Weighing the usefulness of social tags for content discovery

Khasfariyati Razikin; Goh, Dion Hoe-Lian; Lee, Chei Sian; Chua, Alton Yeow Kuan

2008

Khasfariyati, R., Goh, D. H. L., Lee, C. S., & Chua, A. Y. K. (2008). Weighing the Usefulness of Social Tags for Content Discovery. 11th International Conference on Asian Digital Libraries: ICADL, 5362, pp.51–60.

<https://hdl.handle.net/10356/94464>

https://doi.org/10.1007/978-3-540-89533-6_6

© 2008 Springer-Verlag Berlin Heidelberg. This is the author created version of a work that has been peer reviewed and accepted for publication by 11th International Conference on Asian Digital Libraries: Universal and Ubiquitous Access To information (ICADL), Springer-Verlag Berlin Heidelberg. It incorporates referee' s comments but changes resulting from the publishing process, such as copyediting, structural formatting, may not be reflected in this document. The published version is available at: DOI [<http://dx.doi.org/10.1007/978-3-540-89533-6>].

Downloaded on 26 Apr 2023 22:27:53 SGT

Weighing the Usefulness of Social Tags for Content Discovery*

Khasfariyati Razikin, Dion Hoe-Lian Goh, Chei Sian Lee, Alton Y.K. Chua

Wee Kim Wee School of Communication & Information, Nanyang Technological University, 31 Nanyang Link, Singapore 637718, Singapore
{khasfariyati, ashlgo, leecs, altonchua}@ntu.edu.sg

Abstract. A new wave of social computing applications has empowered users to create and share a variety of content. This upsurge of user-generated data involves a paradigm shift in terms of the management, searching and accessing of information. Social tagging is one of these ways. This paper serves as an extension to the existing work done on investigating the effectiveness of tags for content discovery using text categorization techniques. In particular, we explored how different tag weighting schemes affect classifier performance. Six text categorization experiments were conducted using a dataset drawn from del.icio.us. The results suggest that not all tags are useful for content discovery even with different weights associated with them. Content analysis was done to understand the relationships between the use of a tag on a document and the document's terms. Implications of this research are discussed.

Keywords: Social tagging, Social computing, Web 2.0, Content discovery and organization, Effectiveness, Text categorization

1 Introduction

Social computing/Web 2.0 applications empower users to create and share a variety of multimedia content including text, images and video. As this new avenue for content creation becomes increasingly popular through services such as blogs, wikis and media sharing, the resulting upsurge of user-generated content requires new approaches for their management, search and access. Beyond established tools such as search engines and directories, social tagging is a promising approach for organizing and discovering content on the Web. It allows users to annotate useful content with uncontrolled keywords (tags), facilitating their future access by the tag creator [12]. In addition, tags may also be shared by other users of the social tagging system, providing an alternative way to discovering and accessing content on the Web.

Social tagging differs from conventional methods of content organization based on taxonomies, controlled vocabularies, faceted classification and ontologies. These methods require experts with domain knowledge, and are rule-bound to ensure that the classification schemes created are consistent [14]. In contrast, tags are “flat”, lacking a predefined taxonomic structure, and their use relies on shared, emergent

* This work is partly funded by A*STAR grant 062 130 0057.

social structures and behaviors, as well as a common conceptual and linguistic understanding within the community [13]. Hence, tags are also known as “folksonomies”, short for “folk taxonomies”, suggesting that they are created by lay users, as opposed to domain experts or information professionals such as librarians.

This unstructured approach to organizing content has drawn criticism by some researchers. For example, the flexibility of free keyword assignment may result in ambiguity of meaning due to a lack of controlled vocabulary [12], resulting in the problem of vocabulary mismatch [5] between tag creators and users [6, 15]. Further, content may be tagged with subjective or ego-centric terms (e.g. “cool”, “todo”, “me”, “toread”) that have meaning only for the tag creator or a select few within a group of users. Tags may also sometimes be driven by the tag creator’s self-serving agenda [3], leading to problems such as tag spamming where popular but irrelevant tags are deliberately used to attract traffic to certain Web sites [9]. Taken together, these issues may hinder the effective use of tags for organizing and sharing content.

Despite these shortcomings, the use of social tagging continues to grow in popularity. Concurrently, there is an emerging body of research that explores their effectiveness for content organization and sharing. For example, from a user’s perspective, work has also been conducted on motivations on behind tagging [1], comparing the use of tags against author assigned index terms in academic papers [8], and on tagging dynamics and usage [4]. Automated, machine learning approaches have also been used to study the ability of tags to classify blogs using text categorization methods [19], and on investigating the effectiveness of tags to classify Web resources in del.icio.us [16].

In this paper, we extend existing work in investigating the effectiveness of tags for content discovery using a machine learning approach (e.g. [17, 19]). In particular, we explore different tag weighting schemes to determine if these affect performance, which is relatively under-researched. Further, to better understand how tags are created, we conduct a content analysis to study the relationships between the use of a tag on a document, and the document’s terms. To the best of our knowledge, there are a limited number of studies that have been done on examining the effectiveness of tags for content discovery using both a machine learning and content analytic approach. The results of our work can be used to tune automated techniques that help users in both seeking content via tags, as well as suggest tags for organizing content. The remainder of this paper is organized as follows. In the next section, we review research related to the present study. We then describe our experimental methodology and present the results. Next, we provide a discussion of the implications of our findings and conclude with opportunities for further work in this area.

2 Related Work

We focus our review on related literature that investigates effectiveness of tags as a means for organizing and sharing content.

From a tag creator’s perspective, work has been done to compare tags with controlled vocabularies. Tags from Connotea (<http://www.connotea.org/>) and Medical Subject Heading (MeSH) terms were evaluated by [11] who found that there was only 11% similarity between MeSH terms and tags. This was because MeSH terms served

as descriptors while tags primarily focused on areas that were of interest to users. Likewise, [8] compared tags with author supplied tags from Cite-U-Like and indexing terms from INSPEC and Library Literature to determine the usage overlap. Results showed that approximately 21% of the tags were the same as the indexing terms. The reason for the divergence was attributed to the different emphases placed on an article by these two groups. For example, tag creators may consider time management information (e.g. “todo”, “toread”, “maybe”) to be important as a tag to indicate a desire to read them in the future, while such information will be disregarded by expert indexers. Together, these findings suggest that tag creators and indexing experts employ vocabularies that have little overlap, potentially causing access problems in social tagging systems.

Research has also been conducted on tag effectiveness using different machine learning approaches. For example, [2] used 350 popular tags and 250 of the most recent blog articles from Technorati. Clustering was done on these articles, and the results suggested that tags were able to organize articles in a broad sense, but not as effective in indicating the specific content for an article. Besides blogs, [17] studied the effectiveness of tags to classify Web content from 100 tags and 20210 documents in del.icio.us. Using Support Vector Machines, experiments were run on two feature sets: document terms only, and document terms plus tags. Surprisingly, results indicated that using document terms only produced better classification results in terms of F-measure than using terms plus tags. Nevertheless, both F-measures from the experiments were relatively low at 0.59 and 0.52, suggesting that not all tags were effective at content discovery, and that the classifier’s performance was likely to be influenced by the tag creator’s motivations, and his/her interpretation of the document content. Finally, [10] investigated tags as a source for metadata to describe music. Using 236974 tags collected for 5722 tracks from last.fm and MyStrands, a Correspondence Analysis was performed to visualize a two-dimensional semantic space defined by the tags. Findings suggest that tags were effective in capturing music similarity, and could be used to describe mood and emotion in music.

While our present study shares the goal of investigating tag effectiveness with the above work, our research is differentiated in several important ways. Firstly, the studies by [11] and [8] were limited to scholarly articles while [2] used blogs. Here, the context of the medium of communication differs in our study. Specifically, in our dataset, the pages that are tagged in del.icio.us are more diverse, and not limited to ordinary Web pages, but also include blogs and academic articles. Put differently, we capture a wider spectrum of content found on the Web. While [17] employed del.icio.us data, it did not examine the influence of different tag weighting schemes, and this may have impacted the results, since the number of tags was significantly fewer than the terms within a document. Taken together, the present study is thus timely as we extend existing work, and perform a more comprehensive analysis of the effectiveness of tags for content discovery.

3 Experimental Setup

The dataset used in these experiments is based on the one used in [17]. The tags and Web documents were harvested from del.icio.us, a popular social tagging service,

from August 2007 and October 2007. During this period, 100 tags and 20210 English language documents were collected. Following [2], the tags were mined starting from the popular tags page. As such our tags will be biased towards the more commonly used ones. Nevertheless, by using popular tags, we were assured that there are a significant number of documents related to each tag that will provide a sufficient dataset size for our experiments.

The tags that were mined consisted of single token terms. The collected documents were processed by removing the HTML elements, JavaScript codes and Cascading Style Sheets elements. This was followed by stop word removal and stemming of the remaining words. For each tag, we selected all the documents that were tagged with the keyword and these were grouped as the positive samples for that particular tag. An equal number of documents, which were associated with a different tag, were randomly selected as negative samples. From this set of positive and negative samples, two-thirds of the documents were used as the training sample while the rest were part of the test set. We made use of the SVM^{light} [7] package for the classification experiments. Binary classifiers were utilized for the study and the default parameters of the SVM package were used during the experiments.

In all, six text categorization experiments were conducted. The first experiment made use of only terms from the documents as features. This experiment served as a baseline for the other experiments due to its use of fundamental features. Here, the TF-IDF values of the terms were used as the feature set. The five subsequent experiments included tags, in addition to terms, as part of its feature vector. The TF-IDF values of the tags were increased by one, three, five, seven and ten times to reflect the differing importance of the tags relatively to the rest of the feature set. The motivation for increasing the TF-IDF values would help determine if this would help in the performance of the classifier.

The output from the classifier, which are precision, recall and F-measure were used to determine the effectiveness of the tags and the different weighting schemes. The performance of the tags was evaluated based on the macro-averaged and micro-averaged precision, recall and F-measure. Macro-averaged values give an indication of the overall performance of the classifier over all tag categories. Micro-averaged values, on the other hand, measure the performance over each document. It emphasizes on the performance of tags that have a larger number of documents.

4 Results

The results obtained for all the experiments are shown in Table 1. The values in bold indicate the highest value obtained among all the six experiments. Surprisingly, as observed from the table, the experiment whose feature vector consisted only of terms obtained the highest scores for the majority of the performance measures. There is a general downward trend for the majority of the performance values as the weighting of the tags were increased. An exception would be the precision values. Apart from this, it appears that the addition of tags as feature vectors and increasing their weights do not help in the improvement for identifying documents associated with the tags. Figures 1 and 2 illustrate the trends for the macro-averaged and micro-averaged values among all experiments respectively.

In terms of macro-averaged values, the precision values obtained were surprisingly almost constant despite increasing the weights of the tags. The increase in precision for the tag weight of three and ten are relatively insignificant with respect to the rest of the precision values obtained. Put differently, increasing the tags' weights does not help in improving the precision of the classification. The recall values show that increasing the weight of the tags does not improve classification performance. Instead, misclassified errors increased as the weights increased. A similar trend is observed for the F-measure values.

Table 1. The macro- and micro-averaged, and standard deviation values for precision, recall and F-measure. The values in bold are the highest values obtained for all experiments.

Experiment	Macro-averaged			Micro-averaged		
	Precision	Recall	F-measure	Precision	Recall	F-measure
Terms only	52.66 (sd = 4.21)	54.86 (sd = 19.05)	52.05 (sd = 10.99)	64.76	54.40	59.14
Tag Weight 1	52.56 (sd = 6.06)	51.60 (sd = 20.75)	50.10 (sd = 13.21)	56.47	54.93	55.69
Tag Weight 3	52.67 (sd = 7.38)	44.04 (sd = 22.94)	44.96 (sd = 15.47)	54.78	48.76	51.60
Tag Weight 5	52.01 (sd = 10.15)	37.93 (sd = 24.90)	39.53 (sd = 18.83)	54.80	42.84	48.09
Tag Weight 7	51.13 (sd = 13.72)	34.76 (sd = 26.26)	36.23 (sd = 21.14)	55.00	39.81	46.19
Tag Weight 10	51.23 (sd = 16.27)	32.38 (sd = 27.08)	33.60 (sd = 22.69)	55.30	37.78	44.89

Micro-averaged performance values, in general, exhibit similar trends as the macro-averaged values. Here, the precision value obtained was at 64.74% in the terms only experiment. The precision values obtained for the subsequent experiments were lower than the first and were almost constant, decreasing when the weight of the tag was set to three, but increasing slightly thereafter. For recall, there was a slight improvement when tags were part of the feature vector. However, increasing the tag weights only made the recall values to deteriorate thereafter. Likewise, the results for micro-averaged F-measure values were similar to recall.

On the whole, these experiments provide a clearer picture on the effect of different tag weighting schemes and their influence on the performance of the classification process. The results show that the addition of tags does not help in improving precision, recall and F-measure values, suggesting further that not all tags could be effectively used for content discovery. There are three possible reasons for this. First, tag creators may have goals other than information discovery by other users [1], and therefore some tags could not be used for this purpose. For instance the documents that are associated with a user-centric, or subjective, tag such as “funny” would vary between tag creators. It was found that some of the documents that had been tagged with “funny” might be humorous to the tag creator, but tag consumers might find it otherwise. Next, social tagging does not have well-defined rules, and the vocabularies employed by different users could vary significantly, leading to inconsistencies in associating documents with tags, impeding effective retrieval of documents [12]. A blog that focuses its contents on cooking and recipes was found to be associated with a diversity of tags such as “blogs”, “blog”, “foodblog” and “foodblogs”. Lastly,

another related reason to account for the poor performance of the classifiers is due to the polysemic nature of tags [6]. As a result, users may impute different interpretations to the same tag so that documents that are associated with a particular tag might not be semantically related at all. For instance, the tag “ruby” might refer to the gem for the uninitiated. However, to a Web applications developer, the tag may refer to a programming language instead.

Fig. 1. The macro-averaged performance for all the experiments

Fig. 2. The micro-averaged performance for all the experiments

5 Tag and Document Analysis

To better understand the performance of the classifiers, we manually analyzed the content of 20 tags to uncover distinguishing patterns. This set comprised the top ten best and poorest performing tags by averaging the F-measure values obtained for all the six experiments (see Table 2).

Table 2. The ten top ten and bottom ten tags and their average F-measure values.

Top Ten Tags	Average F-measure	Bottom Ten Tags	Average F-measure
news	69.10	adobe	20.89
politics	66.78	server	19.45
blogs	65.94	portfolio	19.45
software	64.89	windows	18.62
web2.0	64.45	journal	18.35
interesting	64.42	actionscript	17.68
site	64.17	comic	16.98
resources	63.95	xml	16.42
california	63.62	download	14.83
foodblog	62.20	ajax	13.79

At a glance, the ten best performing tags seemed to have rather broad meanings. In contrast, the bottom ten tags appear to have a narrower definition. It is also interesting to note that at least seven of these tags are computer related. The selected documents were analyzed based on generic/distinctive, subjective/objective characteristics, and the type of medium the content was found in. Our findings are elaborated in the following paragraphs.

Among the top 10 tags, those such as “news”, “site” and “resources” seemed to be more generic in comparison with the rest of the tags. Here, it appears that there is a common understanding among the community of tag creators on the meanings behind these generic tags [11], while the technical tags seem to suffer from more uncertainty in their meanings, given the results in Table 2. Further, the meanings of the top 10 tags appear to be more easily understood by a majority of users as opposed to the poorer performing technical tags whose meanings are presumably known only by those familiar with the field (e.g. “actionscript” and “ajax”).

In addition, a generic tag implies a wider variety of topics associated with it, which also suggests that the terms found in these documents would be varied. Yet, the classifiers performed well on average. A reason for this might be that the tags did not occur frequently in the entire corpus as terms, causing them to have significant TF-IDF values. For example, the subject content associated with the “news” tag was diverse, including the New York Times. While the terms were diverse owing to variety of topics the newspaper accommodates, the term “news” itself did not occur frequently. In contrast, the bottom ten tags had much narrower definitions and were reflected in the associated documents as well. For example, some documents that were tagged with “server” included a tutorial on Microsoft SQL Server and an introduction to Mono. The tag appeared quite frequently in the documents for these two cases. In fact, the tag is the most common word in the dataset, and could be one of the reasons why it performed poorly.

Next, subjective tags are those which could be adjectives, verbs or have other intrinsic qualities, while objective tags are nouns and exhibit extrinsic characteristics [6]. “Interesting” is the only subjective tag that performed well in our experiments. The other nine tags were objective. As found by a prior analysis on subjective tags [17], documents associated with this tag were not topic-specific. Additionally, the documents rarely contained the tag “interesting”, thus making the tag a useful term, leading to a good classifier performance. The bottom ten tags also only had one subjective tag that is “download”. An analysis of the documents associated with the tag revealed that the term appears frequently in the dataset. This finding is in contrast to the best performing subjective tag. Perhaps it is due to the commonness of the tag that “download” performed rather poorly. Similar patterns of findings were observed for the objective tags ranked in the top and bottom ten. However, the documents associated with objective tags were found to be more subject-specific.

Finally, tag creators often associate a document with a tag that defines the type of media the document can be found in [6]. Some examples in the best performing tags were “blogs” and “foodblog”. This observation is true in a social tagging system like del.icio.us where the documents are not restricted to just HTML pages. While the tag “blogs” suggests a diverse range of topics due to its broad meaning, “foodblog” tends to concentrate more on the notion of food. Some examples of documents which are associated with “foodblog” include those which provide recipes, reviews and recommendations of eating establishments, and recipe books and cooking tips. A similar trend was observed in the bottom ten tags, namely, “comic”, “journal” and “portfolio”. A reason for the difference in performance could be due to the amount of text found. In “blogs” and “foodblog”, a significant amount of the content was made up of text. For the others (“comic” and “portfolio”), there was less textual content and more multimedia content, thus affecting classifier performance. Here, other content-based techniques would be needed to address this shortcoming.

In sum, four observations may be drawn on using tags for content discovery: (1) generic tags might surprisingly be better for retrieving relevant documents than specific ones; but more possibly, (2) tags with more accessible meanings are likely to be understood and adopted by the community of tag creators and users; (3) tags which do not appear frequently in the corpus should be recommended for use over those that appear commonly; and (4) automated techniques that suggest tags to users (e.g. [4]) should consider the need for analyzing multimedia content as well.

6 Discussion and Conclusion

To reiterate, the aim of this paper is to investigate the effectiveness of social tags for discovery of relevant content. Six text categorization experiments were conducted in the present study using different tag weighting schemes. Our results indicate that the terms only experiment performed the best in terms of precision, recall and F-measure, further reinforcing existing research that suggest that tags may serve purposes other than content discovery, and that caution must be exercised by tag consumers in their search for relevant information [1, 17].

Our work has yielded the following conclusions. First, introducing tags as feature vectors and increasing their weights does not help in content discovery as the values

for precision, recall and F-measure deteriorate, implying that not all tags are useful for content discovery. This appears to be in contradiction with the Web 2.0 notion of the benefits of harnessing the collective intelligence of users. A reason for this could be related to the level of familiarity of the tag creators with the different aspects of social tagging and/or the nature of the topic itself. A novice tag creator might use ineffective tags or those that are not drawn from the vocabulary of the tagging community, impeding future retrieval by others [18]. Another factor could be due to the polysemic, synonymous and homonymous nature of tags [6]. Here, there would be difficulty in disambiguating the tags without the need for looking at the context of the associated document. As such, measures should be put in place when using tags to search relevant documents to ensure that ambiguity in meanings are addressed. For example, the system could prompt the user to determine the exact context of the ambiguous tag, thus helping them obtain relevant documents that meet their information needs. Next, the best performing tags are seemingly generic tags with broad meanings. Our content analysis suggests that their meanings appear to be widely understood. Put differently, despite their generality, there is some stabilization in terms of tag usage in which tag creators tend to select tags which are commonly understood by other members of the del.icio.us community [18]. For example, a tag creator may apply similar tags to a document previously given by other tag creators. Thus, simply recommending popular tags to users might not be an appropriate method. Instead, recommendations based on the patterns of knowledgeable members of the tagging community would be preferable. This would ensure that the tags have more relevancy and semantic relation to the document in question [4].

There are limitations in our study that could be addressed in future work. First, all terms appearing in the documents, with the exception of stopwords, were deemed useful. There could be a possibility that the classifier could incorrectly generalize its model based on the existence of a rare term that does not have any meaning to the tag. Perhaps it might be helpful to identify which features are useful for a tag to eliminate noisy terms. Examples might include χ^2 or information gain [20]. Secondly, this study made use of only popular tags. However, the number of such tags is proportionately smaller than the entire collection of tags in del.icio.us. Future work could utilize a wider variety of tags to determine if performance may be affected, as less popular tags may be associated with more esoteric, but more specific concepts, and therefore could result in better classifier performance. In addition, research could be done on using tags to filter or rank documents associated with them. Lastly, future work could look into replicating the same experiments for different datasets such as Cite-U-Like or Amazon, where the intent of the tag creators could be different from that of del.icio.us, as well as extending the analyses to multimedia content such as those found in YouTube and Flickr.

References

1. Ames, M., Naaman, M.: Why we tag: motivations for annotation in mobile and online media. In: SIGCHI Conference on Human Factors in Computing Systems p. 971-980. ACM, New York (2007).

2. Brooks, C.H., Montanez, N.: Improved annotation of the blogosphere via autotagging and hierarchical clustering. In: Proceedings of the 15th international conference on World Wide Web p. 625-632. ACM, New York (2006).
3. Chua, A.: Knowledge sharing: A game people play. *Aslib Proceedings*. 55(3) 117 - 129 (2003)
4. Farooq, U., Kannampallil, T.G., Song, Y., Ganoe, C.H., Carroll, J.M., Giles, L.: Evaluating tagging behavior in social bookmarking systems: metrics and design heuristics. In: Proceedings of the 2007 international ACM Conference on Supporting Group Work p. 351-360. ACM, New York (2007).
5. Furnas, G.W., Landauer, T.K., Gomez, L.M., Dumais, S.T.: The vocabulary problem in human-system communication. *Commun. ACM*. 30(11) 964 - 971 (1987)
6. Golder, S.A., Huberman, B.A.: Usage patterns of collaborative tagging systems. *Journal of Information Science*. 32(2) 198-208 (2006)
7. Joachims, T.: Text Categorization with Support Vector Machines: Learning with Many Relevant Features. In: Proceedings of the 10th European Conference on Machine Learning p. 137-142. Springer-Verlag, Berlin (1998).
8. Kipp, M.E.: Exploring the context of user, creator and intermediate tagging. In: *ASIS&T 2006 Information Architecture Summit*, (2006).
9. Koutrika, G., Effendi, F.A., Gyöngyi, Z., Heymann, P., Garcia-Molina, H.: Combating spam in tagging systems. In: Proceedings of the 3rd international workshop on Adversarial information retrieval on the web p. 57-64. ACM, New York (2007).
10. Levy, M., Sandler, M.: A semantic space for music derived from social tags. In: Proceedings of the 8th International Conference on Music Information Retrieval, ISMIR 2007, (2007).
11. Lin, X., Beaudoin, J.E., Bui, Y., Desai, K.: Exploring characteristics of social classification. In: Proceedings 17th Workshop of the American Society for Information Science and Technology Special Interest Group in Classification Research, (2006).
12. Macgregor, G., McCulloch, E.: Collaborative tagging as a knowledge organisation and resource discovery tool. *Library Review*. 55(5) 291 - 300 (2006)
13. Marlow, C., Naaman, M., Boyd, d., Davis, M.: HT06, tagging paper, taxonomy, Flickr, academic article, to read. In: Proceedings of the seventeenth conference on Hypertext and hypermedia p. 31-40. ACM, New York (2006).
14. Morville, P.: *Ambient findability*. 2005, Beijing: O'Reilly, Sebastopol, CA.
15. Olsen, K.A., Sochats, K.M., Williams, J.G.: Full text searching and information overload. *International Information and Library Review*. 30(2) 105-122 (1998)
16. Razikin, K., Goh, D.H.-L., Cheong, E.K.C., Ow, Y.F.: Social efficacy of tags in social tagging system, In: 10th International Conference on Asian Digital Libraries, ICADL 2007. Lecture Notes in Computer Science 4822, p. 506-507. Springer, Berlin (2007).
17. Razikin, K., Goh, D.H.-L., Chua, A.Y.K., Lee, C.S.: Can social tags help you find what you want? In: *Research and Advanced Technology for Digital Libraries*. 12th European Conference on Digital Libraries, ECDL 2008. Lecture Notes in Computer Science 5173, p. 50-61. Springer, Berlin (2008).
18. Sen, S., Lam, S., Rashid, A.M., Cosley, D., Frankowski, D., Osterhous, J., Harper, F.M., Riedl, J.: Tagging, communities, vocabulary, evolution. In: Proceedings of the 2006 ACM Conference on Computer Supported Cooperative Work p. 181-190. ACM, New York (2006).
19. Sun, A., Suryanto, M.A., Liu, Y.: Blog classification using tags: an empirical study, In: 10th International Conference on Asian Digital Libraries, ICADL 2007. Lecture Notes in Computer Science 4822, p. 307-316. Springer, Berlin (2007).
20. Yang, Y., Pedersen, J.O.: A comparative study on feature selection in text categorization. In: Proceedings of ICML-97, 14th International Conference on Machine Learning p. 412-420. Morgan Kaufmann Publishers, San Francisco (1997).