

# Recommendation in Internet Forums and Blogs

**Jia Wang**  
Southwestern Univ.  
of Finance &  
Economics China  
wj96@sina.cn

**Qing Li**  
Southwestern Univ.  
of Finance &  
Economics China  
liq\_t@swufe.edu.cn

**Yuanzhu Peter Chen**  
Memorial Univ. of  
Newfoundland  
Canada  
yzchen@mun.ca

**Zhangxi Lin**  
Texas Tech Univ.  
USA  
zhangxi.lin  
@ttu.edu

## Abstract

The variety of engaging interactions among users in social media distinguishes it from traditional Web media. Such a feature should be utilized while attempting to provide intelligent services to social media participants. In this article, we present a framework to recommend relevant information in Internet forums and blogs using user comments, one of the most representative of user behaviors in online discussion. When incorporating user comments, we consider structural, semantic, and authority information carried by them. One of the most important observation from this work is that semantic contents of user comments can play a fairly different role in a different form of social media. When designing a recommendation system for this purpose, such a difference must be considered with caution.

## 1 Introduction

In the past twenty years, the Web has evolved from a framework of information dissemination to a social interaction facilitator for its users. From the initial dominance of static pages or sites, with addition of dynamic content generation and provision of client-side computation and event handling, Web applications have become a prevalent framework for distributed GUI applications. Such technological advancement has fertilized vibrant creation, sharing, and collaboration among the users (Ahn et al., 2007). As a result, the role of Computer Science is not as much of designing or implementing certain data communication techniques, but more of enabling a variety of creative uses of the Web.

In a more general context, Web is one of the most important carriers for “social media”, e.g. In-

ternet forums, blogs, wikis, podcasts, instant messaging, and social networking. Various engaging interactions among users in social media differentiate it from traditional Web sites. Such characteristics should be utilized in attempt to provide intelligent services to social media users. One form of such interactions of particular interest here is *user comments*. In *self-publication*, or *customer-generated media*, a user can publish an article or post news to share with others. Other users can read and comment on the posting and these comments can, in turn, be read and commented on. Digg ([www.digg.com](http://www.digg.com)), Yahoo!Buzz ([buzz.yahoo.com](http://buzz.yahoo.com)) and various kinds of blogs are commercial examples of self-publication. Therefore, reader responses to earlier discussion provide a valuable source of information for effective recommendation.

Currently, self-publishing media are becoming increasingly popular. For instance, at this point of writing, Technorati is indexing over 133 million blogs, and about 900,000 new blogs are created worldwide daily<sup>1</sup>. With such a large scale, information in the *blogosphere* follows a Long Tail Distribution (Agarwal et al., 2010). That is, in aggregate, the not-so-well-known blogs can have more valuable information than the popular ones. This gives us an incentive to develop a recommender to provide a set of relevant articles, which are expected to be of interest to the current reader. The user experience with the system can be immensely enhanced with the recommended articles. In this work, we focus on recommendation in Internet forums and blogs with discussion threads.

Here, a fundamental challenge is to account for topic divergence, i.e. the change of gist during the process of discussion. In a discussion thread, the original posting is typically followed by other readers' opinions, in the form of comments. Inten-

<sup>1</sup><http://technorati.com/>

tion and concerns of active users may change as the discussion goes on. Therefore, recommendation, if it were only based on the original posting, can not benefit the potentially evolving interests of the users. Apparently, there is a need to consider topic evolution in adaptive content-based recommendation and this requires novel techniques in order to capture topic evolution precisely and to prevent drastic topic shifting which returns completely irrelevant articles to users.

In this work, we present a framework to recommend relevant information in Internet forums and blogs using user comments, one of the most representative recordings of user behaviors in these forms of social media.

It has the following contributions.

The relevant information is recommended based on a balanced perspective of both the authors and readers.

We model the relationship among comments and that relative to the original posting using graphs in order to evaluate their combined impact. In addition, the weight of a comment is further enhanced with its content and with the authority of its poster.

## 2 Related Work

In a broader context, a related problem is content-based information recommendation (or filtering). Most information recommender systems select articles based on the contents of the original postings. For instance, Chiang and Chen (Chiang and Chen, 2004) study a few classifiers for agent-based news recommendations. The relevant news selections of these work are determined by the textual similarity between the recommended news and the original news posting. A number of later proposals incorporate additional metadata, such as user behaviors and timestamps. For example, Claypool et al. (Claypool et al., 1999) combine the news content with numerical user ratings. Del Corso, Gullí, and Romani (Del Corso et al., 2005) use timestamps to favor more recent news. Cantador, Bellogin, and Castells (Cantador et al., 2008) utilize domain ontology. Lee and Park (Lee and Park, 2007) consider matching between news article attributes and user preferences. Anh et al. (Ahn et al., 2007) and Lai, Liang, and Ku (Lai et al., 2003) construct explicit user profiles, respectively. Lavrenko et al. (Lavrenko et al., 2000) propose

the e-Analyst system which combines news stories with trends in financial time series. Some go even further by ignoring the news contents and only using browsing behaviors of the readers with similar interests (Das et al., 2007).

Another related problem is topic detection and tracking (TDT), i.e. automated categorization of news stories by their themes. TDT consists of breaking the stream of news into individual news stories, monitoring the stories for events that have not been seen before, and categorizing them (Lavrenko and Croft, 2001). A topic is modeled with a language profile deduced by the news. Most existing TDT schemes calculate the similarity between a piece of news and a topic profile to determine its topic relevance (Lavrenko and Croft, 2001) (Yang et al., 1999). Qiu (Qiu et al., 2009) apply TDT techniques to group news for collaborative news recommendation. Some work on TDT takes one step further in that they update the topic profiles as part of the learning process during its operation (Allan et al., 2002) (Leek et al., 2002).

Most recent researches on information recommendation in social media focus on the blogosphere. Various types of user interactions in the blogosphere have been observed. A prominent feature of the blogosphere is the *collective wisdom* (Agarwal et al., 2010). That is, the knowledge in the blogosphere is enriched by such engaging interactions among bloggers and readers as posting, commenting and tagging. Prior to this work, the linking structure and user tagging mechanisms in the blogosphere are the most widely adopted ones to model such collective wisdom. For example, Esmaili et al. (Esmaili et al., 2006) focus on the linking structure among blogs. Hayes, Avesani, and Bojars (Hayes et al., 2007) explore measures based on blog authorship and reader tagging to improve recommendation. Li and Chen further integrate trust, social relation and semantic analysis (Li and Chen, 2009). These approaches attempt to capture accurate similarities between postings without using reader comments. Due to the interactions between bloggers and readers, blog recommendation should not limit its input to only blog postings themselves but also incorporate feedbacks from the readers.

The rest of this article is organized as follows. We first describe the design of our recommendation framework in Section 3. We then evaluate the performance of such a recommender using two




Figure 1: Design scheme

different social media corpora (Section 4). This paper is concluded with speculation on how the current prototype can be further improved in Section 5.

### 3 System Design

In this section, we present a mechanism for recommendation in Internet forums and blogs. The framework is sketched in Figure 1. Essentially, it builds a topic profile for each original posting along with the comments from readers, and uses this profile to retrieve relevant articles. In particular, we first extract structural, semantic, and authority information carried by the comments. Then, with such collective wisdom, we use a graph to model the relationship among comments and that relative to the original posting in order to evaluate the impact of each comment. The graph is weighted with postings’ contents and the authors’ authority. This information along with the original posting and its comments are fed into a synthesizer. The synthesizer balances views from both authors and readers to construct a topic profile to retrieve relevant articles.

#### 3.1 Incorporating Comments

In a discussion thread, comments made at different levels reflect the variation of focus of readers. Therefore, recommended articles should reflect their concerns to complement the author’s opinion. The degree of contribution from each comment, however, is different. In the extreme case, some of them are even advertisements which are completely irrelevant to the discussion topics. In this work, we use a graph model to differentiate the importance of each comment. That is, we model the authority, semantic, structural relations of comments to determine their combined impact.

##### 3.1.1 Authority Scoring Comments

Intuitively, each comment may have a different degree of authority determined by the status of its author (Hu et al., 2007). Assume we have  $n$  users

in a forum, denoted by  $U = \{u_1, \dots, u_n\}$ . We calculate the authority  $A(u_i)$  for user  $u_i$ . To do that, we employ a variant of the PageRank algorithm (Brin and Page, 1998). We consider the cases that a user replies to a previous posting and that a user quotes a previous posting separately. For user  $u_i$ , we use  $r_i(u_j)$  to denote the number of times that  $u_i$  has replied to user  $u_j$ . Similarly, we use  $q_i(u_j)$  to denote the number of times that  $u_i$  has quoted user  $u_j$ . We combine them linearly:

Further, we normalize the above quantity to record how frequently a user refers to another:

Inline with the PageRank algorithm, we define the authority of user  $u_i$  as

##### 3.1.2 Differentiating comments with Semantic and Structural relations

Next, we construct a similar model in terms of the comments themselves. In this model, we treat the original posting and the comments each as a text node. This model considers both the content similarity between text nodes and the logic relationship among them.

On the one hand, the semantic similarity between two nodes can be measured with any commonly adopted metric, such as cosine similarity and Jaccard coefficient (Baeza-Yates and Ribeiro-Neto, 1999). On the other hand, the structural relation between a pair of nodes takes two forms as we have discussed earlier. First, a comment can be made in response to the original posting or at most one earlier comment. In graph theoretic terms, the hierarchy can be represented as a tree  $T = (V, E)$ , where  $V$  is the set of all text nodes and  $E$  is the edge set. In particular, the original posting is the root and all the comments are ordinary nodes. There is an arc (directed edge)  $e_{ij}$  from node  $u_i$  to node  $u_j$ , denoted  $e_{ij}$ , if the corresponding comment  $c_j$  is made in response to comment (or original posting)  $c_i$ . Second, a comment can quote from one or more earlier comments. From this perspective, the hierarchy can be modeled using a directed acyclic graph (DAG),

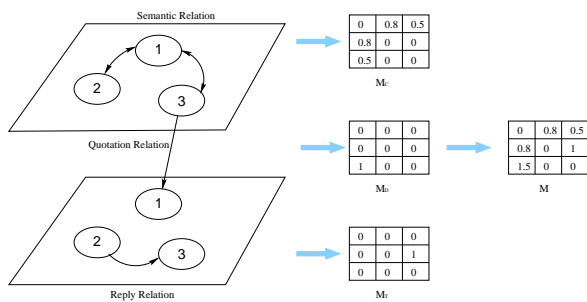


Figure 2: Multi-relation graph of comments based on the structural and semantic information

denoted  $M_s$ . There is an arc from node  $i$  to node  $j$ , denoted  $M_{ij}$ , if the corresponding comment  $i$  quotes comment (or original posting)  $j$ . As shown in Figure 2, for either graph or  $M_s$ , we can use a adjacency matrix, denoted  $M_s$  and  $M_s$ , respectively, to record them. Similarly, we can also use a matrix defined on  $M_q$  to record the content similarity between nodes and denote it by  $M_q$ . Thus, we combine these three aspects linearly:

The importance of a text node can be quantized by the times it has been referred to. Considering the semantic similarity between nodes, we use another variant of the PageRank algorithm to calculate the weight of comment  $i$ :

where  $\alpha$  is a damping factor, and  $\beta$  is the normalized weight of comment  $j$  referring to  $i$  defined as

where  $M_{ij}$  is an entry in the graph adjacency matrix  $M$  and  $\epsilon$  is a constant to avoid division by zero.

In some social networking media, a user may have a subset of other users as “friends”. This can be captured by a matrix of  $F$ , whose entries are denoted by  $F_{ij}$ . Thus, with this information and assuming poster  $k$  has made a comment  $k$  for user  $i$ ’s posting, the final weight of this comment is defined as

### 3.2 Topic Profile Construction

Once the weight of comments on one posting is quantified by our models, this information along with the entire discussion thread is fed into a synthesizer to construct a topic profile. As such, the perspectives of both authors and readers are balanced for recommendation.

The profile is a weight vector of terms to model the language used in the discussion thread. Consider a posting  $i$  and its comment sequence  $C_i$ . For each term  $t$ , a compound

weight  $w_t$  is calculated. It is a linear combination of the contribution by the posting itself,  $w_{t,i}$ , and that by the comments,  $w_{t,C_i}$ . We assume that each term is associated with an “inverted document frequency”, denoted by  $\text{idf}(t)$ , where  $N$  is the corpus size and  $n_t$  is the number of documents in corpus containing term  $t$ . We use a function  $\text{tf}(t, d)$  to denote the number of occurrences of term  $t$  in document  $d$ , i.e. “term frequency”. Thus, when the original posting and comments are each considered as a document, this term frequency can be calculated for any term in any document. We thus define the weight of term  $t$  in document  $d$ , be the posting itself or a comment, using the standard TF/IDF definition (Baeza-Yates and Ribeiro-Neto, 1999):

The weight contributed by the posting itself,  $w_{t,i}$ , is thus:

The weight contribution from the comments  $w_{t,C_i}$  incorporates not only the language features of these documents but also their importance in the discussion thread. That is, the contribution of comment score is incorporated into weight calculation of the words in a comment.

Such a treatment of compounded weight is essentially to recognize that readers’ impact on selecting relevant articles and the difference of their influence. For each profile, we select the top-highest weighted words to represent the topic.

With the topic profile thus constructed, the retriever returns an ordered list of articles with decreasing relevance to the topic. Note that our approach to differentiate the importance of each comment can be easily incorporated into any generic retrieval model. In this work, our retriever is adopted from (Lavrenko et al., 2000).

### 3.3 Interpretation of Recommendation

Since interpreting recommended items enhances users' trusting beliefs (Wang and Benbasat, 2007), we design a creative approach to generate hints to indicate the relationship (generalization, specialization and duplication) between the recommended articles and the original posting based on our previous work (Candan et al., 2009).

Article  $A$  being more general than  $B$  can be interpreted as  $A$  being less constrained than  $B$  by the keywords they contain. Let us consider two articles,  $A$  and  $B$ , where  $A$  contains keywords,  $k_1$  and  $k_2$ , and  $B$  only contains  $k_1$ .

If  $A$  is said to be more general than  $B$ , then the additional keyword,  $k_2$ , of article  $A$  must render  $A$  less constrained than  $B$ . Therefore, the content of  $A$  can be interpreted as  $A \supset B$ .

If, on the other hand,  $B$  is said to be more specific than  $A$ , then the additional keyword,  $k_2$ , must render  $B$  more constrained than  $A$ . Therefore, the content of  $B$  can be interpreted as  $B \subset A$ .

Note that, in the two-keyword space  $(k_1, k_2)$ , can be denoted by a vector  $(w_1, w_2)$  and  $B$  can be denoted by  $(w_1, 0)$ . The origin  $(0, 0)$  corresponds to the case where an article does contain neither  $k_1$  nor  $k_2$ . That is,  $(0, 0)$  corresponds to an article which can be interpreted as  $A = B$ . Therefore, if  $A$  is said to be more general than  $B$ ,  $\frac{w_2}{w_1}$  should be greater than  $\frac{0}{w_1}$ . This allows us to measure the degrees of generalization and specialization of two articles. Given two articles,  $A$  and  $B$ , of the same topic, they will have a common keyword base, while both articles will also have their own content, different from their common base. Let us denote the common part of  $A$  by  $C$  and common part of  $B$  by  $D$ . Note that  $C$  and  $D$  are usually unequal because the same words in the common part have different term weights in article  $A$  and  $B$  respectively. Given these and the generalization concept introduced above for two similar

articles  $A$  and  $B$ , we can define the degree of generalization ( $G$ ) and specialization ( $S$ ) of  $A$  with respect to  $B$  as

To alleviate the effect of document length, we revise the definition as

$$G = \frac{|C|}{|A|} \quad S = \frac{|D|}{|B|}$$

The relative specialization and generalization values can be used to reveal the relationships between recommended articles and the original posting. Given original posting  $A$  and recommended article  $B$ , if  $G > S$ , for a given generalization threshold  $T$ , then  $B$  is marked as a generalization. When this is not the case, if  $S > T$ , for a given specialization threshold,  $T$ , then  $B$  is marked as a specialization. If neither of these cases is true, then  $B$  is duplicate of  $A$ .

Such an interpretation provides a control on delivering recommended articles. In particular, we can filter the duplicate articles to avoid recommending the same information.

## 4 Experimental Evaluation

To evaluate the effectiveness of our proposed recommendation mechanism, we carry out a series of experiments on two synthetic data sets, collected from Internet forums and blogs, respectively. The first data set is called Forum. This data set is constructed by randomly selecting 20 news articles with corresponding reader comments from the Digg Web site and 16,718 news articles from the Reuters news Web site. This simulates the scenario of recommending relevant news from traditional media to social media users for their further reading. The second one is the Blog data set containing 15 blog articles with user comments and 15,110 articles obtained from the Myhome Web site<sup>2</sup>. Details of these two data sets are shown in Table 1. For evaluation purposes, we adopt the traditional pooling strategy (Zobel, 1998) and apply to the TREC data set to mark the relevant articles for each topic.

<sup>2</sup><http://blogs.myhome.ie>

Table 1: Evaluation data set

Synthetic Data Set		Forum	Blog
Topics	No. of postings	20	15
	Ave. length of postings	676	236
	No. of comments per posting	81.4	46
	Ave. length of comments	45	150
Target	No. of articles	16718	15110
	Ave. length of articles	583	317

The recommendation engine may return a set of essentially the same articles re-posted at different sites. Therefore, we introduce a metric of novelty to measure the topic diversity of returned suggestions. In our experiments, we define precision and novelty metrics as

$$\text{Precision} = \frac{|S \cap T|}{|S|} \quad \text{and} \quad \text{Novelty} = \frac{|S \cap T|}{|T|}$$

where  $S$  is the subset of the top- $k$  articles returned by the recommender,  $T$  is the set of manually tagged relevant articles, and  $U$  is the set of manually tagged relevant articles excluding duplicate ones to the original posting. We select the top 10 articles for evaluation assuming most readers only browse up to 10 recommended articles (Karypis, 2001). Meanwhile, we also utilize mean average precision (MAP) and mean average novelty (MAN) to evaluate the entire set of returned article.

We test our proposal in four aspects. First, we compare our work to two baseline works. We then present results for some preliminary tests to find out the optimal values for two critical parameters. Next, we study the effect of user authority and its integration to comment weighting. Fourth, we evaluate the performance gain obtained from interpreting recommendation. In addition, we provide a significance test to show that the observed differences in effectiveness for different approaches are not incidental. In particular, we use the  $\chi^2$ -test here, which is commonly used for significance tests in information retrieval experiments (Hull, 1993).

#### 4.1 Overall Performance

As baseline proposals, we also implement two well-known content-based recommendation methods (Bogers and Bosch, 2007). The first method, Okapi, is commonly applied as a representative of the classic probabilistic model for relevant information retrieval (Robertson and Walker, 1994). The second one, LM, is based on statistical language models for relevant information retrieval (Ponte and Croft, 1998). It builds a proba-

Table 2: Overall performance

Data	Method	Precision		Novelty	
Forum	<b>Okapi</b>	0.827	0.833	0.807	0.751
	<b>LM</b>	0.804	0.833	0.807	0.731
	<b>Our</b>	0.967	0.967	0.9	0.85
Blog	<b>Okapi</b>	0.733	0.651	0.667	0.466
	<b>LM</b>	0.767	0.718	0.70	0.524
	<b>Our</b>	0.933	0.894	0.867	0.756

bilistic language model for each article, and ranks them on query likelihood, i.e. the probability of the model generating the query. Following the strategy of Bogers and Bosch, relevant articles are selected based on the title and the first 10 sentences of the original postings. This is because articles are organized in the so-called *inverted pyramid* style, meaning that the most important information is usually placed at the beginning. Trimming the rest of an article would usually remove relatively less crucial information, which speeds up the recommendation process.

A paired  $t$ -test shows that using  $\chi^2$  and  $\text{MAP}$  as performance measures, our approach performs significantly better than the baseline methods for both Forum and Blog data sets as shown in Table 2. In addition, we conduct  $\chi^2$ -tests using MAP and MAN as performance measures, respectively, and the  $p$ -values of these tests are all less than 0.05, meaning that the results of experiments are statistically significant. We believe that such gains are introduced by the additional information from the collective wisdom, i.e. user authority and comments. Note that the retrieval precision for Blog of two baseline methods is not as good as that for Forum. Our explanation is that blog articles may not be organized in the inverted pyramid style as strictly as news forum articles.

#### 4.2 Parameters of Topic Profile

There are two important parameters to be considered to construct topic profiles for recommendation. 1) the number of the most weighted words to represent the topic, and 2) combination coefficient  $\alpha$  to determine the contribution of original posting and comments in selecting relevant articles. We conduct a series of experiments and find out that the optimal performance is obtained when the number of words is between 50 and 70, and  $\alpha$  is between 0.65 and 0.75. When  $\alpha$  is set to 0, the recommended articles only reflect the author’s opinion. When  $\alpha = 1$ , the suggested articles represent the concerns of readers exclusively. In the

Table 3: Performance of four runs

	Method	Precision		Novelty	
Forum	RUN1	0.88	0.869	0.853	0.794
	RUN2	0.933	0.911	0.9	0.814
	RUN3	0.94	0.932	0.9	0.848
	RUN4	0.967	0.967	0.9	0.85
Blog	RUN1	0.767	0.758	0.7	0.574
	RUN2	0.867	0.828	0.833	0.739
	RUN3	0.9	0.858	0.833	0.728
	RUN4	0.933	0.894	0.867	0.756

following experiments, we set topic word number to 60 and combination coefficient to 0.7.

### 4.3 Effect of Authority and Comments

In this part, we explore the contribution of user authority and comments in social media recommender. In particular, we study the following scenarios with increasing system capabilities. Note that, lacking friend information (Section 3.1.2) in the Forum data set, is set to zero.

RUN 1 (Posting): the topic profile is constructed only based on the original posting itself. This is analogous to traditional recommenders which only consider the focus of authors for suggesting further readings.

RUN 2 (Posting+Authority): the topic profile is constructed based on the original posting and participant authority.

RUN 3 (Posting+Comment): the topic profile is constructed based on the original posting and its comments.

RUN 4 (All): the topic profile is constructed based on the original posting, user authority, and its comments.

Here, we set . Our -test shows that using and as performance measures, RUN4 performs best in both Forum and Blog data sets as shown in Table 3. There is a step-wise performance improvement while integrating user authority, comments and both. With the assistance of user authority and comments, the recommendation precision is improved up to 9.8% and 21.6% for Forum and Blog, respectively. The opinion of readers is an effective complementarity to the authors' view in suggesting relevant information for further reading.

Moreover, we investigate the effect of the semantic and structural relations among comments, i.e. semantic similarity, reply, and quotation. For this purpose, we carry out a series of experiments based on different combinations of these relations.




Figure 3: Effect of content, quotation and reply relation

Content Relation (CR): only the content relation matrix is used in scoring the comments.

Quotation Relation (QR): only the quotation relation matrix is used in scoring the comments.

Reply Relation (RR): only the reply relation matrix is used in scoring the comments.

Content+Quotation Relation (CQR): both the content and quotation relation matrices is used in scoring the comments.

Content+Reply Relation(CRR): both the content and reply relation matrices are used in scoring the comments.

Quotation+Reply Relation (QRR): both the quotation and reply relation matrices are used in scoring the comments.

All: all three matrices are used.

The MAP yielded by these combinations for both data sets is plotted in Figure 3. For the case of Forum, we observe that incorporating content information adversely affects recommendation precision. This concurs with what we saw in our previous work (Wang et al., 2010). On the other hand, when we test the Blog data set, the trend is the opposite, i.e. content similarity does contribute to retrieval performance positively. This is attributed by the text characteristics of these two forms of social media. Specifically, comments in news forums usually carry much richer structural information than blogs where comments are usually “flat” among themselves.

### 4.4 Recommendation Interpretation

To evaluate the precision of interpreting the relationship between recommended articles and the

original posting, the evaluation metric of success rate is defined as

where  $n$  is the number of recommended articles,  $w$  is the error weight of recommended article  $i$ . Here, the error weight is set to one if the result interpretation is mis-labelled.

From our studies, we observe that the success rate at top-10 is around 89.3% and 87.5% for the Forum and Blog data sets, respectively. Note that these rates include the errors introduced by the irrelevant articles returned by the retrieval module. To estimate optimal thresholds of generalization and specialization, we calculate the success rate at different threshold values and find that neither too small nor too large a value is appropriate for interpretation. In our experiments, we set generalization threshold to 3.2 and specialization threshold to 1.8 for the Forum data set, and to 3.5 and to 2.0 for Blog. Ideally, threshold values would need to be set through a machine learning process, which identifies proper values based on a given training sample.

## 5 Conclusion and Future Work

The Web has become a platform for social networking, in addition to information dissemination at its earlier stage. Many of its applications are also being extended in this fashion. Traditional recommendation is essentially a push service to provide information according to the profile of individual or groups of users. Its niche at the Web 2.0 era lies in its ability to enable online discussion by serving up relevant references to the participants. In this work, we present a framework for information recommendation in such social media as Internet forums and blogs. This model incorporates information of user status and comment semantics and structures within the entire discussion thread. This framework models the logic connections among readers and the innovativeness of comments. By combining such information with traditional statistical language models, it is capable of suggesting relevant articles that meet the dynamic nature of a discussion in social media. One important discovery from this work is that, when integrating comment contents, the structural information among comments, and reader relationship, it is crucial to distinguish the characteristics of various forms of social media. The reason is that the

role that the semantic content of a comment plays can differ from one form to another.

This study can be extended in a few interesting ways. For example, we can also evaluate its effectiveness and costs during the operation of a discussion forum, where the discussion thread is continually updated by new comments and votes. Indeed, its power is yet to be further improved and investigated.

## Acknowledgments

Li's research is supported by National Natural Science Foundation of China (Grant No.60803106), the Scientific Research Foundation for the Returned Overseas Chinese Scholars, State Education Ministry, and the Fok Ying-Tong Education Foundation for Young Teachers in the Higher Education Institutions of China. Research of Chen is supported by Natural Science and Engineering Council (NSERC) of Canada.

## References

- Nitin Agarwal, Magdiel Galan, Huan Liu, and Shankar Subramanya. 2010. Wiscoll: Collective wisdom based blog clustering. *Information Sciences*, 180(1):39–61.
- Jae-wook Ahn, Peter Brusilovsky, Jonathan Grady, Daqing He, and Sue Yeon Syn. 2007. Open user profiles for adaptive news systems: help or harm? In *Proceedings of the 16th International Conference on World Wide Web (WWW)*, pages 11–20.
- James Allan, Victor Lavrenko, and Russell Swan. 2002. Explorations within topic tracking and detection. Topic detection and tracking: event-based information organization *Kluwer Academic Publishers*, pages 197–224.
- Ricardo Baeza-Yates and Berthier Ribeiro-Neto. 1999. Modern information retrieval. *Addison Wesley Longman Publisher*.
- Toine Bogers and Antal Bosch. 2007. Comparing and evaluating information retrieval algorithms for news recommendation. In *Proceedings of 2007 ACM conference on Recommender Systems*, pages 141–144.
- Sergey Brin and Lawrence Page. 1998. The anatomy of a large-scale hypertextual web search engine. *Computer networks and ISDN systems*, 30(1-7):107–117.
- K. Selçuk Candan, Mehmet E. Dönderler, Terri Hedgpeth, Jong Wook Kim, Qing Li, and Maria Luisa Sapino. 2009. SEA: Segment-enrich-annotate paradigm for adapting dialog-based content for improved accessibility. *ACM Transactions on Information Systems (TOIS)*, 27(3):1–45.



- Ivan Cantador, Alejandro Bellogin, and Pablo Castells. 2008. Ontology-based personalized and context-aware recommendations of news items. In *Proceedings of IEEE/WIC/ACM international Conference on Web Intelligence and Intelligent Agent Technology (WI)*, pages 562–565.
- Jung-Hsien Chiang and Yan-Cheng Chen. 2004. An intelligent news recommender agent for filtering and categorizing large volumes of text corpus. *International Journal of Intelligent Systems*, 19(3):201–216.
- Mark Claypool, Anuja Gokhale, Tim Miranda, Pavel Murnikov, Dmitry Netes, and Matthew Sartin. 1999. Combining content-based and collaborative filters in an online newspaper. In *Proceedings of the ACM SIGIR Workshop on Recommender Systems*.
- Abhinandan S. Das, Mayur Datar, Ashutosh Garg, and Shyam Rajaram. 2007. Google news personalization: scalable online collaborative filtering. In *Proceedings of the 16th International Conference on World Wide Web (WWW)*, pages 271–280.
- Gianna M. Del Corso, Antonio Gullí, and Francesco Romani. 2005. Ranking a stream of news. In *Proceedings of the 14th International Conference on World Wide Web (WWW)*, pages 97–106.
- Kyumars Sheykh Esmaili, Mahmood Neshati, Mohsen Jamali, Hassan Abolhassani, and Jafar Habibi. 2006. Comparing performance of recommendation techniques in the blogosphere. In *ECAI 2006 Workshop on Recommender Systems*.
- Conor Hayes, Paolo Avesani, and Uldis Bojars. 2007. An analysis of bloggers, topics and tags for a blog recommender system. In *Workshop on Web Mining (WebMine)*, pages 1–20.
- Meishan Hu, Aixin Sun, and Ee-Peng Lim. 2007. Comments-oriented blog summarization by sentence extraction. In *Proceedings of the sixteenth ACM Conference on Conference on Information and Knowledge Management (CIKM)*, pages 901–904.
- David Hull. 1993. Using statistical testing in the evaluation of retrieval experiments. In *Proceedings of the 16th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 329–338.
- George Karypis. 2001. Evaluation of item-based Top-N recommendation algorithms. In *Proceedings of the 10th International Conference on Information and Knowledge Management (CIKM)*, pages 247–254.
- Hung-Jen Lai, Ting-Peng Liang, and Yi Cheng Ku. 2003. Customized internet news services based on customer profiles. In *Proceedings of the 5th International Conference on Electronic commerce (ICEC)*, pages 225–229.
- Victor Lavrenko and W. Bruce Croft. 2001. Relevance based language models. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 120–127.
- Victor Lavrenko, Matt Schmill, Dawn Lawrie, Paul Ogilvie, David Jensen, and James Allan. 2000. Language models for financial news recommendation. In *Proceedings of the 9th International Conference on Information and Knowledge Management (CIKM)*, pages 389–396.
- Hong Joo Lee and Sung Joo Park. 2007. MONERS: A news recommender for the mobile web. *Expert Systems with Applications*, 32(1):143–150.
- Tim Leek, Richard Schwartz, and Srinivasa Sista. 2002. Probabilistic approaches to topic detection and tracking. *Topic detection and tracking: event-based information organization*, pages 67–83.
- Yung-Ming Li and Ching-Wen Chen. 2009. A synthetic approach for blog recommendation: Combining trust, social relation, and semantic analysis. *Expert Systems with Applications*, 36(3):6536 – 6547.
- Jay Michael Ponte and William Bruce Croft. 1998. A language modeling approach to information retrieval. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 275–281.
- Jing Qiu, Lejian Liao, and Peng Li. 2009. News recommender system based on topic detection and tracking. In *Proceedings of the 4th Rough Sets and Knowledge Technology*.
- Stephen E. Robertson and Stephen G Walker. 1994. Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval. In *Proceedings of the 17th ACM SIGIR conference on Research and Development in Information Retrieval*, pages 232–241.
- Weiquan Wang and Izak Benbasat. 2007. Recommendation agents for electronic commerce: Effects of explanation facilities on trusting beliefs. *Journal of Management Information Systems*, 23(4):217–246.
- Jia Wang, Qing Li, and Yuanzhu Peter Chen. 2010. User comments for news recommendation in social media. In *Proceedings of the 33rd ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 295–296.
- Yiming Yang, Jaime Guillermo Carbonell, Ralf D. Brown, Thomas Pierce, Brian T. Archibald, and Xin Liu. 1999. Learning approaches for detecting and tracking news events. *IEEE Intelligent Systems*, 14(4):32–43.
- Justin Zobel. 1998. How reliable are the results of large-scale information retrieval experiments? In *Proceedings of the 21st International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 307–314.