

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

News Recommender Based on Rich Feedback

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1521614> since 2018-01-15T09:33:05Z

Publisher:

Springer

Published version:

DOI:10.1007/978-3-319-20267-9_27

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)



UNIVERSITÀ DEGLI STUDI DI TORINO

This is an author version of the contribution published on:

Questa è la versione dell'autore dell'opera:

News recommendation based on rich user feedback, DOI 10.1007/978-3-319-20267-9_27

The definitive version is available at:

La versione definitiva è disponibile alla URL:

http://link.springer.com/chapter/10.1007/978-3-319-20267-9_27

News Recommender Based on Rich Feedback

Liliana Ardissono, Giovanna Petrone, and Francesco Vigliaturo

Dipartimento di Informatica, Università di Torino, Italy
{liliana.ardissono, giovanna.petrone}@unito.it

Abstract. This paper proposes to exploit author-defined tags and social interaction data (commenting and sharing news items) in news recommendation. Moreover it presents a hybrid news recommender which suggests news items on the basis of the reader's short and long-term reading history, taking reading trends and short-term interests into account. The experimental results we carried out provided encouraging results about the accuracy of the recommendations.

Keywords: Hybrid news recommender, tag-based news specification.

1 Introduction

The news domain challenges recommender systems with peculiar issues to be faced with respect to the suggestion of items in online catalogs [1]: e.g., the pool of stories continuously changes; freshness and diversity of recommendations are crucial to keep the reader's interest; moreover, as people are not expected to rate articles, user interests have to be learned unobtrusively; finally, people's interests dynamically change and they can be influenced by reading trends.

Existing news recommendation work tracks page visualization. However modern online newspapers also enable registered users to add comments to news stories and to share them in social networks. Moreover, they support synthetic descriptions of the content of news items based on author-defined tags. In order to investigate the impact of these features on recommendation we developed the Social News Recommender (SAND), which exploits these types of information for the selection of the news stories to suggest. SAND integrates a short-term and a long-term content-based recommender system with a collaborative filtering one to suggest stories on the basis of the user's past history, her/his recent reading history and peers' reading behavior. Moreover the system exploits a popularity-based recommender to enrich suggestions with successful items.

We tested SAND online by involving a small number of readers and using the feeds of newspaper "Il Fatto Quotidiano" (www.ilfattoquotidiano.it/). The test results show that the exploitation of rich information about news items and user actions supports a fairly accurate recommendation of news stories satisfying individual user interests.

Section 2 outlines the related work; Section 3 describes SAND; Section 4 summarizes our test results and Section 5 concludes the paper.

2 Related Work

Most news recommenders are based on collaborative filtering (e.g., [2, 3]), content-based filtering ([4, 5]), or integrated approaches ([1, 6]). SAND follows this approach but it differs in the following aspects: first, it analyzes new types of user behavior, which were not previously considered (article sharing and commenting), to enhance the estimation of readers' interests. Second, in order to promote recommendations based on rather different principles (with the aim of enhancing diversity, serendipity, etc.), it proposes a fair merging technique to propose the best suggestions identified by each of the integrated recommenders. Furthermore it exploits author-defined tags for characterizing the content of news items in an accurate and scalable way (taking inspiration from [8]) without analyzing the whole text of news stories. Different from our work, semantic news recommenders (e.g., [9]) depend on the definition and maintenance of large ontologies for the representation of concepts and semantic relations. We prefer to adopt a lightweight knowledge representation approach for scalability purposes. Navigation based approaches (e.g., [10]) are complementary to our work: they are suitable for learning interests of anonymous readers; however they only personalize the recommendation of news categories.

3 The SAND Hybrid News Recommender

SAND analyzes the news feeds published by online newspapers and provides personalized recommendations by learning individual user models and by combining different techniques with the aim of enhancing the quality of the suggestions it generates. It can integrate multiple news feeds but, being developed as a web service, it can also be used to enrich an individual online newspaper with personalized news suggestion. The system presents recent stories and recommendations enabling readers to share the stories in a social network and to add comments to them. Our current prototype downloads news stories from the feed of "Il Fatto Quotidiano" (<http://www.ilfattoquotidiano.it/feed/>). SAND includes two components: a news feed collector which downloads the stories published by the source newspapers and extracts information about their content (see Section 3.1) and a hybrid recommender which analyzes information about user actions on news stories and generates personalized suggestions (Section 3.3). We assume that readers authenticate themselves and can be monitored across sessions.

3.1 Analysis of the Content of News Items

The content of a news item i is represented as a bag of words $BW = \{ \langle k_1, r_1 \rangle, \dots, \langle k_m, r_m \rangle \}$ where $k_1 \dots, k_m$ is a set of keywords extracted from i and r_x is the relevance of k_x for the story, computed by applying TF-IDF [11]. The feed collector extracts the list of keywords of a news item from its title and author-defined tags by stop words removal and stemming. Different from previous work, which also uses the summary or the complete text of news items for keyword

extraction, SAND uses on title and tags because they focus on the topic of the articles, enabling a fast analysis of news items. We compared the two keyword extraction methods on a dataset from “Il Fatto Quotidiano”: for each document we compared the lists including the 5 most similar documents, generated by using title+summary and title+tags. The average intersection between the pairs of lists is about 31%. However, looking at a subset of the dataset in detail we noticed that the documents selected using title+tags are the most similar ones (we obtained comparable results considering the 10 most similar documents).

3.2 Representation of User Interests

The system manages two user models for each registered user U to describe short-term and long-term user interests. The models have the same format but refer to observations collected during the last 10 days / 60 days, respectively.

Each user model is a vector of observations $UM = \{ \langle i_1, e_1, t_1 \rangle, \dots, \langle i_n, e_n, t_n \rangle \}$ where i_1, \dots, i_n are the news items which U inspected during the relevant time interval, e_x is the evidence of U 's interest in i_x (inferred by tracking U 's behavior), and t_x is the timestamp of U 's last observed action concerning i_x . We model evidence as a numerical value in $[0, 1.5]$: (i) Clicking on the title of a news item carries weak information about the reader's interests because it does not prove that (s)he read the story. For this reason, we attribute this type of activity an evidence of interest in the item equal to 0.5. (ii) Sharing a news item in a social network or adding a comment to it do not mean that the actor approves its content but we can interpret these actions as strong evidence of interest towards the story. Thus we attribute them an evidence equal to 1.5.

The system stores in the user model an entry for each news item visualized by the user. The interest value associated to the item is the maximum value which can be attributed on the basis of the user's actions; i.e., 0.5 if (s)he only clicked on it, 1.5 if (s)he also shared/commented it. The timestamps of the observations allow forgetting older ones.

3.3 News Recommendation

According to previous works (e.g., [1, 12, 13]) recommenders based on collaborative filtering have more chances than content-based ones to suggest variegated items due to the fact that they exploit the ratings provided by other users; moreover they can reflect reading trends. However they tend to suggest highly popular items which the user might not be interested in, and they are affected by the novel item issue. In contrast, content-based recommenders can suggest new entries but they are not sensitive to quality and they tend to recommend items which are rather similar to each other. In order to combine the benefits of both approaches we designed SAND as a hybrid system: the system estimates a user U 's interest in the items to be analyzed by merging the suggestions of 4 different recommenders: a collaborative filtering one, a short-time content-based one, a long-term content-based one and a popularity-based one.

Let's consider a fresh item i and vector $interest_i$ representing the evidence values about the interest in i stored in the long-term user models of all the registered readers:

- The collaborative filtering recommender follows the item-to-item approach [12]. It predicts U 's preference for i by exploiting U 's observed interest in items similar to i from the viewpoint of the other readers' interests:

$$preference_{CF} = \frac{\sum_{j \in N(i)} \sigma(i, j) * e_j}{\sum_{j \in N(i)} \sigma(i, j)} \quad (1)$$

$\sigma(i, j)$ is the similarity between the interest vectors of i and j , evaluated as the Euclidean distance between the two vectors;¹ e_j is the evidence of interest in j stored in U 's long-term user model, and $N(i)$ is the set of neighbors of i , i.e., those news items whose similarity with i is over a threshold.

- The short-term content-based recommender predicts U 's preference for i by exploiting her/his recent interests in news stories whose content is similar to that of i . In order to emphasize the impact of individual keywords, which may refer to the occurrence of specific events (e.g., sport ones), the short-term recommender estimates the preference for a fresh news item by directly exploiting the observed interest in news stories having similar keywords. This is done by analyzing the items referenced in U 's short-term model and by selecting the most similar ones:

$$preference_{ST} = \sum_{j \in N(i)} \sigma(i, j) * e_j \quad (2)$$

Given the bag of words representations of i (BW_i) and j (BW_j), $\sigma(i, j)$ is computed as the cosine similarity between BW_i and BW_j . Moreover $N(i)$ is determined by selecting the items whose similarity is over a threshold, and e_j is the evidence of interest in j stored in U 's short-term user model.

- The long-term content-based recommender predicts U 's preference for i by taking into account U 's preferred topics emerging from the keywords which occur in the stories (s)he read in the recent past:
 1. First it estimates U 's interest in the keywords associated to the news items which (s)he inspected during the last 60 days by analyzing her/his observed interest in those items. For each keyword k the recommender computes the cumulative evidence of interest of k as follows:
$$cumulativeEvidence(k) = \sum_{j \in Stories} e(j)$$
 $Stories = \{j_1, \dots, j_h\}$ are the news items in U 's long-term model having k as a relevant keyword and $e(j)$ is the evidence of interest for item j .
 2. Then it exploits the 10 keywords having the highest cumulative evidence of interest ($MostInt$) for estimating U 's preference for i :

$$preference_{LT} = \sum_{k \in Keywords(i)} score(k) \quad (3)$$

¹ This is the implementation offered by the Apache Mahout framework exploited for implementing this recommender (<http://mahout.apache.org/>).

$Keywords(i)$ are the keywords in the bag of words representation of i .
 $score(k) = cumulativeEvidence(k)$ if $k \in MostInt$ (i.e., the keyword is a most interesting one), 0 otherwise.

- The popularity-based recommender suggests the news items of the current day that received the highest cumulative interest from registered readers.

The contributions of the recommenders are integrated into a list as parallel suggestion flows to include the best items of each flow. Given the lists of items separately ranked by the recommenders, and excluding those which the user has read, the final recommendation list includes 9 entries: typically, 3 items selected by the collaborative filtering one, 3 by the short-term content-based one, 2 by the long-term content-based one and 1 by the popularity-based one. However the first three recommenders start to predict a user’s preferences after having collected a minimum amount of evidence about her/his interests (10 user actions). Moreover, the user might be idle for a while and thus the short-term recommender might be unable to estimate her/his interests in that interval. In these cases the contributions of the long-term recommender and of the popularity-based one are extended, enabling them to suggest a larger number of items.

4 Evaluation Results

We evaluated the accuracy of SAND by means of an online test which involved 25 participants (12 males and 13 females; age: min=22, max=55, avg=38.7, median=37) having heterogeneous backgrounds (computer science, humanities, physics, electronics, architecture, social sciences, education science). Participants were asked to overview the news downloaded from the feed of “Il Fatto Quotidiano” for 30 days and they could click on news titles in order to read them; moreover they could make comments or share stories in Facebook. On day 31 the system recommended 9 news stories per participant and asked her/him to specify which ones were very interesting (if any) by clicking on a “Like” button. Moreover the system asked to answer 3 questions (selecting a value in [1, 5] where 1 is the worst value and 5 is the best one) to evaluate whether the recommended news items (i) matched the participant’s interests, (ii) dealt with heterogeneous topics, (iii) stimulated her/his curiosity even though they did not correspond to her/his general interests. Furthermore the system offered a free-text field to let participants provide further comments. The test provided the following results:

- 49.5% of suggested stories were evaluated as very interesting by clicking on the “Like button”, with an average of 4.5 clicks per user (std. dev. = 1.14).
- The questions aimed at assessing the recommendation accuracy received fairly positive answers: (1) long-term interest matching: 3.8; (ii) topic diversity; 3.7; (iii) stimulate curiosity: 3.7. All values are significant (one-sided T-test with $\alpha=0.05$).
- The free-text comments revealed that most participants felt that the system suggested relevant articles. The worst evaluations were provided by people who do not like “Il Fatto Quotidiano” because it focuses on topics out of their interests (the newspaper focuses on politics, economics and big events).

These results are encouraging but suggest that other online newspapers should be integrated to improve the diversity of topics and the system's capability of recommending rich pools of news stories. Moreover, an online test with a larger number of readers is needed to collect more extensive feedback and evaluations.

5 Conclusions

We described a hybrid news recommender that exploits information about readers' interaction with news stories and author-defined tags to provide personalized suggestions. The integration of different recommendation techniques in the generation of suggestions has provided encouraging experimental results. In our future work we will extend the analysis of user behavior to further enrich the evidence on user interest which can be collected thanks to the convergence of online newspapers with social networks. Moreover we will carry out further experiments in order to evaluate the recommendation accuracy which can be achieved by employing alternative integration strategies for the recommender systems, and to extend the evaluation to a larger set of users.

References

1. Liu, J., Donlan, P., Ronby Peders, E.: Personalized news recommendation based on click behavior. Proc. of IUI 2010, Hong Kong, China (2010) 31–40
2. Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Riedl, J.: Grouplens: an open architecture for collaborative filtering of netnews. Proc. of CSCW 1994, Chapel Hill, North Carolina, ACM (1994) 175–186
3. Das, A., Datar, M., Garg, A., Rajaram, S.: Google news personalization: scalable online collaborative filtering. In: Proc. of WWW'2007, Banff (2007) 271–280
4. Billsus, D., Pazzani, M.: A hybrid user model for news story classification. In: Proc. 7th Int. Conf. on User Modeling, Banff, Canada (1999) 99–108
5. Kompan, M., Bieliková, M.: Content-based news recommendation. In: E-commerce and Web Technologies, EC-Web 2010, LNBIP N. 61, Bilbao (2010) 61–72
6. Doychev, D., Lawlor, A., Rafter, R., Smyth, B.: An analysis of recommender algorithms for online news. Proc. of CLEF 2014, Sheffield, UK (2014)
7. Agrawal, R., Karimzadehgan, M., Zhai, C.: An online recommender system for social networks. Proc. of ACM SIGIR workshop on Search in Social Media, (2009)
8. Gemmel, J., Schimoler, T., Mobasher, B., Burke, R.: Resource recommendation in social annotation systems: a linear-weighted hybrid approach. Journal of computer and system sciences **78** (2012) 1160–1174
9. Cantador, I., Castells, P., Bellogin, A.: An enhanced semantic layer for hybrid recommender systems: Application to news recommendation. Int. Journal on Semantic Web and Information Systems **7**(1) (2011) 44–77
10. Garcin, F., Dimitrakakis, C., Faltings, B.: Personalized news recommendation with context trees. Proc. of (RecSys 2013), Honk Kong, China (2013) 105–112
11. Salton, G.: Automatic text processing. Addison-Wesley (1989)
12. Desrosiers, C., Karypis, G.: A comprehensive survey of neighborhood-based recommendation methods. Recommender systems handbook. Springer (2011) 107–144
13. Lops, P., de Gemmis, M., Semeraro, G.: Content-based recommender systems: state of the art and trends. Recommender systems handbook. Springer (2011)