A Semantic-Aware Profile Updating Model for Text Recommendation

Hossein Rahmatizadeh Zagheli[1], Hamed Zamani[2], and Azadeh Shakery[1,3]
[1] School of Electrical and Computer Engineering, College of Engineering, University of Tehran, Iran
[2] Center for Intelligent Information Retrieval, University of Massachusetts Amherst, MA 01003
[3] School of Computer Science, Institute for Research in Fundamental Sciences (IPM), Iran
{rahmatizadeh,shakery}@ut.ac.ir,zamani@cs.umass.edu

ABSTRACT

Content-based recommender systems (CBRSs) rely on user-item similarities that are calculated between user profiles and item representations. Appropriate representation of each user profile based on the user's past preferences can have a great impact on user's satisfaction in CBRSs. In this paper, we focus on text recommendation and propose a novel profile updating model based on previously recommended items as well as semantic similarity of terms calculated using distributed representation of words. We evaluate our model using two standard text recommendation datasets: TREC-9 Filtering Track and CLEF 2008-09 INFILE Track collections. Our experiments investigate the importance of both past recommended items and semantic similarities in recommendation performance. The proposed profile updating method significantly outperforms the baselines, which confirms the importance of incorporating semantic similarities in the profile updating task.

KEYWORDS

Content-based recommender systems, text recommendation, adaptive filtering, word embedding, semantic similarity

1 INTRODUCTION

Text recommendation is the task of delivering sets of documents to users in order to satisfy their information needs. There exist several real-world applications for text recommendation, including applications for recommending blog and social media posts [9], news articles [3, 25, 27], scientific papers [2, 31], and products (based on their reviews) [18]. In content-based text recommender systems, the recommendation decision is made based on the similarity of user profile to the candidate documents. It has been shown that accurately modeling user profiles can lead to significant improvements in recommendation performance [7, 24].

The Rocchio's relevance feedback algorithm [30] has been extensively used for profile updating within the vector space model framework [24]. The language modeling framework for information retrieval [26] has shown theoretical and empirical advantages over the traditional vector space model [7, 16]. Relevance models proposed by Lavrenko and Croft [15] have been successfully used for profile updating in the language modeling framework [32]. All of these profile updating models use recommendation history and users' feedback in order to construct a profile for each user.

There is a large overlap between the profile updating task in text recommendation and query modeling in search engines. Recently, exploiting semantic similarities captured by distributed representation of words has attracted a lot of attention for query modeling in the information retrieval literature [13, 28, 34–36]. In this paper, we first propose to update user profiles based on recommendation history using a log-logistic model. We further propose to enrich user profiles via semantic similarity of words. The intuition behind our approach is that each user profile should be constructed based on the words extracted from the previous successful recommendations to the user that are semantically similar to the initial user profile.

We evaluate our model with the adaptive filtering setup using two standard datasets: TREC-9 Filtering Track (containing scientific publications in the field of medicine) and CLEF 2008-2009 Information Filtering Evaluation (INFLIE) Track (containing news articles) collections. The results and analysis demonstrate that the proposed log-logistic model for profile updating outperforms state-of-the-art baselines, including the mixture model and the relevance model. Our experiments also show that although updating user profiles solely based on semantic similarity of terms does not outperform the profile updating approaches based on recommendation history, constructing a semantic-aware profile updating model based on recommendation history significantly outperforms all the baselines. Our results suggest to construct semantic-aware user profiles for text recommendation.

2 RELATED WORK

Word embedding models, such as word2vec [19], have been shown to be highly effective in many natural language processing tasks. Several attempts have been made for improving the recommendation performance by employing word embedding vectors. Musto et al. [21, 22] used word embeddings for content-based recommender systems. The authors used Wikipedia as an external resource for item recommendation. Huang [11] employed word embeddings for representing job postings in a job recommendation scenario. These studies do not focus on text recommendation. Ozsoy [23] and Krishnamurthy et al. [12] used the word embedding idea to model non-textual items for the item recommendation task. Recently, Bansal et al. [2] used recurrent neural networks for text representation in a supervised manner.

Furthermore, employing semantic similarities calculated based on word embedding vectors for the query expansion task has shown significant improvements in the retrieval performance [13, 34–36].

RecSys '17, Como, Italy

DOI: 10.1145/3109859.3109904

However, the effectiveness of employing such semantic similarities in profile updating for text recommendation is still relatively unstudied, which is the focus of this paper.

3 TEXT RECOMMENDATION FRAMEWORK

Text recommendation frameworks consist of three major components: text representation, filtering component, and profile learner [17]. In this section, we present how we implement the first two components in our recommendation process.

The text representation component in our recommender system is based on the language modeling framework in information retrieval [26]. Similar to prior work [16, 27, 33], each document D is represented with a unigram language model (called θ_D) smoothed using the Dirichlet prior smoothing method [38] as follows:

$$p(w|\theta_D) = \frac{|D|}{|D| + \mu} p_{ML}(w|D) + \frac{\mu}{|D| + \mu} p(w|C) \tag{1}$$

where $p_{ML}(w|D)$ and C represent the maximum likelihood estimation for document D and a reference language model, respectively. μ is the smoothing parameter.

To implement the filtering component, we use the KL-divergence retrieval model [14]. In more detail, the similarity score of each user profile P and each candidate document D is calculated as:

$$score(P, D) = -D(\theta_P || \theta_D) = -\sum_{w \in P} p(w|\theta_P) \log \frac{p(w|\theta_P)}{p(w|\theta_D)} \tag{2}$$

where $D(\cdot||\cdot)$ denotes the KL-divergence formula, and θ_P represents the profile language model. In Section 4, we present our methodology to model and update user profiles.

The recommendations are based on the similarity scores computed by the filtering component and a dissemination threshold τ. Updating the dissemination threshold plays a key role in improving the recommendation performance in text recommendation [39]. To update the threshold, we use the linear auto-adjust threshold optimization (LAUTO) algorithm [33], a simple yet effective algorithm that has recently shown relatively good performance [27, 33]. The algorithm starts with an initial threshold and is adjusted based on each user's recommendation history. Based on the LAUTO algorithm, passing a number of non-relevant documents to the user means that the dissemination threshold is probably lower than the optimal value and the threshold should be increased. In addition, rejecting a large number of continuous arrival documents by the filtering component is a signal to shows that the dissemination threshold is higher than the optimal value and should be decreased.

4 SEMANTIC-AWARE PROFILE UPDATING

In this section, we present our semantic-aware profile updating model. To do so, we first explain how we construct a profile model from the user's recommendation history. We further explain how we inject semantic similarity calculated based on distributed representation of words into our model.

In recommender systems, including this paper, the user's profile may be updated after each successful recommendation to the user. Consider a document D recommended to user u, for which an implicit or explicit feedback shows that this was a successful recommendation. We extract the weight of each term from the recommended document D using a log-logistic model. Although log-logistic has shown good performance in the pseudo-relevance feedback task in information retrieval [8, 20], to the best of our knowledge, this is the first attempt to update user profiles using this approach. In fact, the goal is to realize which term in the document is a good candidate for updating the user profile. The log-logistic score for each term w in the recommended document D is calculated as:

$$LL(w, D) = \log \left(\frac{\text{count}(w, D) * \log(1 + c\frac{avgl}{|D|}) + \lambda_w}{\lambda_w} \right) \tag{3}$$

where $\text{count}(w, D)$ and $avgl$ denote the frequency of term w in document D and the average document length, respectively. c is a free hyper-parameter that controls the weight of document length normalization component. λ_w shows how frequent the term w is in the reference collection. In other words, λ_w is equal to N_w/N where N_w and N respectively denote the total number of documents that contain w and the total number of documents in the collection. This part of the formula models how general the term w is. The logarithm function is used to satisfy the concavity constraint for term frequency.

One of the main shortcomings of the log-logistic formulation (i.e., Equation (3)) is that it is independent of the user's original profile. In other words, $LL(w, D)$ only depends on the recommended document, while the word might not be related to the initial user profile. To address this issue, we propose to involve the semantic similarity of words to the original user profile in the profile updating model. To this aim, the profile language model θ_{P^*} is calculated as follows:

$$p(w|\theta_{P^*}) \propto score(w, F^+) p_{sem}(w|P) \tag{4}$$

where F^+ and P denote the set of past successful recommended documents and the original user profile, respectively. $score(w, F^+)$ is calculated based on the independence assumption of the recommended documents as follows:

$$score(w, F^+) \propto \sum_{D \in F^+} LL(w, D) \tag{5}$$

where $p_{sem}(w|P)$ represents the semantic similarity between the term and the original user's profile. $p_{sem}(w|P)$ can be calculated in different ways. A simple and well-known calculation is based on the similarity of the embedding vectors of w and P. We propose a translation-based probabilistic model [4] to compute this probability:

$$p_{sem}(w|P) = \sum_{p_i \in P} p_{sem}(w|p_i) p_{sem}(p_i|P) \tag{6}$$

where p_i is a profile term. The intuition behind this model is that the semantic similarities to the words that are close to the whole user's profile should get higher weights. p_{sem} is calculated based on the softmax function as follows:

$$p_{sem}(w|p_i) = \frac{\exp(\vec{w}.\vec{p_i})}{\sum_{w' \in V} \exp(\vec{w'}.\vec{p_i})} \tag{7}$$

where V denotes the vocabulary set. The vectors in Equation (7) denote the distributed representations of the given terms. The probability $p_{sem}(p_i|P)$ is estimated similarly where $\vec{P}$ is the centroid vector of all profile terms.

Table 1: Performance of the proposed method and the baselines. The best result in each column is boldfaced. The superscripts 1/2/3/4/5/6/7/8 denote that the improvements over NoUpdate/MIXTURE/RM3/LL/Cent/CombSUM/CombMNZ/CombMAX are statistically significant.

Method	OHSUMED			INFILE		
	F1-measure	Precision	Recall	F1-measure	Precision	Recall
NoUpdate	0.2007	0.301	0.3281	0.3472	0.4696	0.4932
MIX	0.2251	0.3019	0.3611	0.3598	0.4779	0.5213
RM3	0.2263	0.3031	0.3605	0.3581	0.4758	**0.5282**
LL	0.2219	0.3003	0.3706	0.3631	0.4777	0.5160
Cent	0.2134	0.3015	0.3398	0.3525	0.4897	0.5001
CombSUM	0.2052	0.3018	0.329	0.3491	0.4915	0.4985
CombMNZ	0.205	0.3032	0.3363	0.3507	**0.4927**	0.4994
CombMAX	0.2114	0.2979	0.3516	0.3451	0.4746	0.515
LL-WE	**0.2344**^{145678}	**0.3161**^{12345678}	**0.3912**^{12345678}	**0.3771**^{12345678}	0.4871^{1}	0.5272^{15678}

The final language model is estimated using the linear interpolation of the estimated profile model and the original profile model as follows:

$$p(w|\theta_P) = \alpha p_{ML}(w|P) + (1-\alpha)p(w|\theta_{P^*}) \quad (8)$$

where the parameter α controls the weight of the original profile model $p_{ML}(w|P)$ computed based on maximum likelihood estimation.

5 EXPERIMENTS

In this section, we first explain our experimental design, including the datasets that we used, the pre-processing steps, and the experimental setup details. We further introduce our evaluation metrics and finally report and discuss the results.

5.1 Experimental Design

Datasets. In our experiments, we use two standard collections: The first one is the OHSUMED collection [10] used in TREC-9 Filtering Track [29]. This collection consists of 348,566 documents collected from the United States National Library of Medicine's bibliographic database between 1987 and 1991. The TREC-9 Filtering Track's collection contains 63 topics. The second one is the INFILE collection [5] that includes around 1.5 million news articles published in a three-year period (2004-2006) by Agence France Presse (AFP). We only considered the English documents in this collection. This collection was used in CLEF 2008-2009 INFILE (INformation FILtering Evaluation) Track [6]. The INFILE collection contains 50 topics covering two different categories: 30 topics focus on general news and events and the other 20 topics include scientific and technological issues. We considered the title and keywords of each topic as the initial profile of each user.

Experimental Setup. All the experiments were carried out using the Lemur toolkit[1]. All documents were stemmed using the Porter stemmer and stopped using the INQUERY stopword list. Similar to the CLEF 2009 Filtering Track, feedback is not allowed on discarded documents and limited number of feedback documents (200 feedback documents) is allowed in our experiments [5]. The embedding vectors were trained using the word2vec model[2] [19] on the target collections.

Parameters Setting. The Dirichlet prior smoothing parameter μ is set to the average document length of each collection. The free hyper-parameters α and maximum profile size were set using 5-fold cross validation over the topics of each collection and were selected from $\{0.1, \cdots, 0.9\}$ and $\{10, 30, 50\}$, respectively. The embedding dimension was set to 100 and the parameter c of the log-logistic model was set to 1.

5.2 Evaluation Metrics

Following the evaluation methodology in CLEF 2008-2009 Filtering Track [6], we use F1-measure as the main evaluation metric. To have complete sets of evaluations, we also report precision and recall. The reported results for each metric are based on 5-fold cross validation where that metric is optimized. Statistical significance of the differences between the corresponding measures were calculated using the two-tailed paired t-test at a 95% confidence level.

5.3 Results and Discussion

In this subsection, we first introduce our baselines. We further report and discuss the results achieved by the proposed method and the baselines.

Baselines. The baselines can be categorized as follows:

- **No profile updating**: this is our simple text recommendation model with no profile updating (NoUpdate).
- **Profile updating based on recommendation history**: these baselines use the past successful recommendations to adaptively update the user profiles. We use the mixture model (MIX) [37], the relevance model (RM3) [1, 15], and the log-logistic model (LL) [8] as different profile updating models in this category.

[1] http://www.lemurproject.org/

[2] https://code.google.com/p/word2vec/

Table 2: The updated profiles for the initial topic "fight against climate change". Note that the terms were stemmed.

total number of feedback documents	top 10 profile terms in descending order
2	kyoto, climat, greenhous, gase, toothless, protocol, reduct, emiss, meaning, ecolog
5	greenhous, kyoto, climat, gase, emiss, warm, protocol, acia, reduct, ratifi
10	climat, greenhous, emiss, kyoto, gase, warm, protocol, acia, dioxid, carbon
22	emiss, climat, greenhous, kyoto, gase, protocol, warm, carbon, dioxid, environment
34	emiss, climat, greenhous, kyoto, gase, warm, protocol, carbon, environment, dioxid

- **Profile updating based on semantic similarity**: these baselines only consider the semantic similarity of vocabulary terms to the initial user profiles. All of these models are based on word embedding vectors learned from the target collection. The first baseline in this category is based on the similarity of terms to the centroid vector of the profile terms (Cent). We also consider three other baselines in this category, including CombSUM, CombMNZ, and CombMAX, that were previously used by Kuzi et al. [13] for query expansion. We used these methods for updating user profile using previously successful recommendations.

Results. The experimental results are reported in Table 1. The highest value in each column is boldfaced. According to the table, the results achieved by all the profile updating methods are higher than those obtained by the NoUpdate model (except the results of CombMAX and LL in terms of Precision in OHSUMED collection). This shows the importance of profile updating in text recommendation. Among the profile updating methods based on recommendation history (i.e., MIX, RM3, and LL), LL outperforms the other methods in terms of F1-measure in INFILE dataset. This indicates the effectiveness of the log-logistic model for profile updating. The profile updating methods solely based on semantic similarities (i.e., Cent, CombSUM, CombMNZ, and CombMAX) perform worse than those based on recommendation history. This emphasizes on the fact that the past successful recommendations provide strong signals for representing the users' information needs. Among the profile updating methods based on semantic similarities, Cent achieves the highest F1-measures in both collections. Zamani and Croft [35] have theoretically proved that this model would result in the global optimum representation in terms of likelihood.

According to Table 1, the proposed semantic-aware profile updating model, i.e., LL-WE, outperforms all the baselines in terms of F1-measure in both collections. The method also achieves higher precision and recall values in the TREC-9 Filtering Track collection. The improvements are statistically significant in nearly all cases. This shows that incorporating semantic similarity in profile updating based on recommendation history leads to significant improvements.

To have an insight into the updated profiles, Table 2 reports the updated profile obtained by the LL-WE method for a random topic from the INFILE collection. According to this table, after recommending sufficient number of relevant documents, the profile becomes stable. The updated profiles contain relevant terms that can better represent the user's information need[3].

[3]Note that the narrative for this topic is "Relevant documents will describe the climatic policy implemented at the national or international level to fight against climatic change. They will speak about international agreements relative to air pollution control and more specifically to human induced greenhouse gases emissions reduction, responsible for climate warming".

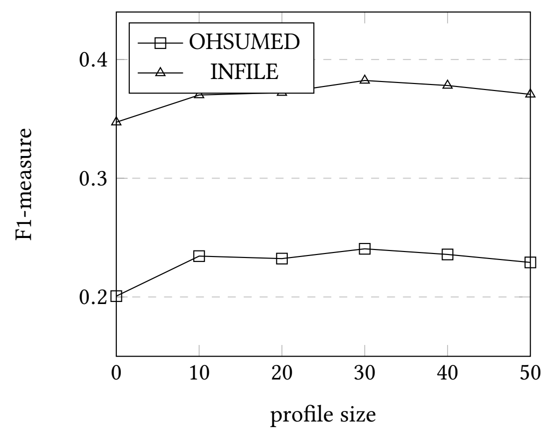

Figure 1: Sensitivity of the LL-WE method to the profile size in terms of F1-measure.

Sensitivity to the profile size. In this set of experiments, we study the sensitivity of the LL-WE method to the profile size. We sweeped the profile size from 0 to 50 terms. The results in terms of F1-measure are plotted in Figure 1. According to this figure, 30 is a reasonable profile size for the proposed profile updating model. Note that increasing the profile size does not lead to dramatic decrease in the recommendation performance.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we focused on the profile updating task in text recommendation. We first proposed a log-logistic model for profile updating and further enriched our model by incorporating semantic similarities using distributed representations of words. We implemented our model in the language modeling framework and evaluated our models using two standard datasets: TREC-9 Filtering Track and CLEF 2008-2009 INFILE Track collections. The experiments demonstrated that the proposed semantic-aware log-logistic profile updating model significantly outperforms all the other profile updating models in both collections.

An interesting future direction is to consider semantic similarities in using negative feedback. In addition, this paper focused on the word embedding vectors that were learned based on term proximity. We also intend to investigate how to learn embedding vectors that are specific to the text recommendation task.

Acknowledgements. This work was supported in part by a grant from the Institute for Research in Fundamental Sciences (No. CS1396-4-51) and in part by the Center for Intelligent Information Retrieval. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsors.

REFERENCES

[1] Nasreen Abdul-jaleel, James Allan, W. Bruce Croft, Fernando Diaz, Leah Larkey, Xiaoyan Li, Donald Metzler, Mark D. Smucker, Trevor Strohman, Howard Turtle, and Courtney Wade. 2004. UMass at TREC 2004: Novelty and HARD. In *Proceedings of the 2004 Text Retrieval Conference (TREC '04)*. National Institute of Standards and Technology, Special Publication.

[2] Trapit Bansal, David Belanger, and Andrew McCallum. 2016. Ask the GRU: Multi-task Learning for Deep Text Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*. ACM, New York, NY, USA, 107–114.

[3] Trapit Bansal, Mrinal Das, and Chiranjib Bhattacharyya. 2015. Content Driven User Profiling for Comment-Worthy Recommendations of News and Blog Articles. In *Proceedings of the 9th ACM Conference on Recommender Systems (RecSys '15)*. ACM, New York, NY, USA, 195–202.

[4] Adam Berger and John Lafferty. 1999. Information Retrieval As Statistical Translation. In *Proceedings of the 22Nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '99)*. ACM, New York, NY, USA, 222–229.

[5] Romaric Besançon, Stéphane Chaudiron, Djamel Mostefa, Ismaïl Timimi, and Khalid Choukri. 2008. The INFILE Project: a Crosslingual Filtering Systems Evaluation Campaign. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC '08)*. ELRA, Marrakech, Morocco.

[6] Romaric Besançon, Stéphane Chaudiron, Djamel Mostefa, Ismaïl Timimi, Khalid Choukri, and Meriama Laïb. 2009. Information Filtering Evaluation: Overview of CLEF 2009 INFILE Track. In *Proceedings of the 10th Cross-language Evaluation Forum Conference on Multilingual Information Access Evaluation: Text Retrieval Experiments (CLEF '09)*. Springer-Verlag, Berlin, Heidelberg, 342–353.

[7] Toine Bogers and Antal van den Bosch. 2007. Comparing and Evaluating Information Retrieval Algorithms for News Recommendation. In *Proceedings of the 2007 ACM Conference on Recommender Systems (RecSys '07)*. ACM, New York, NY, USA, 141–144.

[8] Stéphane Clinchant and Eric Gaussier. 2010. Information-based Models for Ad Hoc IR. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '10)*. ACM, New York, NY, USA, 234–241.

[9] Ido Guy, Naama Zwerdling, Inbal Ronen, David Carmel, and Erel Uziel. 2010. Social Media Recommendation Based on People and Tags. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '10)*. ACM, New York, NY, USA, 194–201.

[10] William Hersh, Chris Buckley, T. J. Leone, and David Hickam. 1994. OHSUMED: An Interactive Retrieval Evaluation and New Large Test Collection for Research. In *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '94)*. Springer-Verlag New York, Inc., New York, NY, USA, 192–201.

[11] Yanbo Huang. 2016. *Exploiting Embedding in Content-Based Recommender systems*. Master's thesis. TU Delf.

[12] Balaji Krishnamurthy, Nikaash Puri, and Raghavender Goel. 2016. Learning Vector-space Representations of Items for Recommendations Using Word Embedding Models. *Procedia Computer Science* 80 (2016), 2205 – 2210.

[13] Saar Kuzi, Anna Shtok, and Oren Kurland. 2016. Query Expansion Using Word Embeddings. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management (CIKM '16)*. ACM, New York, NY, USA, 1929–1932.

[14] John Lafferty and ChengXiang Zhai. 2001. Document Language Models, Query Models, and Risk Minimization for Information Retrieval. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '01)*. ACM, New York, NY, USA, 111–119.

[15] 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 (SIGIR '01)*. ACM, New York, NY, USA, 120–127.

[16] Victor Lavrenko, Matt Schmill, Dawn Lawrie, Paul Ogilvie, David Jensen, and James Allan. 2000. Language Models for Financial News Recommendation. In *Proceedings of the Ninth International Conference on Information and Knowledge Management (CIKM '00)*. ACM, New York, NY, USA, 389–396.

[17] Pasquale Lops, Marco de Gemmis, and Giovanni Semeraro. 2011. Content-based Recommender Systems: State of the Art and Trends. In *Recommender Systems Handbook*. Springer, 73–105.

[18] Julian McAuley and Jure Leskovec. 2013. Hidden Factors and Hidden Topics: Understanding Rating Dimensions with Review Text. In *Proceedings of the 7th ACM Conference on Recommender Systems (RecSys '13)*. ACM, New York, NY, USA, 165–172.

[19] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems (NIPS '13)*. 3111–3119.

[20] Ali Montazeralghaem, Hamed Zamani, and Azadeh Shakery. 2016. Axiomatic Analysis for Improving the Log-Logistic Feedback Model. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '16)*. ACM, New York, NY, USA, 765–768.

[21] Cataldo Musto, Giovanni Semeraro, Marco de Gemmis, and Pasquale Lops. 2015. Word Embedding Techniques for Content-based Recommender Systems: An Empirical Evaluation. In *Proceedings of the 9th ACM Conference on Recommender Systems (Posters) (RecSys '15)*. ACM, New York, NY, USA.

[22] Cataldo Musto, Giovanni Semeraro, Marco de Gemmis, and Pasquale Lops. 2016. Learning Word Embeddings from Wikipedia for Content-Based Recommender System. In *Proceedings of the 38th European Conference on Information Retrieval (ECIR '16)*. Springer International Publishing.

[23] Makbule Gulcin Ozsoy. 2016. From Word Embeddings to Item Recommendation. *CoRR* abs/1601.01356 (2016). http://arxiv.org/abs/1601.01356

[24] MichaelJ. Pazzani and Daniel Billsus. 2007. Content-Based Recommendation Systems. In *The Adaptive Web*. Lecture Notes in Computer Science, Vol. 4321. Springer Berlin Heidelberg, 325–341.

[25] Owen Phelan, Kevin McCarthy, and Barry Smyth. 2009. Using Twitter to Recommend Real-time Topical News. In *Proceedings of the Third ACM Conference on Recommender Systems (RecSys '09)*. ACM, New York, NY, USA, 385–388.

[26] Jay M. Ponte and W. 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 (SIGIR '98)*. ACM, New York, NY, USA, 275–281.

[27] Hossein Rahmatizadeh Zagheli, Mozhdeh Ariannezhad, and Azadeh Shakery. 2017. Negative Feedback in the Language Modeling Framework for Text Recommendation. In *Proceedings of the 39th European Conference on Information Retrieval (ECIR '17)*. Springer International Publishing.

[28] Navid Rekabsaz, Mihai Lupu, Allan Hanbury, and Hamed Zamani. 2017. Word Embedding Causes Topic Shifting; Exploit Global Context!. In *Proceedings of the 40th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '17)*. ACM, New York, NY, USA.

[29] Stephen Robertson and David A. Hull. 2000. The TREC-9 filtering track final report. In *Proceedings of Ninth Text REtrieval Conference (TREC-9)*. National Institute of Standards and Technology, Special Publication, 25–40.

[30] Joseph John Rocchio. 1971. Relevance feedback in information retrieval. *The Smart retrieval system-experiments in automatic document processing* (1971), 313–323.

[31] Chong Wang and David M. Blei. 2011. Collaborative Topic Modeling for Recommending Scientific Articles. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '11)*. ACM, New York, NY, USA, 448–456.

[32] Jia Wang, Qing Li, Yuanzhu Peter Chen, Jiafen Liu, Chen Zhang, and Zhangxi Lin. 2010. News Recommendation in Forum-Based Social Media. In *Proceedings of AAAI Conference on Artificial Intelligence*. AAAI Press.

[33] Hamed Zamani. 2015. *Recommender Systems for Multi-Publisher Environments*. Master's thesis. University of Tehran.

[34] Hamed Zamani and W. Bruce Croft. 2016. Embedding-based Query Language Models. In *Proceedings of the 2016 ACM International Conference on the Theory of Information Retrieval (ICTIR '16)*. ACM, New York, NY, USA, 147–156.

[35] Hamed Zamani and W. Bruce Croft. 2016. Estimating Embedding Vectors for Queries. In *Proceedings of the 2016 ACM International Conference on the Theory of Information Retrieval (ICTIR '16)*. ACM, New York, NY, USA, 123–132.

[36] Hamed Zamani and W. Bruce Croft. 2017. Relevance-based Word Embedding. In *Proceedings of the 40th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '17)*. ACM, New York, NY, USA.

[37] ChengXiang Zhai and John Lafferty. 2001. Model-based Feedback in the Language Modeling Approach to Information Retrieval. In *Proceedings of the Tenth International Conference on Information and Knowledge Management (CIKM '01)*. ACM, New York, NY, USA, 403–410.

[38] ChengXiang Zhai and John Lafferty. 2004. A Study of Smoothing Methods for Language Models Applied to Information Retrieval. *ACM Trans. Inf. Syst.* 22, 2 (April 2004), 179–214.

[39] Yi Zhang and Jamie Callan. 2001. Maximum Likelihood Estimation for Filtering Thresholds. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '01)*. ACM, New York, NY, USA, 294–302.