

Optimizing Academic Conference Classification using Social Tags

Jing Xia

Huazhong University of
Science and Technology,
Wuhan, China
0809xiajiejing@gmail.com

Kunmei Wen

Huazhong University of
Science and Technology,
Wuhan, China
kmwen@hust.edu.cn

Ruixuan Li

Huazhong University of
Science and Technology,
Wuhan, China
rxli@hust.edu.cn

Xiwu Gu

Huazhong University of
Science and Technology,
Wuhan, China
guxiwu@hust.edu.cn

Abstract- Automatically classifying academic conference into semantic topic promises improved academic search and browsing for users. Social tagging is an increasingly popular way of describing the topic of academic conference. However, no attention has been devoted to academic conference classification by making use of social tags. Motivated by this observation, this paper proposes a method which utilizes social tags as well as the content of academic conference in order to improve automatically identifying academic conference classification. The proposed method applies different automatic classification algorithms to improve classification quality by using social tags. Experimental results show that this method mentioned above performs better than the method which only utilizes the content to classify academic conference with 1% Precision measure score increase and 1.64% F1 measure score increase, which demonstrates the effectiveness of the proposed method.

Keywords- classification; academic conference; feature selection.

I. INTRODUCTION

Nowadays, the websites and systems which include academic social resources have recently emerged as a popular way or platform for users to discover and classify academic conference (e.g., DBLP, Google Scholar and ArnetMiner). Moreover, there is a trend to make more and more use of user-generated tags to describe the topic of academic conference (e.g., WikiCFP). The academic information can benefit many applications. For example, Tang et al. [1] develop an academic search system ArnetMiner by extracting and mining academic information. Ramage et al. [2] try to improve automatic clustering of web pages by using user-generated tags. Aliakbary et al [3] investigate the use of social tags for web page classification. However, to the best of our knowledge, there is no any work directly using social tags to automatically identify academic conference classification. Academic conference classification has traditionally been solved by using the content or sub-content of academic conference. It is necessary for users to identify the topic of their interested academic conference as accurately as possible. In academic field, the importance of academic conference classification can help scholars to submit their papers to more relevant

conferences, etc. Therefore, it is meaningful to make use of social tags as a complementary data source to academic conferences classification.

In this paper, we try to address one central question: How can social tags be used to improve academic conference classification? WikiCFP is used as the source of academic conferences and social tags. In general, the detailed conference information of WikiCFP is mainly provided in the section of Call For Papers (CFPs). Therefore CFPs is selected as the content of academic conference approximatively. We can extract the information directly from the section of WikiCFP.

First we make use of CFPs as the main data source in feature selection processing of five common classification algorithms, which is referred to as Classifier1. By analyzing, we found that social tags of academic conferences of WikiCFP are provided in the section of named Categories. Next we use the corresponding tags of academic conference as a complementary data source in feature selection processing of five common classification algorithms, which is referred to as Classifier2. Therefore Classifier 1 represents the common classification method which does not use tags. Classifier 2 represents the proposed method in this paper. We adopt mutual information (MI) as the feature selection method for the two classifiers. The classification results of Classifier1 and Classifier2 are evaluated by comparing the scores of Precision measure and the scores of F1 measure. Experimental results show that Classifier 2 clearly outperforms Classifier 1.

The main contributions of this paper are the following. (1) Social tags are explored as a novel evidence to classify academic conference. (2) We propose a unified framework for feature selection using MI algorithm based on the content and social tags of academic conference. (3) Extensive experiments are conducted to demonstrate the effectiveness of social tags in optimization of academic conference classification.

The rest of the paper is organized as follows. Section 2 reviews related work on social tags and their application in classification. The core of our paper, Section 3, proposes a method that is used to classify academic conference using the content and social tags. Social tags can optimize classification results to some extent. Section 4 shows the experimental results, and we draw some conclusions. Finally express our thanks in Section 5.

II. RELATED WORK

To the best of our knowledge, this paper is the first work to explore social tags for academic conference classification. Some research works have been done in related areas.

Social tags, as a new way of user-generated data, can benefit many applications, such as information retrieval [4, 5], semantic web [6], web page clustering [7], user interest mining [8]. Bao et al. [9] have utilized social tags to optimize web search and Tang et al. [10] have tried to learn an ontology from the tags generated by authors freely.

Web page classification has been extensively studied. Many researchers resort to making use of more information to improve classification results because web pages usually contain rich heterogeneous features. Hyperlink [11], html meta data [12] and query log [13] are explored to improve classification results of web pages. Glover et al. [14] also used anchor text and hyperlink surrounding text to classify Web pages. However, there is little work on exploring social tags for academic conference classification. This paper combines CFPs and social tags to classify academic conference. CFPs often consist of important information of conference, such as the topic of conference, workshops, important dates of conference. Therefore CFPs can be used to substitute for the content of conference. K. M. Schneider [15] has extracted information of conferences from CFPs.

In recent years, exploring social tags for classification has attracted considerable attentions. Noll et al. [16] systematically have analyzed exploring social tags for web document classification. Sakurai et al. [17] have proposed a method that classifies bloggers using social tags. Yin et al. [18] have explored social tags for web object classification and validated that the exploration of social tags effectively boosts web object classification. Social tagging has become a popular activity of web users to annotate web pages. So web page classification using social tags has been investigated for a long time. Aliakbary et al [3] have utilized social tags to classify common web pages. Unlike the works mentioned above, we propose a general framework for feature selection processing by using of social tags as a complementary data source for the purpose of optimizing the classification results with high accuracy.

III. PROPOSED METHOD

In this section, we first describe the problem of definition. Then how to obtain the representation of academic conference are explained. Finally, applying five common classification algorithms to classify academic conference are explained based on the representation of academic conference.

A. Problem of Definition

1) *Definition of Representation of Academic Conference:* $D = \{(r_1, t_1), \dots, (r_n, t_n)\}$ is a collection of tagged resources, where $r_i \in R$ is a resource (CFPs of academic conference) and t_i is a set of tags assigned to resource r_i . The tagged resources are used as the training set in this paper. We suppose that the training set is

classified manually to K classes which constructed the class set denoted by $Class = \{Class_1, Class_2, \dots, Class_K\}$, where $Class_i$ denotes a specific topic of academic conference.

V_T is denoted as the tags vocabulary and V_C as the CFPs vocabulary. Based on the mentioned vocabularies, using MI algorithm, CFPs feature vector and tags feature vector are obtained for each class. w is denoted as CFPs feature vector, and t is denoted as tags feature vector, for each class in the training set respectively. The tags feature vector and the CFPs feature vector for each class in the class set are showed in Table I.

TABLE I. THE REPRESENTATION OF ACADEMIC CONFERENCE FOR EACH CLASS

Class	tags feature vector for class	CFPs feature vector for class
$Class_1$	t_1	w_1
$Class_2$	t_2	w_2
...
$Class_K$	t_K	w_K

We can obtain the representation of academic conference from the K tags feature vectors and the K CFPs feature vectors. In this paper, we present an annotated academic conference as a vector:

$$C = \langle W, T \rangle \quad (1)$$

In this paper, the strategy of selecting tags feature vector t of each class is to select the top-M tags with the top-M MI scores. So the tags feature vector of academic conferences, denoted as:

$$T = \langle t_1, t_2, \dots, t_K \rangle \quad (2)$$

The vector T is constructed by the K tags feature vectors for different K classes. Similarly, the strategy of selecting CFPs feature vector w of each class is to select the top-N terms with the top-N MI scores. So the CFPs feature vector of academic conferences, denoted as:

$$W = \langle w_1, w_2, \dots, w_K \rangle \quad (3)$$

After feature selection, the representation of academic conference can be used the vector $C = \langle W, T \rangle$. Through assigning term frequency and tag frequency to the vector $C = \langle W, T \rangle$, academic conferences are represented as vectors:

$$C = \langle tf_1, \dots, tf_n, af_1, \dots, af_m \rangle \quad (4)$$

According to the vector $C = \langle W, T \rangle$, where $|W| = n$ denote that there are n terms in the vector W and $|T| = m$ denote that there are m tags in the vector T . In detail, we denote the vector $C = \langle tf_1, \dots, tf_n, af_1, \dots, af_m \rangle$ as the representation of a academic conference, where tf_i is the term frequency and af_j is the tag frequency. Given this, the goal of this work is to identify academic conference classification according to the representation of academic conference $C = \langle tf_1, \dots, tf_n, af_1, \dots, af_m \rangle$ automatically by using common classification algorithms.

Therefore, in Classifier 2, the representation of academic conference can be used the vector $C = \langle tf_1, \dots, tf_n, af_1, \dots, af_m \rangle$. In Classifier 1, the representation of academic conference can be used t vector:

$$C = \langle tf_1, \dots, tf_n \rangle \quad (5)$$

The vector is not contained tag frequency. The difference between the Classifier 1 and the Classifier 2 is the representation of academic conference.

2) *Definition of Classification of Academic Conference:* Under the academic conference presentation, the goal of optimizing academic conference classification using social tags can be defined as: Given a set of tagged academic conference, each of which is represented in a vector $C = \langle tf_1, \dots, tf_n, af_1, \dots, af_m \rangle$. And we want to use a classification algorithm F:

$$F: C \rightarrow \{Class_1, Class_2, \dots, Class_K\} \quad (6)$$

The function maps each academic conference C to a specific class $Class_i \in \{Class_1, Class_2, \dots, Class_K\}$.

B. Feature Selection of Academic Conferences using Social Tags

In our proposed method, academic conference is represented by its *CFPs feature vector* and *tags feature vector*. The whole process of obtaining the representation of academic conference is showed in Figure 1.

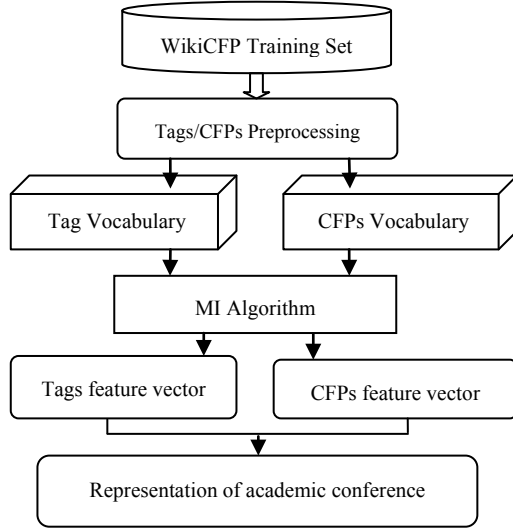


Fig. 1. The whole process of feature selection for academic conference

As the figure shows, in the first stage, two vocabularies are constructed and the two vocabularies could not be merged. In the second stage, two feature vectors are obtained by using MI algorithm to select the subset of the two vocabularies respectively. Then in the last stage, we use term frequency assigned to the corresponding CFPs feature vector and tag frequency assigned to the corresponding tags feature vector. And finally each academic conference in the training set can be represented by using the two vectors assigned the corresponding term frequency and tag frequency.

1) *Preprocessing of Social Tags and CFPs:* Through the analysis of the features of social tags, we found that social tags have some kind of noise or ambiguity caused by annotators. Therefore, we tried to decrease the effect of these noises in the preprocessing phase. Preprocessing includes three steps shown in Figure 2.




Fig. 2. Preprocessing of social tags

Lowercase conversion refers to converting the social tags to lowercase. Bad tags removing refers to removing commonly used stop words (e.g. numbers and special symbols). Stemming refers to replacing the tags with word stems using the API interface of WordNet [19]. The process of stemming unifies different forms of a common word (e.g. “semantic” and “semantics”).

We also use the same preprocessing phase to process CFPs of academic conferences. After preprocessing of CFPs and social tags, feature selection processing can be implemented subsequently.

2) *Feature Selection Processing:* The aim of feature selection is to obtain the subset of the training set and use this subset as the features. In this paper, we try to find relevant features from CFPs and social tags of academic conference using MI method. In this paper, both tags feature selection and CFPs feature selection are implemented respectively. The process of tags feature selection is described in Figure 3. The process mainly contains three steps.




Fig. 3. The process of tags feature selection

Tags vocabulary construction refers to constructing the vocabulary from the whole social tags obtained from the training set. Applying MI algorithm refers to measuring the degree in which the presence/absence of each tag in the vocabulary contributes to making the correct class decision. The outputs of this step are MI scores for each class. Tags feature selection refers to selecting representative tags relatively according to the MI scores. The strategy of selecting features for each class is to select top-M tags having top-M MI scores correspondingly in each class.

In the process of CFPs feature selection, we also adopt almost same steps mentioned in the tags feature selection. Different from tags feature selection, top-N terms with top-N MI scores are selected correspondingly for each class.

C. Automatic Academic Conference Classification

By using the steps given above, each academic conference in the training set is represented by its corresponding vector $C = \langle W, T \rangle$. We adopt five

common classification algorithms to classify academic conference based on the representation of academic conference. Specifically, in accordance with the file format, we input vector $C = \langle tf_1, \dots, tf_n, af_1, \dots, af_m \rangle$. Then classification result can be outputted for each academic conference together with some result-analyzing parameters. In addition, for the purpose of identifying the effect of the proposed method quantitatively, we also only utilize CFPs to classify academic conference. We compare the two classification results and find that using social tags is a good way of optimizing academic conference classification.

IV. EVALUATION OF EXPERIMENTAL RESULTS AND DISCUSSION

In this paper, we adopt the third-party classification toolkit Weka-3-6. Specifically, five common algorithms which have been implemented in the toolkit are utilized in the experiment. The five typical classification algorithms are NaïveBayes algorithm, Support Vector Machine algorithm denoted by SMO in Weka-3-6, J48 algorithm, ClassificationViaRegression algorithm and IBk algorithm.

A. Dataset

In the paper, we have conducted experiments on the data set extracted from the famous academic conference website WikiCFP. First, we extract 6936 academic conferences with triples of (link, CFPs, tags). After selecting, our data set contains 4465 academic conferences with non-empty CFPs and non-empty social tags. After the preprocessing by using the method mentioned in the section 3.2, the tag vocabulary contains 58740 terms. It means that each academic conference is tagged using more than 13 terms.

B. Criteria for Manual Classification

TABLE II. DETAIL INFORMATION OF 11 CLASSES

Class Number	Class Name	Class Number	Class Name
1	Algorithms and Theory of Computation	7	Distributed and Parallel Computing
2	Application	8	Hardware and Architecture
3	Artificial Intelligence	9	Security
4	Communication and Computer Networks	10	Programming Languages and Software Engineering
5	Computer Graphics, Image and Multimedia	11	Miscellaneous
6	Database, Data Mining and Information Retrieval		

In this paper, manual classification results for academic conferences in the data set are viewed as the ground truth. We invite six postgraduates and PhDs whose research interests are information retrieval to identify the classification for academic conference. This paper intends to make use of the authoritative literature DBLP, Citeseer and the literature information Google Directory as the reference criteria of classification. Table II shows the detail information about 11 classes in the dataset.

As the table shows, totally, we divide the dataset into 11 classes. Each class name means a specific topic semantically. Each class name is defined by us according to academic resources mentioned above. Class number is in front of the corresponding class name. In the follow-up section, class number is used to substitute for the corresponding class name.

C. Evaluation of the Proposed Method

First, in order to evaluate how well the proposed method predicted a certain percentage of the data set which is held out for testing totally, we select some different amounts of data set to evaluate. The evaluation metric we used is Precision measure and F1 measure. Second, in order to show the parameters of the results of the proposed method in detail, we partition the data set to Training Set (40% of the dataset) and Test Set (60% of the dataset), which viewed as a specific case. The Training Set is used to build the representation of academic conference. The Test Set is used to examine the proposed method. The social tags of the Test Set are not included in the process of constructing the vocabularies and feature vectors. Then we compare the proposed method adopted in Classifier 2 with the method which only utilizes CFPs to classify academic conferences adopted in Classifier 1.

1) *Evaluation of the Proposed Method Totally:* We partition the data set into two sets. One is used to as the Training Set. The other is used to as the Test Set. Table III shows how well the proposed method predicted a certain percentage of the data set which is held out for testing based on Precision measure and F1 measure totally.

TABLE III. DETAIL EVALUATION RESULTS FOR DIFFERENT PERCENTAGES OF THE DATA SET BASED ON PRECISION

the percentage of data set for testing classification algorithm	30%	50%	70%	90%
Naïve Bayes	0.882	0.934	0.955	0.935
Support Vector Machine (SMO)	0.931	0.93	0.952	0.984
J48 algorithm	0.907	0.945	0.966	0.965
ClassificationViaRegression	0.785	0.84	0.934	0.873
IBk algorithm	0.914	0.917	0.93	0.947

Table III shows the detail Precision scores for different percentages of data set for testing. The row of the table is the percentage of data set for testing. The column of the table is the classification algorithm. The value of the table is the Precision score.

Table IV shows the detail F1 measure scores for different percentages of data set for testing. The row of the table is the percentage of data set for testing. The column of the table is the classification algorithm. The value of the table is the F1 measure score.

TABLE IV. DETAIL EVALUATION RESULTS FOR DIFFERENT PERCENTAGES OF THE DATA SET BASED ON F1 MEASURE

the percentage of data set for testing classification algorithm	30%	50%	70%	90%
Naïve Bayes	0.885	0.928	0.952	0.929
Support Vector Machine (SMO)	0.915	0.927	0.944	0.977
J48 algorithm	0.913	0.932	0.959	0.953
ClassificationViaRegression	0.778	0.838	0.934	0.897
IBk algorithm	0.907	0.914	0.92	0.927

By analyzing the results from the two tables, we can totally conclude that the proposed method in this paper can be used to classify academic conference. In detail, we find that with the increase of the amount of the data set for training, the effectiveness of the proposed method may be worse. The reason is that it would occur the overfitting problem with the increase of the amount of the training set.

2) *Comparing Classifier 1 with Classifier 2 using Precision Measure:* Based on precision measure, comparison results of the Classifier 1 with the Classifier 2 are shown in Figure 4. Precision estimates the correct classified results accounted for the manual classification results in the Test Set.

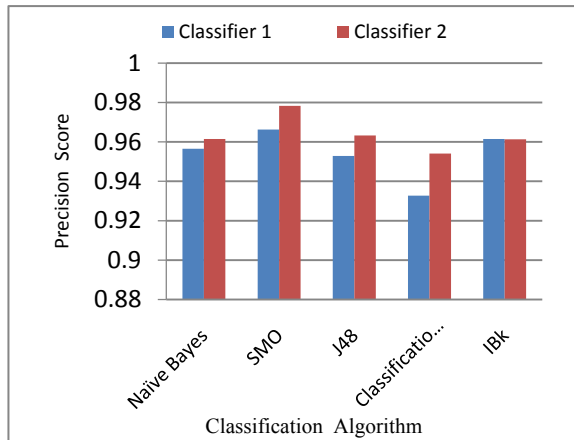


Fig. 4. Average precision comparison between the two classifiers

The higher the precision score, the performance of Classifier is better. In order to measure the performance of Classifier 2, Table V shows the detail Precision scores for each class. The row of the table is the sequence of classification algorithms, which refer to Naïve Bayes, SMO, J48, ClassificationViaRegression and IBk. The column of the table is the class number. The value of the table is the Precision score. By analyzing the comparison results of the two classifiers, we conclude that the Classifier 2 performs better than the Classifier 1. Specifically, the Classifier 2 optimizes the classification results by 1% in Precision measure averagely comparing with the Classifier 1.

TABLE V. DETAIL PRECISION SCORES FOR EACH CLASS

Algorithm Class	1	2	3	4	5
1	0.939	1	1	0.967	1
2	1	1	1	0.947	0.947
3	0.982	1	0.982	0.981	1
4	0.98	0.98	0.979	0.959	0.959
5	0.963	0.962	0.857	0.926	0.893
6	0.927	0.963	0.907	1	0.963
7	0.889	0.889	1	0.875	0.9
8	0.9	0.9	0.925	0.854	0.946
9	1	1	1	0.917	0.958
10	0.962	1	1	0.98	0.962
11	1	1	1	1	1

3) *Comparing Classifier 1 with Classifier 2 using F1 Measure:* In statistics, the F1 score can be interpreted as a weighted average of the precision and recall, where F1 score reaches its best value at 1 and the worst score at 0. Based on F1 measure, comparison results of the Classifier1 with the Classifier 2 are shown in Figure 5.




Fig. 5. Average F1 measure comparison between the two classifiers

Table VI shows the detail F1 scores for each class in Classifier 2. The row and the column of Table VI are the same as the Table V. The value of the Table VI is the F1 measure score. By analyzing the comparison results of the two classifiers, we can conclude that the Classifier 2 performs better than the Classifier 1. Specifically, the Classifier 2 optimizes the classification results by 1.64% in F1 measure averagely comparing with the Classifier 1.

TABLE VI. DETAIL F1 SCORES FOR EACH CLASS

Algorithm Class	1	2	3	4	5
1	0.969	1	1	0.951	1
2	0.971	1	1	0.973	0.973
3	0.982	1	0.991	0.964	0.963
4	0.97	0.99	0.958	0.959	0.959
5	0.963	0.943	0.873	0.926	0.909
6	0.944	0.972	0.916	1	0.972
7	0.8	0.8	0.842	0.737	0.857
8	0.935	0.935	0.961	0.897	0.946
9	0.958	0.957	0.979	0.917	0.958
10	0.98	1	1	0.98	0.98
11	0.977	1	1	0.977	0.952

D. Conclusions and Discussions

Using social tags is rapidly increasing in academic field, providing potentially useful information to improve academic conference classification. This paper studied how to make use of social tags as a complementary data source to optimize the classification of academic conference. Social tags can actively benefit the academic conference classification.

However, for further improving the performance of the proposed method in this paper, there are still several problems to further address. First, we must consider the sparsity of social tags and may be correspondingly propagating the tags between academic conferences. Secondly, as the semantic similarity of social tags, we can incorporate the semantic similarity to the proposed method. Finally, we found that a few spam exists in the content of academic conference. These spam information may have negative impact to the experiment result. Now we only manually remove the spam information. Therefore we may filter out spam tags automatically by using linguistic analysis before executing the method proposed in this paper.

V. ACKNOWLEDGEMENT

This paper is supported by National Natural Science Foundation of China under Grant No. 60873225, 60773191, 70771043, Natural Science Foundation of Hubei Province (NSF-HB) under grant No. 2009CDB298, Open Foundation

of State Key Laboratory of Software Engineering under Grant No. SKLSE20080718 and Innovation Fund of Huazhong University of Science and Technology under Grant No. Q2009021.

REFERENCES

- [1] J. Tang, J. Zhang, L.M. Yao, and J.Z. Li, "ArnetMiner: Extraction and mining of an academic social network," In: Proc. of the International Conf. on Knowledge Discovery and Data Mining (KDD 2008), pp.990-998, 2008.
- [2] D. Ramage, P. Heymann, C.D. Manning, and H. Garcia-Molina, "Clustering the tagged web," WSDM 2009, pp.54-63, 2009.
- [3] S. Aliakbary, H. Abolhassani, H. Rahmani, and B. Nobakht, "Web Page Classification Using Social Tags," In: Proc. of the International Conf. on Computation Science and Engineering (CSE 2009), pp.588-593, 2009.
- [4] D. Zhou, J. Bian, S. Zheng, H. Zha, and C.L. Giles, "Exploring social annotations for information retrieval," In: Proc. of the International Conf. on World Wide Web (WWW 2008), pp.715-724, 2008.
- [5] R. Schenkel, T. Crecelius, M. Kacimi, S. Michel, T. Neumann, J. X. Parreira, and G. Weikum, "Efficient top-k querying over social tagging networks," In: Proc. of the International ACM SIGIR Conf. on Research and Development in Information Retrieval, pp.523-530, 2008.
- [6] X. Wu, L. Zhang, and Y. Yu, "Exploring social annotations for the semantic web," In: Proc. of the International Conf. on World Wide Web (WWW 2006), pp.417-426, 2006.
- [7] C.H. Brooks, and N. Montanez, "Improved annotation of the blogosphere via autotagging and hierarchical clustering," In: Proc. of the International Conf. on World Wide Web (WWW 2006), pp.625-632, 2006.
- [8] X. Li, L. Guo, and Y.E. Zhao, "Tag-based social interest discovery," In: Proc. of the International Conf. on World Wide Web (WWW 2008), pp.675-684, 2008.
- [9] S.H. Bao, G.R. Xue, X. Wu, Y. Yu, B. Fei, and Z. Su, "Optimizing web search using social annotations," In: Proc. of the International Conf. on World Wide Web (WWW 2007), pp.501-510, 2007.
- [10] J. Tang, H.F. Leung, Q. Luo, D. Chen, and J. Gong, "Towards ontology learning from folksonomies," In: Proc. of the International Joint Conf. on Artificial Intelligence (IJCAI 2009), pp.2089-2094, 2009.
- [11] S. Chakrabarti, B. Dom, and P. Indyk, "Enhanced hypertext categorization using hyperlinks," SIGMOD 1998, pp.307-318, 1998.
- [12] R. Ghani, S. Slattery, and Y. Yang, "Hypertext categorization using hyperlink patterns and meta data," In: Proc. of the International Conf. on Machine Learning (ICML 2001), pp.178-185, 2001.
- [13] D. Shen, K.T. Sun, Q. Yang, and Z. Chen, "A comparison of implicit and explicit links for web page classification," WWW 2006, pp.643-650, 2006.
- [14] E. J. Glover, K. Tsioutsoulouklis, S. Lawrence, D.M. Pennock, and G.W. Flake, "Using Web structure for classifying and describing Web Pages," In: Proc. of the International Conf. on World Wide Web (WWW 2002), pp.562-569, 2002.
- [15] K.M. Schneider, "Information extraction from calls for papers with conditional random fields and layout features," Journal of Artif. Intell. Rev, vol.25(1-2), pp.67-77, 2006.
- [16] M.G. Noll, and C. Meinel, "Exploring Social Annotations for Web Document Classification," SAC 2008, pp.2315-2320, 2008.
- [17] S. Sakurai, H. Tsutsui, and R. Orihara, "Classification of Bloggers using Social Annotations," AINA 2009, pp.929-936, 2009.
- [18] Z.J. Yin, R. Li, Q.Z. Mei, and J.W. Han, "Exploring social tagging graph for web object classification," In: Proc. of the International Conf. on Knowledge Discovery and Data Mining (KDD 2009), pp.957-966, 2009.
- [19] WordNet, <http://wordnet.princeton.edu/>.