Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)

News Recommendation in Forum-Based Social Media

Jia Wang, Qing Li *
Southwestern University of Finance and Economics, China

Yuanzhu Peter Chen
Memorial University of Newfoundland, Canada

Jiafen Liu, Chen Zhang, Zhangxi Lin†
Southwestern University of Finance and Economics, China

Abstract

Self-publication of news on Web sites is becoming a common application platform to enable more engaging interaction among users. Discussion in the form of comments following news postings can be effectively facilitated if the service provider can recommend articles based on not only the original news itself but also the thread of changing comments. This turns the traditional news recommendation to a "discussion moderator" that can intelligently assist online forums. In this work, we present a framework to implement such adaptive news recommendation. In addition, to alleviate the problem of recommending essentially identical articles, the relationship (*duplication*, *generalization* or *specialization*) between suggested news articles and the original posting is investigated. Experiments indicate that our proposed solutions provide an enhanced news recommendation service in forum-based social media.

Introduction

In its brief two-decade history, the Web has evolved from a technical framework for information dissemination to a social interaction enabler among its users. Nowadays, Web is one of the most important vehicles for "social media", e.g. Internet forums, blogs, wikis, podcasts, instant messaging, and social networking. One form of social media of particular interest here is *self-publishing*, or *customer-generated media*. In self-publishing, a user can publish an article or post news to share with other users. Other users can read and comment on the posting and these comments can, in turn, be read and commented on. Digg (www.digg.com) and Yahoo!Buzz (buzz.yahoo.com) are commercial examples of self-publishing.

A useful extension of this self-publishing application is to add a recommendation feature to the current discussion thread. That is, based on the original posting and various levels of comments, the system can provide a set of relevant articles, which are expected to be of interest of the active users of the thread. The users learning experience with the system can be immensely enhanced with the recommended articles. Due to the evolving nature of the thread and the inherent small sizes of the comments, such contextual recommendation will be different from traditional news recommendation to a large extent.

Here, we explore the problem of news recommendation for dynamic discussion threads. A fundamental challenge in adaptive news recommendation is to account for topic divergence, i.e. the change of gist during the process of discussion. In a forum, the original news is typically followed by other readers' opinions, in the form of comments. Concerns and intention of active users may change as the discussion continues. Therefore, news recommendation, if it were only based on the original posting, can not benefit the potentially changing interests of the users. Apparently, there is a need to consider topic evolution in adaptive news recommendation and this requires novel techniques that can help to capture topic evolution precisely to prevent wild topic shifting which returns completely irrelevant news to users. In this work, we propose a framework of adaptive news recommendation in social media. It has the following contributions.

- It is the first attempt of incorporating reader comments for adaptive news recommendation. We model the relationship among comments and that relative to the original posting in order to evaluate their overall impact.
- We determine the relationship (*generalization* or *specialization*) between suggested news articles and the original posting and present it to users via a novel user interface.

In a broader context, a related problem is content-based information filtering (or recommendation). Most information recommender systems select articles based solely on the contents of the original postings. For instance, Lavrenko et al. (2000) propose the e-Analyst system utilizing language models to recommend interesting news stories that are likely to affect the market behavior. Chiang and Chen (2004) study several classification approaches 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 information, such as user behaviors and timestamps. For example, Claypool et al. (1999) combine the news content with numerical user ratings. Del Corso et al. (2005) use times-

*This research is supported by National Natural Science Foundation of China Grant No.60803106.

†Also at Texas Tech University, USA.

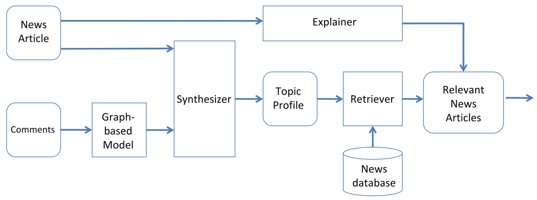

Figure 1: Design scheme

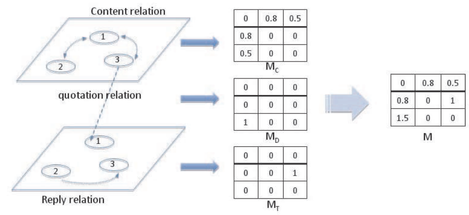

Figure 2: Multi-relation graph of comments

tamps to favor more recent news. Cantador et al. (2008) utilize domain ontology. Lee and Park (2007) consider matching between news article attributes and user preferences. Anh et al. (2007) and Lai et al. (2003) construct explicit user profiles, respectively. Some go even further by ignoring the news contents and only using browsing behaviors of the readers with similar interests (Das et al. 2007).

A common issue with traditional news recommendation is result redundancy. That is, the recommendation engine may return a set of essentially the same articles re-posted at different sites. In forum-oriented recommendation, we argue that the novelty of recommendation returns is crucial. In addition, due to the sheer diversity and the number of the recommended articles, the users will feel much easier navigating them even if we simply categorize them as generalization or specialization of the original article.

System Design

In this section, we present a mechanism for adaptive news recommendation. The framework is sketched in Figure 1. Essentially, it builds a topic profile for each original news posting along with the comments from readers, and uses this profile to retrieve relevant news. In particular, we first model the relationship among comments and that relative to the original posting in order to evaluate their overall impact. This information along with the news and comments are fed into a synthesizer. The synthesizer balances views of both authors and readers to construct a topic profile to retrieve relevant news. These recommended news articles are further analyzed by an explainer to reveal their relationship to the original posting.

Incorporating Comments

In a discussion thread, comments made at different levels reflect the variation of focus of its participants. Therefore, comments should be incorporated when building recommendation models. In our model, we treat the original posting and the comments each as a text node. This model both considers the content similarity between text nodes and the logic relationship among them.

On one hand, the content similarity between two nodes can be measured by any commonly adopted metric, such as cosine similarity and Jaccard coefficient. This metric is taken over every node pair in the discussion thread. On the other hand, the logic relation between nodes takes two forms. First, a comment is always made in response to the original posting or an earlier comment. In graph theoretic terms, the hierarchy can be represented as a tree $T = (V, E_T)$, where V is the set of all text nodes and E_T is the edge set. In particular, the original posting is the root and all the comments are ordinary nodes. There is a directed edge $e \in E_T$ from node u to node v, denoted (u, v), if the corresponding comment v is made in response to comment (or original posting) u. 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), denoted $D = (V, E_D)$. There is a directed edge $e \in E_D$ from node u to node v, denoted (u, v), if the corresponding comment v quotes from comment (or original posting) u. As shown in Figure 2, for either graph T or D, we can use a $|V| \times |V|$ adjacency matrix, denoted M_T and M_D, respectively, to record them. Inline with the adjacency matrices, we can also use a $|V| \times |V|$ matrix defined on $[0, 1]$ to record the content similarity between nodes and denote it by M_C. Thus, we can combine these three aspects linearly:

$$M = \gamma_1 \times M_C + \gamma_2 \times M_T + \gamma_3 \times M_D.$$

where M is a $|V| \times |V|$ adjacency matrix capturing both content and logic relations among text nodes.

Intuitively, the important comments are those whose topics are discussed by a large number of other important comments. Therefore, we propose to apply the PageRank algorithm (Brin and Page 1998) to rank the comments as

$$s_j = \lambda/|V| + (1 - \lambda) \times \sum_{c_i} r(c_i, c_j) \times s_i,$$

where λ is the damping factor as in PageRank and this value is recommended to be 0.85, i and j are node indices, and $|V|$ denotes the number of text nodes in the thread. In addition, $r(c_i, c_j)$ is the normalized weight of comment c_i referring to c_j defined as

$$r(c_i, c_j) = \frac{M_{i,j}}{\sum_{c_j} M_{i,j} + \epsilon},$$

where $M_{i,j}$ is an entry in the graph adjacency matrix and ϵ is a constant to avoid division by zero.

Topic Profile Construction

Once the importance of comments on one news posting is quantified by our model, this information along with the news itself are fed into a synthesizer to construct a topic profile of this news discussion thread. As way, the views of both authors and readers are balanced for news recommendation.

The profile is a weight vector of terms to model the language used in the thread. Consider a news posting d_0 and its comment sequence $\{d_1, d_2, \cdots, d_m\}$. For each term t, a compound weight $W(t) = \alpha \times W_1(t) + (1-\alpha) \times W_2(t)$ is calculated. It is a linear combination of the contribution by the news posting itself, $W_1(t)$, and that by the comments, $W_2(t)$. We assume that each term is associated with an "inverted document frequency", denoted by $I(t) = \log \frac{N}{n(t)}$, where N is the corpus size and $n(t)$ is the number of documents in corpus containing term t. We use a function $f(t,d)$ to denote the number of occurrences of term t in document d, i.e. "term frequency". Thus, when news posting and comments are each considered as a document, this term frequency value can be calculated for any term in any document. We thus define the weight of term t in document d, be the news posting itself or a comment, using the standard TF/IDF definition (Ricardo et al. 1999):

$$w(t,d) = \Big(0.5 + 0.5 \times f(t,d)/\max_{t'} f(t',d)\Big) \times I(t)$$

The weight contributed by the news itself, d_0, is thus:

$$W_1(t) = w(t,d_0)/\max_{t'} w(t',d_0)$$

The weight contribution from the comments $\{d_1, d_2, \cdots, d_m\}$ incorporates not only the language features of these documents but also their importance of leading a discussion in related topics. That is, the contribution of comment score is incorporated into weight calculation of the words in a text node.

$$W_2(t) = \sum_{i=1}^{m} w(t,d_i)/\max_{t'} w(t',d_i) \times s_i/\max_{i'} s_{i'}$$

Such a treatment of compounded weight $W(t)$ is essentially to recognize that readers' impact on selecting relevant news and the difference of their influence strength. For each profile, we select top-N high weighted words to represent the topic.

News Retrieval with Relevance Language Models

With the topic profile constructed as above, we can use it to select relevant news for recommendations. That is, the retriever returns an order list of news with decreasing relevance to the topic. Our model to differentiate the importance of each comment can be easily incorporated into any good retrieval model. In this work, our retrieval model is derived from (Lavrenko et al. 2000). Specifically, the score of the candidate document d with regard to profile p is defined as

$$s(p,d) = (p \cdot d)$$

It is the similarity of the language vectors can be measured by those commonly adopted metrics, such as cosine similarity and Kullback-Leibler (KL) divergence. In this work, to measure the content similarity, $(p \cdot d)$, we take the relevance language model approach of Lavrenko et al. (2001). The similarity between a news article and a topic profile is measured by the KL divergence between the document model and the topic model. Articles with smaller divergence are considered more relevant. The divergence is defined as:

$$(p \cdot d) \approx KL(M_p \parallel M_d) = \sum_{w} P(w|M_p) \log \frac{P(w|M_p)}{P(w|M_d)}$$

where M_d is a language model for document d in the collection, which is a probability distribution that captures the statistical regularities of natural language use. M_p is a language model for topic p. $P(w|M_p)$ can be estimated as :

$$P(w|M_p) \approx \frac{P(w, m_1 \cdots m_k)}{P(m_1 \cdots m_k)} = \frac{P(w, m_1 \cdots m_k)}{\sum\limits_{v \in vocabulary} P(v, m_1 \cdots m_k)}$$

$$P(w, m_1 \cdots m_k) = \sum_{d \in D_r} P(M_d) \left[P(w|M_d) \prod_{i=1}^{k} P(m_i|M_d) \right]$$

where m_i denotes i-th word of the topic profile. $P(M_d)$ denotes some prior distribution over the relevant document set D_r which is usually taken to be uniform. To accelerate execution, we further restrict to only contain 50 documents corresponding to the top-ranked documents retrieved by the topic profile. $P(w|M_d)$ specifies the probability of observing w if we sample a random word from M_d. $P(w|M_d)$ is estimated by

$$P(w|M_d) = \eta f_{w,d}/|d| + (1-\eta) f_w/F$$

where, $f_{w,d}$ is the number of times word w occurs in document d, f_w is the number of times word w in the corpus. F is the total number of tokens (i.e. words with repetition) in the corpus, $|d|$ is the number of tokens in d. Note that we use a parameter η to control the contribution of the term frequency to this probability. $P(m_i|M_d)$ is estimated in the same way.

With such a scoring mechanism, the retriever can return a number of most relevant news with regard to any evolving discussion in a self-publication system.

Explanation for Recommendation

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

News 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 news 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 news article A must render A less constrained than B. Therefore, the content of A can be interpreted as $k_1 \cup k_2$.
- If, on the other hand, A is said to be more specific than B, then the additional keyword, k_2, must render A more constrained than B. Therefore, the content of A can be interpreted as $k_1 \cap k_2$.

Note that, in the two-keyword space $\langle k_1, k_2 \rangle$, A can be represented by a vector $\langle a_A, b_A \rangle$ and B can be represented by $\langle a_B, 0 \rangle$. The extreme point $O = \langle 0, 0 \rangle$ corresponds to the case where a news article does contain neither k_1 nor k_2; in other words, O corresponds to a news article which can be interpreted as $\neg k_1 \cap \neg k_2 \equiv \neg (k_1 \cup k_2)$. Therefore, if A

Figure 3: User interface of recommending items

is said to be more general than B, $\Delta A = dist(A, O)$ should be greater than $\Delta B = dist(B, O)$. This gives a way to measure the degrees of generalization and specialization of two news articles. Given two news articles, A and B, *of the same topic*, they will have a common keyword base, while both news articles will also have their own content, different from their common base. Let us denote the common part of A with A^c and common part of B with B^c. Note that ΔA^C and ΔB^C are usually unequal because the same words in the common part have different term weights in news A and B respectively. Given these and the generalization concept introduced above for two similar news A and B, we could define the degree of generalization (G_{AB}) and specialization (S_{AB}) of B with respect to A as

$$G_{ab} = \Delta A/\Delta B^c, S_{ab} = \Delta B/\Delta A^c$$

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

$$G_{AB} = \frac{\Delta A/\log(\Delta A)}{\Delta B^c/\log(\Delta A + \Delta B)}, S_{AB} = \frac{\Delta B/\log(\Delta B)}{\Delta A^c/\log(\Delta A + \Delta B)}$$

The relative specialization and generalization values can be used to reveal the relationships of recommended news articles and the original posting. Given original news posting A and recommended news article B, if $G_{AB} > \Theta_g$, for a given generalization threshold, Θ_g, then B is marked as a generalization. When this is not the case, if $S_{AB} > \Theta_s$, for a given specialization threshold, Θ_s, then B is marked as a specialization. If neither of these cases is true, then B is duplicated to A.

Such kinds of explanation provide a free control on delivering recommended news. In particular, we can block these duplicated news to avoid recommending the same information without any dissimilarity.

Visualization of Recommended News

In our system design, there are several uncertain factors such as the weight α. While they can be set experimentally, we design a novel user interface to give readers more control. The efficiency and effectiveness of this UI is being evaluated in a separate work. We here only present the final design as shown in Figure 3. The recommendation-result interface consists of two parts:

- Recommended news bubble pool — It contains a large number of candidate news articles to be recommended. Each bubble represents one news article. By clicking on a bubble, a reader will be directed to the corresponding article. The size and color of a bubble indicates the rank of this article in the recommendation list.
- Display control panel — There are three sliding controls and two selection controls. The sliding controls allow the user to select different parameter values for α, Θ_g, and Θ_s. The selection controls provide options of filtering news based on different categories and news sources, respectively.

Table 1: Evaluation data set

Synthetic Data set		
Topics	No. of Postings	20
	Ave. Length of Postings	676
	No. of Comment per Posting	81.4
	Ave. Length of Comments	45
News	No. of News Articles	16,718
	Ave. Length of News	583

Experimental Evaluation

To gauge how well the proposed recommendation approach performs, we carry out a series of experiments on a synthetic data set collected from Digg and Reuters news website. We randomly select 20 news articles with corresponding reader comments from Digg website. These news articles with different topics are treated as the original news postings, recommended news are selected from a corpus of articles collected from Reuters news website. This simulates the scenario of recommending relevant news from traditional media to social media readers for their further reading. For evaluation purpose, we adopt the traditional pooling strategy (Zobel 1998) applied in TREC dataset to mark the relevant articles for each topic. In particular, we use three different retrieval engines to produce a pool of potentially relevant news articles. For each posting along with comments, the rankings from the three retrieval engines are merged into one ranking list. Three judgers go through the top 200 results of this list to evaluate their relationships (generalization, specialization and duplication) to the original posting. The non-judged documents are considered irrelevant. Details of the dataset are shown in Table 1.

In addition to the proposed news recommender, we also implement two well-known news recommendation methods as the baseline (Bogers and Bosch 2007). The first method, Okapi, is commonly applied as a representative of the classical probabilistic model for relevant information retrieval (Robertson and Walker 1994). The second method, LM, is based on statistical language models for relevant information retrieval (Ponte and Croft 1998). Following the strategy of Bogers (Bogers and Bosch 2007), relevant news articles are selected based on the title and first ten sentences of the original postings. This is because news articles are organized in so-called inverted pyramid style, meaning that the most important information is usually at the beginning.

Traditionally, the quality of news recommendation schemes is evaluated by a precision metric. The recommendation redundancy, however, greatly bothers users. Therefore, we argue that a novelty metric is as important for effectiveness evaluation of news recommendations. We group

Table 2: Overall performance

	Precision		Novelty	
Method	P@10	MAP	D@10	MAN
Okapi	0.827	0.833	0.807	0.751
LM	0.804	0.833	0.807	0.731
The Proposed	0.94	0.932	0.9	0.848

relevant news into three subcategories (generalization, specialization and duplication) based on the relationship to the original posting. Thus, we apply precision and novelty metrics in our experiments as

$$P@N = \frac{|C \cap R|}{|R|}, \quad D@N = \frac{|E \cap R|}{|R|},$$

where, R is the subset of the top-n news articles returned by the recommender, C is the set of tagged relevant news, and E is the set of manually tagged relevant news excluding duplicated ones to the original news posting. We select top 10 news articles for evaluation assuming most of readers only browse up to 10 recommended articles (Karypis 2001). Meanwhile, we also apply mean average precision (MAP) and mean average novelty (MAN) for evaluation from a system engineering point of view.

Overall Performance

A paired t-test shows that using P@10 and D@10 as performance measures, our approach performed significantly better than the best baseline methods as shown in Table 2, at p=0.005 and p=0.015, respectively. In addition, we have t-tests using MAP and MAN as performance measures, respectively, and the p values of these tests are all less than 0.05, which means that the results of experiments are statistically significant. We believe that such gains are caused by the additional information from comment contents.

Parameters of Topic Profile

There are two major parameters to be considered carefully to construct topic profile for news recommendations. That is, (1) the number of highest weighted words to represent the topic; (2) coefficient α to determine the contribution of original posting and comments in selecting relevant news.We conduct a series of preliminary experiments to find the optimal performance obtained when the number of words is between 50 to 70 words, and α is between 0.65 to 0.75. Note that, α can be manually set by readers in the GUI to satisfy different purposes.When α is set to 0, the recommended news only reflect the author's opinion. When α is 1, suggested news only declare the concerns of readers. In the following experiments, we empirically set topic file word number to 60 and α to 0.7.

Effect of Comments

To observe the impact of readers' concerns on original news posting in social media, which is reflected by comments on the news postings, we study the following scenarios:

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

Table 3: Performance of three runs

	Precision		Novelty	
Method	P@10	MAP	D@10	MAN
RUN 1 (News)	0.88	0.869	0.853	0.794
RUN 2 (Comments)	0.92	0.891	0.9	0.848
RUN 3 (Both)	0.94	0.932	0.9	0.848

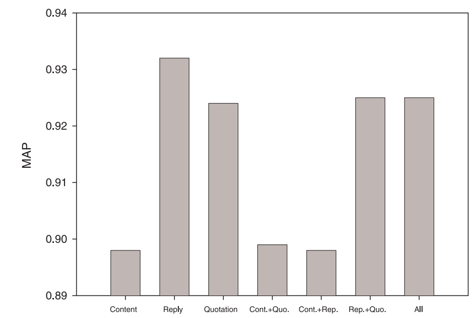

Figure 4: Effect of content, quotation and reply relation

- RUN 2 (Comments): the topic file is constructed only based on the comments.
- RUN3 (Both): the topic file is constructed based on both the content of news and its comments.

Here, we set $\gamma_1 = \gamma_2 = \gamma_3 = 1$. Our t-test shows that using P@10 and D@10 as performance measures, RUN3 performs best as shown in Table 3, at $p = 0.021$ and $p = 0.014$, respectively. These are much less than the critical confidence value (0.05). We believe that such gains are contributed by the additional information from user comments.

Furthermore, we investigate the effect of the three forms of relationship among comments, i.e. content similarity, reply, and quotation. We carry out a series of experiments for this purpose. As shown in Figure 4, we can observe that replies are slightly more effective than quotations and both of these outperform pure content similarity. In other words, the importance of comments can be well evaluated by the logic organization of these comments. Quoting and replying reveal readers' concerns on discussion topics. We also notice that the incorporation of content similarity decreases the system effectiveness. This may seem to contradict our intuition that the textual information should complement the logic-based models. By further investigating our results, we find that content similarity sometimes misleads the decision on the importance of the comments. Besides, the computation cost of calculating the content similarity matrix M_C is very high. Therefore, we only apply the structural information to determine the importance of each comment.

Recommendation Explanation

To evaluate the precision of explaining the relationship of recommended news to the original posting, the evaluation metric of successful rate S is defined as

$$S = \sum_{i=1}^{m} (1 - e_i)/m$$

where m is the number of recommended news articles, e_i is the error weight of recommended article i, which is employed to account for gravity of the error in the computed success rate. In our study, the error weight is set to one if the explanation result is mis-labelled.

Based on our study, we observed that the success rate at top-10 is around 89.3%. Note that this includes the errors caused by the retrieval function. To estimate optimal thresholds of generalization and specialization, we calculated the success rate at different thresholds and found that neither too small nor too large values are good for explanation. In our experiment, we set generalization threshold Θ_g to 3.2 and specialization threshold Θ_s to 1.8. Actually, threshold values need to be set through a machine learning process, which identifies proper values based on a given training sample.

Conclusion and Future Work

The Web has become a platform for information exchange and user interaction, instead of information dissemination at its earlier stage. Many of its applications are also being extended in this fashion. News recommendation is one such example. Traditional application of 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 facilitate online discussion by providing relevant news references to the active participants. In this work, we present a framework for adaptive news recommendation that incorporates information from the entire discussion thread. This framework models the logic structure among the comments by distinguishing the cases of replies and quotations. By combining such logic structural information with traditional statistical language models, it can recommend news articles that meet the dynamic nature of a discussion forum. Our tests indicate that this framework is able to return a wider scope of articles from the Web which reflect the trend in the ongoing discussion among its users. These articles are presented to the user via a novel user interface.

This study can be extended in a few interesting ways. For example, we can use this technique to process personal Web blogs and email archives. The technique itself can also be extended by incorporating such information as reader scores on comments, chronological information of comments, and reputation of users. Indeed, its power is yet to be further improved and investigated.

References

Ahn, J.; Brusilovsky, P.; Grady, J.; He, D. and Syn, S. 2007. Open user profiles for adaptive news systems: help or harm? In *Proc. of WWW*, 11–20.

Bogers, T., and Van Den Bosch, A. 2007. Comparing and evaluating information retrieval algorithms for news recommendation. In *Proc. of ACM conference on Recommender Systems*, 141–144.

Brin, S. and Page, L. 1998. The anatomy of a large-scale hypertextual Web search engine. *Comput. Netw. ISDN Syst.* 30(1-7):107–117.

Candan, K. S.; Dönderler, M. E.; Hedgpeth, T.; Kim, J. W.; Li, Q. and Sapino, M. L. 2009. SEA: Segment-enrich-annotate paradigm for adapting dialog-based content for improved accessibility. *ACM Transactions on Information Systems* 27(3):1–45.

Cantador, I., Bellogin, A. and Castells, P. 2008. Ontology-based personalized and context-aware recommendations of news items. In *Proc. of WI*, 562–565.

Chiang, J.-H. and Chen, Y.-C. 2004. An intelligent news recommender agent for filtering and categorizing large volumes of text corpus. *International Journal of Intelligent Systems* 19(3):201–216.

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

Das, A. S.; Datar, M.; Garg, A. and Rajaram, S. 2007. Google news personalization: scalable online collaborative filtering. In *Proc. of WWW*, 271–280.

Del Corso, G. M.; Gullí, A. and Romani, F. 2005. Ranking a stream of news. In *Proc. of WWW*, 97–106.

Karypis, G. 2001. Evaluation of item-based top-n recommendation algorithms. In *Proc. of CIKM*, 247–254.

Lai, H.-J.; Liang, T.-P. and Ku, Y. C. 2003. Customized internet news services based on customer profiles. In *Proc. of ICEC*, 225–229.

Lavrenko, V. and Croft, W. B. 2001. Relevance based language models. In *Proc. of SIGIR*, 120–127.

Lavrenko, V.; Schmill, M.; Lawrie, D.; Ogilvie, P.; Jensen, D. and Allan, J. 2000. Language models for financial news recommendation. In *Proc. of CIKM*, 389–396.

Lee, H. and Park, S. 2007. MONERS: A news recommender for the mobile web. *Expert Systems with Applications* 32(1):143–150.

Ponte, J. M. and Croft, W. B. 1998. A language modeling approach to information retrieval. In *Proc. of SIGIR*, 275–281.

Ricardo, B. and Berthier, R. 1999. Modern information retrieval. *Addison Wesley Longman Publishing Co. Inc.*.

Robertson, S. E. and Walker, S. 1994. Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval. In *Proc. of SIGIR*, 232–241.

Wang, W., and Benbasat, I. 2007. Recommendation agents for electronic commerce: Effects of explanation facilities on trusting beliefs. *Journal of Management Information Systems* 23(4):217–246.

Zobel, J. 1998. How reliable are the results of large-scale information retrieval experiments? In *Proc. of SIGIR*, 307–314.