	F-Meas.	Recall	Prec.	F-Meas.	Recall	Prec.	F-Meas.
1	0.010	0.032	0.015	0.045	0.158	0.070	0.049	0.169	0.076
2	0.018	0.031	0.022	0.073	0.131	0.094	0.078	0.140	0.100
3	0.022	0.029	0.025	0.092	0.114	0.102	0.099	0.122	0.110
4	0.026	0.028	0.027	0.094	0.112	0.103	0.102	0.120	0.111
5	0.028	0.027	0.028	0.096	0.112	0.103	0.105	0.120	0.112

3.3 Setups and Results for the Challenge Submission

We have submitted tag recommendations for Task 1 and Task 2 in the ECML PKDD Discovery Challenge 2009. Task 1 aims at recommending tags for arbitrary users annotating a resource in 2009 based on tag assignments until 2008. Thus the test data contain tags, resources, and users which are not available in the training data. The topic models have been trained on the full dataset, comprising about 9.3 Mio tokens for 415 K resources. The test set consists of 43002 posts. Table 6 compares the results for 5000 latent topics with the results using the most frequent tags, and the combination of the two approaches³. Because for the posts in the test set there are only about 0.3 posts per resource in the training set, recommending only the most frequent tags does not recommend any tags for most of the resources. Consequently, recall and precision are significantly lower than for the approach based on latent topics. The combination of the two approaches achieves slightly but consistently better recall and precision.

Task 2 operates on the post-core at Level 2, where all tags, users, and resources occur at least twice in the training data, which comprises about 750 K tokens for 22389 resources. The test set consists of 778 posts, for which there exist on average 5.8 posts in the training set. Table 7 again compares the results

³ Our submission to the DC09 challenge was based on 2500 latent topics without combination with most frequent tags, which achieved an F-measure of 0.098.

Table 7. Evaluation results DC09 challenge Task 2 based on most frequent tags, based on 5000 latent topics, and their combination with $\lambda = 0.5$.

No. Tags	Most Frequent Tags			Latent Topics			Combination		
	Recall	Prec.	F-Meas.	Recall	Prec.	F-Meas.	Recall	Prec.	F-Meas.
1	0.147	0.411	0.216	0.133	0.404	0.200	0.156	0.450	0.232
2	0.223	0.341	0.270	0.204	0.326	0.251	0.252	0.386	0.305
3	0.284	0.305	0.294	0.258	0.281	0.269	0.313	0.339	0.326
4	0.325	0.275	0.298	0.298	0.251	0.272	0.352	0.300	0.324
5	0.357	0.256	0.298	0.319	0.224	0.263	0.386	0.276	0.322

Table 8. Evaluation results for DC09 challenge Task 2 for 5000 latent topics without content

No. Tags	Recall	Precision	F-Measure
1	0.128	0.362	0.189
2	0.191	0.293	0.232
3	0.236	0.254	0.245
4	0.267	0.225	0.244
5	0.299	0.207	0.245

for the two individual approaches and their combination ⁴. As is to be expected, recall and precision are much better than for Task 1, because there is more knowledge available about the tagging practices of users. Like in our internal tests tag recommendation based on most frequent tags outperforms the approach based on LDA, and the combination outperforms the individual approaches.

Table 8 shows the results when only tags are used to elicit latent topics. Recall and precision are consistently lower. Thus taking into account the content of resources leads to more effective latent topics for tag recommendation. However, this does not hold for tag recommendation based on most frequent tags. Recommending the most frequent content terms or tags consistently leads to lower precision and recall.

4 Related Work

Tag recommendation has received considerable interest in recent years. Most work has focused on personalized tag recommendation, suggesting tags to the user bookmarking a new resource: This is often done using collaborative filtering, taking into account similarities between users, resources, and tags. [10] introduces an approach to recommend tags for weblogs, based on similar weblogs tagged by the same user. Chirita et al. [11] realize this idea for the personal desktop, recommending tags for web resources by retrieving and ranking tags from

⁴ Our submission to the DC09 challenge was based on 5000 topics without combination with the most frequent tags and no limit on the number of recommended tags. This achieved an F-measure of 0.258.

similar documents on the desktop. [12] aims at recommending a few descriptive tags to users by rewarding co-occurring tags that have been assigned by the same user, penalizing co-occurring tags that have been assigned by different users, and boosting tags with high descriptiveness (TFIDF).

Sigurbjörnsson and van Zwol [13] also look at co-occurrence of tags to recommend tags based on a user defined set of tags. The co-occurring tags are then ranked and promoted based on e.g. descriptiveness. Jaeschke et al. [14] compare two variants of collaborative filtering and FolkRank, a graph based algorithm for personalized tag recommendation. For collaborative filtering, once the similarity between users on tags, and once the similarity between users on resources is used for recommendation. FolkRank uses random walk techniques on the user-resource-tag (URT) graph based on the idea that popular users, resources, and tags can reinforce each other. These algorithms take co-occurrence of tags into account only indirectly, via the URT graph. Symeonidis et al. [15] employ dimensionality reduction to personalized tag recommendation. Whereas [14] operate on the URT graph directly, [15] use generalized techniques of SVD (Singular Value Decomposition) for n-dimensional tensors. The 3 dimensional tensor corresponding to the URT graph is unfolded into 3 matrices, which are reduced by means of SVD individually, and combined again to arrive at a more dense URT tensor approximating the original graph. Tag recommendation then suggests tags to users, if their weight is above some threshold.

An interactive approach is presented in [16]. After the user enters a tag for a new resource, the algorithm recommends tags based on co-occurrence of tags for resources which the user or others used together in the past. After each tag the user assigns or selects, the set is narrowed down to make the tags more specific. In [17], Shepitsen et al. propose a recommendation system based on hierarchical clustering of the tag space. The recommended resources are identified using user profiles and tag clusters to personalize the recommendation results. Note that they use tag clusters to recommend resources whereas we use LDA topics, which can be considered clusters, to recommend tags.

[3] introduce an approach to tag recommendation using association rules. Resources are regarded as baskets consisting of tags, from which association rules of the form $T_1 \rightarrow T_2$ are mined. On this basis tags in T_2 are recommended whenever the resource contains all tags in T_1 . A comparison of this approach with the approach presented in this paper can be found in [18].

When content of resources is available, tag recommendation can also be approached as a classification problem, predicting tags from content. A recent approach in this direction is presented in [19]. They cluster the document-term-tag matrix after an approximate dimensionality reduction, and obtain a ranked membership of tags to clusters. Tags for new resources are recommended by classifying the resources into clusters, and ranking the cluster tags accordingly.

5 Conclusions and Future Work

In this paper we have presented and evaluated the use of Latent Dirichlet Allocation for collective tag recommendation. Using selected features from the content of resources, tags, and users, we elicit latent topics that comprise typically co-occurring tags and users. On this basis we can recommend tags for new users and resources by mapping them to the latent topics and choosing the most likely tags from the topics. The approach complements simple tag recommendation based on most frequent tags especially for new resources with only few posts. Consequently, combining tag recommendations based on latent topics with tag recommendations based on most frequent tags outperforms the individual approaches.

For future work we want to investigate approaches that take into account individual tagging practices for personalized tag recommendation.

Regarding data sets, we also want to experiment with datasets from different domains, to check whether photo, video, or music tagging sites show different system behavior influencing our algorithms. Another interesting direction we want to follow is to apply LDA not only for tag recommendation but to employ it in the context of recommending resources.

6 Acknowledgments

This work was supported in part by the EU project IST 45035 - Platform for search of Audiovisual Resources across Online Spaces (PHAROS).

References

1. Marlow, C., Naaman, M., Boyd, D., Davis, M.: Ht06, tagging paper, taxonomy, flickr, academic article, to read. In Wiil, U.K., Nürnberg, P.J., Rubart, J., eds.: HYPERTEXT 2006, Proceedings of the 17th ACM Conference on Hypertext and Hypermedia, August 22-25, 2006, Odense, Denmark, New York, NY, USA, ACM (2006) 31–40
2. Golder, S., Huberman, B.A.: Usage patterns of collaborative tagging systems. *Journal of Information Science* **32**(2) (April 2006) 198–208
3. Heymann, P., Ramage, D., Garcia-Molina, H.: Social tag prediction. In: SIGIR '08: Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval, New York, NY, USA, ACM (2008) 531–538
4. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. *Journal of Machine Learning Research* **3** (January 2003) 993–1022
5. Griffiths, T.L., Steyvers, M.: Finding scientific topics. *Proc Natl Acad Sci U S A* **101 Suppl 1** (April 2004) 5228–5235
6. Bhattacharya, I., Getoor, L.: A latent dirichlet model for unsupervised entity resolution. In: SIAM Conference on Data Mining (SDM). (April 2006) 47–58
7. Bíró, I., Siklósi, D., Szabó, J., Benczúr, A.A.: Linked latent dirichlet allocation in web spam filtering. In: AIRWeb '09: Proceedings of the 5th International Workshop on Adversarial Information Retrieval on the Web, New York, NY, USA, ACM (2009) 37–40

8. McCallum, A.K.: Mallet: A machine learning for language toolkit. <http://mallet.cs.umass.edu> (2002)
9. Yao, L., Mimno, D., McCallum, A.: Efficient methods for topic model inference on streaming document collections. In: KDD '09: Proceedings of the 15th ACM SIGKDD conference on Knowledge Discovery and Data Mining, New York, NY, USA, ACM (2009)
10. Mishne, G.: Autotag: a collaborative approach to automated tag assignment for weblog posts. In: WWW '06: Proceedings of the 15th international conference on World Wide Web, New York, NY, USA, ACM (2006) 953–954
11. Chirita, P.A., Costache, S., Nejdl, W., Handschuh, S.: P-tag: large scale automatic generation of personalized annotation tags for the web. In: WWW '07: Proceedings of the 16th international conference on World Wide Web, New York, NY, USA, ACM (2007) 845–854
12. Xu, Z., Fu, Y., Mao, J., Su, D.: Towards the semantic web: Collaborative tag suggestions. In: Proceedings of Collaborative Web Tagging Workshop at 15th International World Wide Web Conference. (2006)
13. Sigurbjörnsson, B., van Zwol, R.: Flickr tag recommendation based on collective knowledge. In: WWW '08: Proceeding of the 17th international conference on World Wide Web, New York, NY, USA, ACM (2008) 327–336
14. Jäschke, R., Marinho, L.B., Hotho, A., Schmidt-Thieme, L., Stumme, G.: Tag recommendations in folksonomies. In Kok, J.N., Koronacki, J., de Mántaras, R.L., Matwin, S., Mladenic, D., Skowron, A., eds.: 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 Lecture Notes in Computer Science., Heidelberg, Germany, Springer (2007) 506–514
15. Symeonidis, P., Nanopoulos, A., Manolopoulos, Y.: Tag recommendations based on tensor dimensionality reduction. In: RecSys '08: Proceedings of the 2008 ACM conference on Recommender systems, New York, NY, USA, ACM (2008) 43–50
16. Garg, N., Weber, I.: Personalized, interactive tag recommendation for flickr. In: RecSys '08: Proceedings of the 2008 ACM conference on Recommender systems, New York, NY, USA, ACM (2008) 67–74
17. Shepitsen, A., Gemmell, J., Mobasher, B., Burke, R.D.: Personalized recommendation in social tagging systems using hierarchical clustering. In Pu, P., Bridge, D.G., Mobasher, B., Ricci, F., eds.: Proceedings of the 2008 ACM Conference on Recommender Systems, RecSys 2008, Lausanne, Switzerland, October 23-25, 2008, New York, NY, USA, ACM (2008) 259–266
18. Krestel, R., Fankhauser, P., Nejdl, W.: Latent dirichlet allocation for tag recommendation. In: 3rd ACM Conference on Recommender Systems, New York City, USA, October 22-25, 2009, to appear. (2009)
19. Song, Y., Zhuang, Z., Li, H., Zhao, Q., Li, J., Lee, W.C., Giles, C.L.: Real-time automatic tag recommendation. In: SIGIR '08: Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval, New York, NY, USA, ACM (2008) 515–522