Do Recommendations Matter? News Recommendation in Real Life

Alan Said
CWI
Amsterdam, The Netherlands
alansaid@acm.org

Alejandro Bellogín
CWI
Amsterdam, The Netherlands
abellogin@acm.org

Jimmy Lin
University of Maryland
College Park, MD, USA
jimmylin@umd.edu

Arjen de Vries
CWI
Amsterdam, The Netherlands
arjen@acm.org

CSCW'14 Companion, February 15–19, 2014, Baltimore, MD, USA.
ACM 978-1-4503-2541-7/14/02.
http://dx.doi.org/10.1145/2556420.2556510

Abstract
We present a study of how recommendations are received *in real life* by users across different news domains (traditional online newspapers, hobbyist websites, forums, etc.). Our analysis shows that readers of websites centered around specific topics are generally less likely to interact with recommendations than readers of traditional news websites.

Author Keywords
Recommender Systems; Evaluation; News Recommendation

ACM Classification Keywords
H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval - Information filtering

General Terms
Design, Experimentation, Human Factors

Introduction
Recommender systems have become a key feature in today's interactive and adaptive Web. Whether browsing research articles (Mendeley), music (Spotify, Last.fm), movies (Netflix, IMDb) or people (LinkedIn, Twitter), users have become accustomed to a personalized data consumption experience. In this work we present an

Type	URL
Business	cfoworld.de
IT	cio.de
IT	cnet.de
IT	computerwoche.de
IT	gulli.com
News	ksta.de
Cars	motor-talk.de
IT	silicon.de
Sports	sport1.de
News	tagesspiegel.de
IT	tecchannel.de
Living	wohnen-und-garten.de
IT	zdnet.de

Table 1: The 13 news providers' URLs and the types of news they deliver. Note that all providers deliver articles in German.

analysis of how recommendations fare in news domains across different topics. We specifically focus on traditional online newspaper articles and topic-focused news articles from hobbyist news sources. Traditional recommender systems (for music, movies, products) are usually based on collaborative filtering where recommendations are generated by looking at which items users have interacted with in the past [1]. In practice, the news recommendation scenario differs significantly from the traditional model as the recommendations must often be generated without access to the user's interaction history or other personal preferences. In this paper, we present an analysis of two basic recommendation approaches in four different news domains. Both approaches (recency-based and document similarity-based) are agnostic to the news readers' (users') previous interactions; instead, they focus either on information about the currently viewed item or on properties of the catalogue of available items. We find differences in how the two recommendation models are received in traditional news (newspapers and sports) versus more topic-focused (hobbyist) news domains.

The analysis in this paper is based on a 10 day data snapshot from the Plista[1] news recommendation service collected while participating in the 2013 News Recommender System Challenge (NRS'13) [3].

News Recommendation

As stated above, news recommendation differs from traditional recommendation due to data known (or not known) about the users, e.g., lack of profiles and interaction history. Additionally, news articles have different features than items in traditional recommendation settings, e.g., a short time span during which they are relevant and low consumption cost[2] which can result in a larger than normal diversity of the consumed items as a newspaper reader is more likely to "consume" less interesting articles without considering the recommendation bad.

Domains

The domains in which this study was performed varied from traditional news providers such as Tagesspiegel (newspaper) and Sport1 (sports) to websites focused on specific topics, e.g., Motor Talk (automotive) and Computerwoche (IT & Business). These were selected from the domains available in the NRS Challenge (see Table 1 for the complete list). In a previous analysis of a similar dataset, interaction differences across domains showed how the recommendation were received (i.e. clicked or not) by the readers [2], but no analysis of how different recommendation algorithms performed across different domains was presented.

Data and Analysis

The data on which this recommendation study is based consisted of all *impressions* and *clicks* performed by users of four news websites during a period of 10 days (Sept. 14 – Sept. 23, 2013). An impression in this context is the act of loading a web page (interpreted as reading the corresponding article), whereas a click is the act of clicking on a recommended article. News article recommendations are created in situ and are shown when the user scrolls down towards the bottom of the currently read article. See Fig. 1 for an example of an article recommendation on one of the online newspapers which uses the Plista recommendation service.

[1] http://www.plista.com

[2] A product purchase has a *monetary value* (a price), a book or movie has a *longer consumption cost* (the time it takes to consume the item). News articles on the other hand have a consumption cost of at most a few minutes and seldom have monetary costs attached.

Figure 1: An article recommendation (bottom right).

The daily number of recommendation requests to each algorithm used in this work ranged from a few hundred (Computerwoche) to tens of thousands (Sport 1, Tagesspiegel); the total number of impressions for the providers and algorithms was 1.2 million. The daily number of clicked recommendations ranged from a few dozen (Computerwoche) to thousands (Sport 1); the total number of clicks over the ten day period was 16.8 thousand. The global click-through rate (CTR) was 1.4%. However, as will be discussed later, there was a large variance in CTR across domains and algorithms.

Recommendation Approaches & Infrastructure
The recommendations in this study were either based on document (article) similarities or on the recency of the articles available on each website. In the first case, an Apache Lucene[3] index was created based on the articles' titles and leading paragraphs. The title and leading paragraph of the currently read article was used as the query. In the case where recency was considered, a random article from a pool of the most recently published articles on the news website was recommended, also taking into account the article's popularity (the number of views). The currently read article is sent through Plista's Open Recommendation Platform[4] to our recommendation server which in turn generates the recommendation and sends it in a reply. The platform serves as the communication hub, e.g., sending new articles as they are published, requesting recommendations, etc.

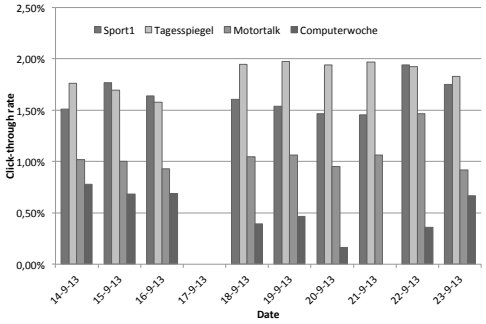

Figure 2: CTR for the document similarity-based recommender. Note that due to a server malfunction there were no recommendations on September 17th.

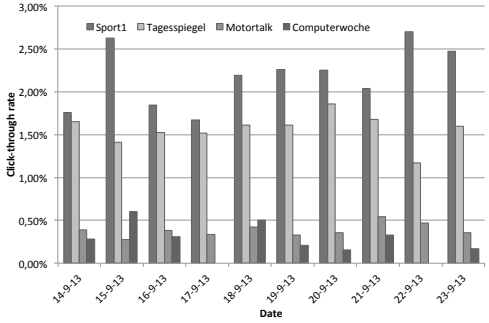

Figure 3: CTR for the article recency-based recommender.

[3] http://lucene.apache.org/

[4] http://orp.plista.com/

Results & Discussion

Analysis shows that users are generally less receptive (i.e., less willing to click on) recommendations when browsing articles on news sources focused on a specific topic. Both of the algorithms used in the study show similar low click-through rates on topic-focused websites. The document similarity recommender performed slightly better on topic-focused websites and slightly worse on general news and sport websites. These observations are shown in Figs. 2 and 3.

Considering the nature of *news*, it is intuitive that the recency-based recommender outperformed its document similarity-based counterpart. Inversely, one could imagine that news articles on a similar topic would be more well-liked by readers than just the most recent articles, regardless of their relation to the currently read article.

As for the slightly better effectiveness of the document similarity-based recommender on Motor Talk, we provide the following interpretation: readers of news sources related to specific topics browse these types of websites in a different context than readers of traditional news, with a different *purpose of consumption*. As this purpose is likely related to the specific interest, the user has a better understanding of what she is looking for (and thus is less receptive to recommendations). In contrast, the purpose of reading a (general) newspaper is to keep up to date with current events, and thus the reader is more willing to browse recommendations.

Conclusions & Future Work

In our analysis we have investigated the click-through rate of two news article recommender algorithms in four different news source (an online newspaper, a sports news website, an automotive news website, and a gardening website). Our analysis suggests that users reading news articles in online news outlets are less likely to interact with recommendations if the news outlet is specifically focused on a certain topic. We are currently performing a larger analysis over a longer period of time, a larger number of websites, and more recommendation algorithms in order to identify which recommendation models are suitable for different consumption purposes.

Acknowledgments

This work was in part carried out during the tenure of an ERCIM "Alain Bensoussan" Fellowship Programme. The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n°246016.

References

[1] Herlocker, J. L., Konstan, J. A., Terveen, L. G., and Riedl, J. T. Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst. 22*, 1 (Jan. 2004), 5–53.

[2] Said, A., Lin, J., Bellogín, A., and de Vries, A. A month in the life of a production news recommender system. In *Proceedings of the CIKM Workshop on Living Labs for Information Retrieval Evaluation*, LivingLab, ACM (2013).

[3] Tavakolifard, M., Gulla, J. A., Almeroth, K. C., Hopfgartner, F., Kille, B., Plumbaum, T., Brodt, T., and Heintz, T. Workshop and challenge on news recommender systems. In *Proceedings of the seventh ACM conference on Recommender systems*, RecSys '13, ACM (New York, NY, USA, 2013).