The Information Value of Context for a Mobile News Service

Toon De Pessemier, Kris Vanhecke and Luc Martens
iMinds - WiCa, Ghent University, Gaston Crommenlaan 8 box 201, B-9050 Ghent, Belgium

Keywords: Recommender System, News, User Study, Context, User Interaction.

Abstract: Traditional recommender systems provide personal suggestions based on the user's preferences, without taking into account any additional contextual information such as time or device type. However, in many applications, this contextual information may be relevant for the human decision process, and as a result, be important to incorporate into the recommendation process, which gave rise to context-aware recommender systems. However, the information value of contextual data for the recommendation process is highly dependent on the application domain and the users' consumption behavior in different contextual situations. This research aims to assess the information value of context for a recommender system of a mobile news service by analyzing user interactions and feedback. A large-scale user study shows that context-aware recommendations outperform traditional recommendations, but also indicates that the accuracy improvement might be limited in a real-life situation. Service usage takes place in a limited number of different contexts due to user habits and repetitive behavior, leaving little room for optimization based on the context. Data fragmentation over different contextual situations strengthens the sparsity problem, thereby limiting the user-perceived accuracy gain obtained by incorporating context in the recommender. These findings are important for news providers when considering to offer context-aware recommendations.

1 INTRODUCTION

Recommender systems are software tools and techniques providing suggestions for items to be of interest to a user such as videos, songs, or news articles. The consumption of these audiovisual media and the accessing of information always happen in a certain context (Ricci, 2012), i.e. conditions or circumstances that significantly affect the decision behavior. This gave rise to the development of context-aware recommender systems (CARS), which take this contextual information into account when providing recommendations.

For various application domains, the user context has gained an increased interest from researchers (Adomavicius and Tuzhilin, 2011). For context-aware music recommendations for example, the user's emotions can be used as input by using support vector machines as emotional state transition classifier (Han et al., 2010). In the application domain of tourism for example, various applications use the current location or activity of the user to personalize and adapt their content offer to the current user needs (Ricci, 2010; De Pessemier et al., 2014). Personal recommendations for points of interest can be provided based on the user's proximity of the venue (Kenteris et al., 2010).

In the domain of audiovisual media, more specifically news content, the influence of context on the consumption behavior and personal preferences is less obvious. However, research (Yu et al., 2006) has shown that the situation of the user (location, activity, time), as well as the device and network capabilities are important contextual parameters for context-aware media recommendations on smartphones.

The growth of the digital news industry and especially the development of mobile products is booming. Mobile has become, especially amongst younger media consumers, the first gateway to most online news brands. In a recent survey (Reuters Institute for the Study of Journalism, 2014), conducted in 10 countries with high Internet penetration, one-fifth of the users now claim that their mobile phone is the primary access point for news.

Despite this shift of news consumption to the mobile platform, the study of Weiss (Weiss, 2013) highlights that a gap exists between what news consumers, particularly young adults, are doing and using on their smartphones and what news organizations are able to provide. In most cases, news organizations disregard contextual data or they are only using geo-location features in their mobile apps for traffic and weather;

De Pessemier T., Vanhecke K. and Martens L..
The Information Value of Context for a Mobile News Service.
DOI: 10.5220/0005408603630369
In *Proceedings of the 11th International Conference on Web Information Systems and Technologies* (WEBIST-2015), pages 363-369
ISBN: 978-989-758-106-9
Copyright © 2015 SCITEPRESS (Science and Technology Publications, Lda.)

they do not anticipate the high use of location-based services by smartphone consumers.

2 RELATED WORK

In the domain of digital news services, various initiatives to personalize the offered news content have been proposed.

SCENE (Li et al., 2011) is such a news service. It stands for a SCalable two-stage pErsonalized News rEcommendation system. The system considers characteristics such as news content, access patterns, named entities, popularity, and recency of news items when performing recommendation. The proposed news selection mechanism demonstrates the importance of a good balance between user interests, the novelty, and diversity of the recommendations.

The News@hand system (Cantador et al., 2008) is a news recommender which applies semantic-based technologies to describe and relate news contents and user preferences in order to produce enhanced recommendations. This news system ensures multimedia source applicability and multi-domain portability. The resultant recommendations can be adapted to the current context of interest, thereby emphasizing the importance of contextualization in the domain of news. However, context is not the main focus of this study and the influence of context on the consumption behavior is not investigated.

The News Recommender Systems Challenge (Said et al., 2013) focused on providing live recommendations for readers of German news media articles. This challenge highlighted why news recommendations have not been as analyzed as some of the other domains such as movies, books, or music. Reasons for this include the lack of data sets as well as the lack of open systems to deploy algorithms in. In the challenge, the deployed recommenders for generating news recommendations are: Recent Recommender (based only on the recency of the articles), Lucene Recommender (a text retrieval system built on top of Apache Lucene), Category-based Recommender (using the article's category), User Filter (filters out the articles previously observed by the current user), and Combined Recommender (a stack or cascade of two or more of the above recommenders).

Although the various initiatives emphasize the importance of a personalized news offer, most of them focus on the recommendation algorithms and ignore the contextual information that is coupled with the information request, the user, and the device. In this study, the focus is not on improving state of the art recommendation algorithms, but rather on investigating the influence of context on the consumption of news content by means of a large-scale user study.

In many cases, the research on CARS remains conceptual (Adomavicius and Tuzhilin, 2011), where a certain method has been developed, but testing is limited to an offline evaluation or a short-term user test with only a handful of people, often students or colleagues who are not representative for the population. In contrast, this research investigates the role of context for news recommendations, based on a large-scale empirical study. Users could utilize a real news service[1] that offers content of four major Flemish news companies on their own mobile devices, in their everyday environment, where and when they wanted, i.e., in a living lab environment.

Living lab experiments are an extension towards more natural and realistic research test environments (Følstad, 2008). Living labs allow to evaluate research hypotheses by users representative for the target population who satisfy their information need in a real context. Since users are following their own agenda, laboratory biases on their behavior can be neglected (Hopfgartner et al., 2014). Although less transparent and predefined, living lab experiments aim to provide more natural settings for studying users' behavior and their experience.

Especially for context-aware applications, in which the user's environment has an influence on the way the application works and/or on the offered content, a realistic setting is essential for a reliable evaluation. Therefore, this paper investigates the influence of context and the benefit of context-aware recommendations for a real news service by means of a large-scale user panel, in a realistic environment, over a longer period of time. Since a user study can provide reliable explanations as to which recommendation method is the best, and why one method is better than the other (Shani and Gunawardana, 2013), three alternative recommendation methods for the news service are compared through such a user study.

3 EXPERIMENTAL SETUP

The news service that was used in this experiment[1], aggregates content of different premium content providers: newspapers, magazines, but also content of television as short video clips. Figure 1 shows a screenshot of the user interface offering the content with a reference of the provider on top of each content item. The aggregated content provides users a

[1] http://www.iminds.be/en/projects/2014/04/17/stream-store

more complete and varied overview of the news than traditional services do. To anticipate the abundance of news content and the associated choices that people have to make, the news service offers personalized recommendations.

During the experiment, the device type (smartphone or tablet) and the time of the day (morning, noon, daytime, or evening) are studied as contextual influences of the news consumption. The news service is accessible through a mobile application, which is available on Android and iOS for tablets as well as smartphones. As a result, the type of device that is used for consuming the news content is an interesting contextual factor.

Compared to the well-established application domains of recommender systems, such as movies or books, news items have a shorter lifespan and are frequently updated. Consequently, consulting the news on a daily basis, or even multiple times a day, can be interesting, which makes the time aspect another important contextual factor. The time is closely related to the location of the user, as was also witnessed during the analysis of users' interactions and consumption behavior. A frequently recurring pattern was as follows: in the morning, users are at home; during daytime, they are at work; and during the evening, they are again at home. Therefore, and to prevent over specification (Section 4.3), the location is not adopted as a separate contextual factor.

For the evaluation of this service and its recommendations, 120 test users were recruited by an experienced panel manager from iMinds-iLab.o[2] (i.e. a research division with a strong expertise in living lab research and panel management). These test users, all interested in news and owning a smartphone and/or tablet, belong to the target group of an online news service. The test users could install the mobile application of the news service on their smartphone and/or tablet and freely use the service during the evaluation period of around 5 weeks. These test users were divided into three groups, each receiving a different type of recommendations, as explained in Section 4.

The test users' interactions with the service were logged to analyze their consumption behavior and to get insight in the actual use of the news service and their overall experience: 10 test users did not install the app, or did not use the news service during the evaluation period. They are excluded from the analysis, so that the number of actual participants was reduced to 110.

[2]http://www.openlivinglabs.eu/livinglab/iminds-ilabo

4 NEWS RECOMMENDATIONS

The experiment takes three different approaches to recommend interesting news. Each user received only one recommendation type during the whole evaluation period. To avoid any bias, test users were not informed about the existence of multiple types of recommendations.

4.1 Recommendations based on Explicit Static Preferences

Before the actual experiment, test users were asked about their preferences for different categories of news content (National, International, Culture, Economy, Lifestyle, Politics, Sports, and Interesting facts) through an online questionnaire. Users could specify their interests on a 5-point rating scale for each category and refine this score for different times of the day (morning, noon, daytime, and evening). The answers on this questionnaire constitute the user profile that is used for generating news recommendations. During the experiment, these preferences are considered static; user profiles are not updated based on explicit or implicit feedback on the content, and the recommender is not learning from the user's behavior.

4.2 Content-based Recommendations

The content-based recommendations are not based on a prior questionnaire but use the implicit and explicit feedback users provide during the evaluation period. A request to read one of the recommended news items is considered as positive implicit feedback. Evaluating the news recommendation by means of the 'Thumbs Up' and 'Thumbs Down' icons in the user interface provides explicit feedback.

This feedback is gradually collected during the usage of the service. As a result, the user profiles of the content-based recommender are dynamic and constantly change as users interact with the news service. As the user is utilizing the news service and provides feedback, the recommender is learning the user's preferences.

A content-based recommendation algorithm was chosen because of the availability of informative metadata about the content items, the sparsity of the data set, and the cold start problem associated with the start-up phase. As content-based solution, the 'InterestLMS algorithm' of the Duine framework (Telematica Instituut / Novay, 2009) was adopted. The InterestLMS algorithm builds a user profile by inferring personal preferences from the metadata describing the news items that are requested (implicit feedback) or

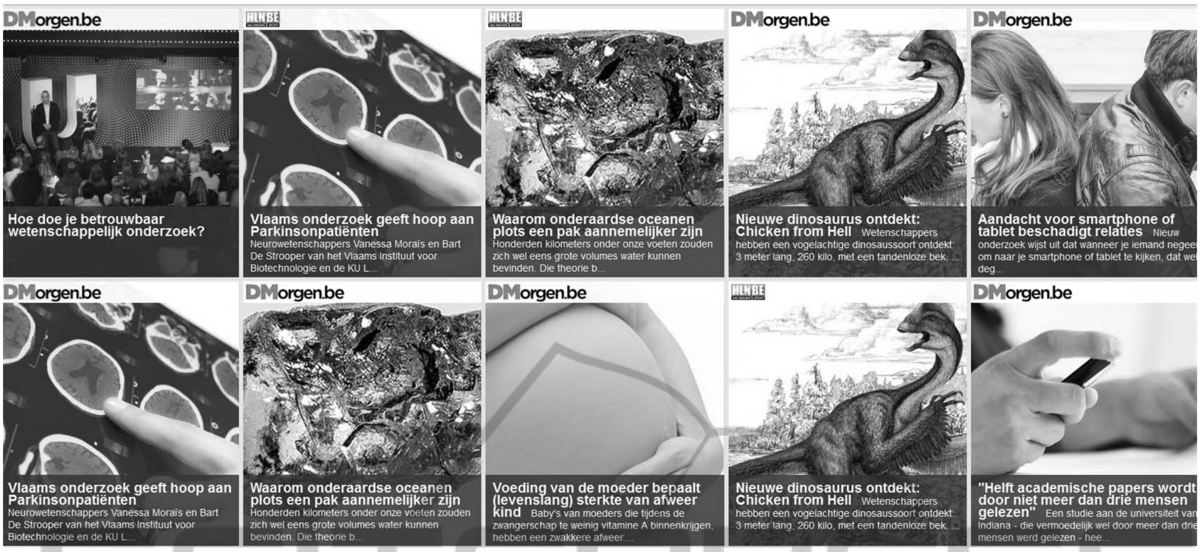

Figure 1: Screenshot of the user interface of the news service.

evaluated (explicit feedback). Subsequently, the algorithm determines the news items that best match the user's profile.

News items are characterized by eight different categories, the same as used in the questionnaire to elicit the explicit static preferences of the users. In addition, keywords provide more detailed info about the news items. These keywords are not predefined but are extracted from the text of the article using OpenCalais (Thomson Reuters, 2013). OpenCalais is a Web service that automatically creates rich semantic metadata for the content. It analyzes the news article and finds the entities within it, but it returns the facts and events hidden within the text as well. This way, the news article is tagged and analyzed with the aim of checking whether it contains information what the user cares about.

For the content-based recommendations, contextual aspects are not taken into account. So, contextual data, such as the device type and the time of the day, are ignored during the creation of the profile and the calculation of the recommendations.

4.3 Context-aware Content-based Recommendations

Just like the content-based recommendations, the context-aware content-based recommendations are not using a prior questionnaire but are self-learning based on the explicit and implicit feedback users provide during the experiment. For this type of recommendations, the InterestLMS algorithm of the Duine framework is extended to take into account the context of the user. Before generating the recommendations, the user feedback is processed by a contextual pre-filter (Adomavicius and Tuzhilin, 2008). Contextual information is used to determine the relevance of the feedback and filter these data based on the current situation. For instance, if a user wants to read news during the evening, an *exact pre-filter* (Adomavicius and Tuzhilin, 2011) selects only feedback gathered during the evening to calculate the recommendations. Therefore, the day is partitioned into four non-overlapping intervals: morning from 6:00 to 11:00, daytime from 11:00 to 12:00 and from 13:00 to 18:00, noon from 12:00 to 13:00 and evening/night from 18:00 to 6:00.

One major advantage of the contextual pre-filtering approach is that it allows deployment of any of the traditional recommendation techniques (Adomavicius and Tuzhilin, 2005). This makes it possible to use the same underlying algorithm for the context-aware content-based recommendations as for the content-based recommendations, which enables the comparison of both types of recommendations and to investigate the influence of contextual information.

Different pre-filtering techniques have proven their efficacy in literature (Baltrunas and Ricci, 2009). They all have to cope with the problem of context over-specification: focusing on the exact context is often a too narrow limitation. An overly specified context may not have enough training examples for accurately estimating the user's interests. For example, if a user rarely utilizes a tablet during noon to read news articles, the exact context (noon + tablet) may not provide enough data (feedback from the user) for accurately calculating the recommendations, which gives rise to the 'sparsity' problem. As a result, insufficient feedback is available for generating reliable recommendations (Papagelis et al., 2005).

An appropriate solution for context over-specification is to use a more general context specification by applying context generalization (Adomavicius and Tuzhilin, 2011). Since certain aspects of the overly specific context may be less significant, the data filtering can be made more general in order to retain more data after the filtering for calculating recommendations.

In this experiment, context generalization is applied in two phases in case of insufficient feedback data. In a first phase, the time frame is broadened. For instance, if recommendations are needed for a user who is reading news on a tablet during noon, the time restriction "noon" is dropped first. The data gathered in that specific context is supplemented with the user's feedback gathered on a tablet during other time periods. If the amount of feedback is still insufficient after this first generalization, the context is further generalized. In a second phase, the device type is broadened. More specifically, the user's feedback provided on a specific type of device (e.g., a tablet) is supplemented with the user's feedback provided on other device types (e.g., a smartphone). We opted to apply the generalization first on the time aspect of the context, and in a second phase on the device type, since many users are utilizing the service during different time periods but on only one type of device.

5 FEEDBACK ANALYSIS

To quantify the added value of a dynamic profile and contextual information, the users' interactions in terms of feedback with each of the three recommenders are analyzed. Figure 2 gives an overview of the user feedback on the news service during the evaluation period. The chart distinguishes implicit feedback, i.e. requesting to 'view' a news item, and explicit feedback, i.e., evaluating a news item by providing a 'Thumbs Up' or 'Thumbs Down' rating.

In Figure 2, this user feedback is aggregated over all users and partitioned by the type of recommendations that the users received. Since some test users dropped out just before the evaluation period, the different recommender types are not evaluated by the same number of test users. Table 1 shows the number of test users assigned to each type of recommendations, which has a direct influence on the total amount of feedback gathered for that type.

During the evaluation period, 2728 positive evaluations ('Thumbs Up') of a news recommendation were registered for all types of recommendations

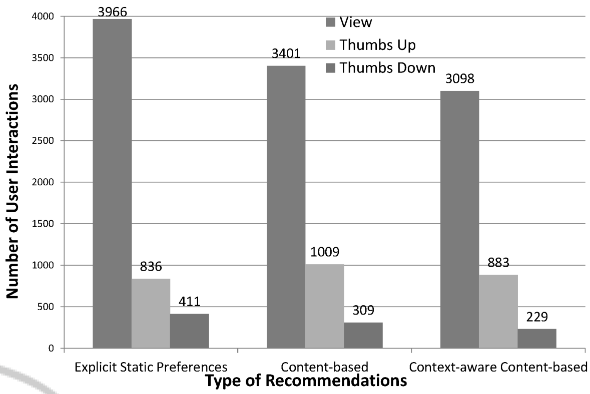

Figure 2: The amount of user interaction with the news service for each type of recommendations, partition according to the type of interaction.

together. In contrast, only 949 times a negative evaluation ('Thumbs Down') was provided by the users. These aggregated values ('Thumbs Up' 74.2% - 'Thumbs Down' 25.8%) are an indication for the general satisfaction of the users with the news that they get recommended.

However, significant differences for the different types of recommendations can be witnessed. Compared to the recommendations that are based on the explicit static preferences of the users, less views (total number, but also number per user) are obtained for the content-based and context-aware content-based recommendations. This might indicate that users get the interesting news more quickly using the more advanced algorithms, since they also spent more time per news item.

Comparing the different types of recommendations in terms of negative feedback ('Thumbs Down') demonstrates the added value of personal feedback and the context of the user for the recommender system. Recommendations based on explicit static preferences received 411 times 'Thumbs Down' from users who are not satisfied with the news content. The content-based recommender uses the implicit feedback (requests to view a news item) and explicit evaluations ('Thumbs Up & Down') as personal feedback during the evaluation period. Compared to the recommendations based on explicit static preferences, less negative (309 times 'Thumbs Down') and more positive evaluations are provided for the content-based recommendations. The lowest number of negative evaluations (229 times 'Thumbs Down') was received for the context-aware content-based recommender, which suggest news items based on the personal feedback of the user and by taking into account the user's current context. A Wilcoxon rank-sum test showed these differences are significant ($p = 0.043 < 0.05$).

Table 1 shows the ratio of the number of positive

Table 1: Comparison of the recommendation types.

Recommender Type	Explicit Static Preferences	Content-based	Context-aware Content-based
Input Data	Preliminary Questionnaire	Personal Feedback	Personal Feedback + Context
Number of Test Users	38	37	35
#Thumbs up / #(Thumbs up + down)	0.67	0.75	0.79

evaluations (# Thumbs Up) and the number of explicit evaluations (# Thumbs Up + Down) for the different types of recommendations. The results confirm the increase in accuracy by making the system dynamic (content-based recommendations) and taking into account the context (context-aware content-based recommendations).

6 DISCUSSION

Although the context-aware content-based recommendations received the most positive evaluation (highest ratio in Table 1), the difference with the traditional content-based recommendations is limited because of two reasons.

Firstly, many test users are not active in different contextual situations. They tend to read the news each day at the same time, using the same device. In this study, 10 users (9.1%) utilized the service during only one time period (morning, noon, daytime or evening), 17 users (15.5%) have viewed news items during two different time periods, 36 users (32.7%) were active during three different time periods, and 47 users (42.7%) have requested news content during morning, noon, daytime, and evening (all four time periods). In terms of time period, users were far most active during the evening (45.0% of the news request), followed by the morning (26.7%). Over the evaluation period of 5 weeks, only 47 users (42.7%) utilized the service at least once during all four time periods (morning, noon, daytime, and evening).

In terms of device types, 91 users (82.7%) utilized only one device to access the news content (either smartphone or tablet) and only 19 users (17.3%) used both devices at least once for reading news. In addition, users that used both device types may have the habit or preference to read news on only one of them. On the smartphone, news is more often consumed during different times of the day (76.8% used a smartphone during 3 or 4 time periods) than on tablets (56.7% used a tablet during 3 or 4 time periods). People generally keep their phone more closely throughout the whole day than their tablet, which is often used only once during the day, typically at home. Because some users consume news in fixed usage patterns, with limited variations in their context, the context-aware recommender cannot fully exploit the contextual information for this news service.

Secondly, the results are based on the evaluation period of approximately 5 weeks, and CARS require sufficient time to learn user preferences in different contextual situations. Because of the cold start problem (i.e. the issue that the system cannot draw any inferences for users or items about which it has not yet gathered sufficient information) and data fragmentation over the different contextual situations, we believe that there is still room for accuracy improvement of the context-aware content-based recommendations by gathering additional user feedback over a longer time period. The required number of ratings to overcome the cold start problem depends on various factors such as the algorithm parameters, the content domain, and the specific items that are rated. Studies have shown that, in general, more than 20 - 30 ratings are necessary for the system to recommend relevant items to the user (Lee et al., 2007). In our user study, some of the users did not achieve enough ratings for each contextual situations.

Additional data can help to learn patterns in the users' behavior and preference differences for various contextual situations, thereby further improving the accuracy of the context-aware content-based recommendations.

7 CONCLUSIONS

In this paper, a start-up news service offering personal recommendations is evaluated by an empirical user study. Three types of recommendations are tested: recommendations based on an explicit static profile, content-based recommendations using the actual user behavior but ignoring the context, and context-aware content-based recommendations incorporating user behavior as well as context. The study aimed to assess the importance of context in the recommender of a real-life mobile news service by focusing on two contextual aspects: device type and time. The results confirm the added value of contextual information for the recommender, but also indicate that the accuracy improvement might be limited in a real-life situation. Two conditions must be met in order to improve traditional recommenders with contextual information. Firstly, users have to utilize the service in different contextual situations. Logged user behavior showed

SCITEPRESS SCIENCE AND TECHNOLOGY PUBLICATIONS

that users typically access the news service with only one device (either smartphone or tablet) at different times of the day, but especially in the evening. Secondly, sufficient training data have to be available to learn user preferences and variations in these preferences over different contextual situations. In context-aware recommender systems, the cold start problem is strengthened by the fragmentation of the consumption data over different contextual situations. Context generalization can be a partial solution, but choosing the right contextual aspect to optimally broaden the context is often difficult.

ACKNOWLEDGEMENTS

This research was performed in the context of the iMinds-MIX Stream Store project, which is cofunded by iMinds, a research institute founded by the Flemish Government.

SCIENCE AND TECHNOLOGY PUBLICATIONS

REFERENCES

Adomavicius, G. and Tuzhilin, A. (2005). Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734–749.

Adomavicius, G. and Tuzhilin, A. (2008). Tutorial on context-aware recommender systems. In *Proceedings of the second ACM conference on Recommender Systems (RecSys '08)*.

Adomavicius, G. and Tuzhilin, A. (2011). Context-aware recommender systems. In Ricci, F., Rokach, L., Shapira, B., and Kantor, P. B., editors, *Recommender Systems Handbook*, pages 217–253. Springer US.

Baltrunas, L. and Ricci, F. (2009). Context-based splitting of item ratings in collaborative filtering. In *Proceedings of the Third ACM Conference on Recommender Systems*, RecSys '09, pages 245–248, New York, NY, USA. ACM.

Cantador, I., Bellogín, A., and Castells, P. (2008). News@hand: A semantic web approach to recommending news. In Nejdl, W., Kay, J., Pu, P., and Herder, E., editors, *Adaptive Hypermedia and Adaptive Web-Based Systems*, volume 5149 of *Lecture Notes in Computer Science*, pages 279–283. Springer Berlin Heidelberg.

De Pessemier, T., Dooms, S., and Martens, L. (2014). Context-aware recommendations through context and activity recognition in a mobile environment. *Multimedia Tools and Applications*, 72(3):2925–2948.

Følstad, A. (2008). Living labs for innovation and development of information and communication technology: A literature review. *Electronic Journal of Organizational Virtualness*, 10:99–131.

Han, B.-J., Rho, S., Jun, S., and Hwang, E. (2010). Music emotion classification and context-based music recommendation. *Multimedia Tools Appl.*, 47(3):433–460.

Hopfgartner, F., Kille, B., Lommatzsch, A., Plumbaum, T., Brodt, T., and Heintz, T. (2014). Benchmarking news recommendations in a living lab. In Kanoulas, E., Lupu, M., Clough, P., Sanderson, M., Hall, M., Hanbury, A., and Toms, E., editors, *Information Access Evaluation. Multilinguality, Multimodality, and Interaction*, volume 8685 of *Lecture Notes in Computer Science*, pages 250–267. Springer International Publishing.

Kenteris, M., Gavalas, D., and Mpitziopoulos, A. (2010). A mobile tourism recommender system. In *Proceedings of the The IEEE symposium on Computers and Communications*, ISCC '10, pages 840–845, Washington, DC, USA. IEEE Computer Society.

Lee, H., Kim, J., and Park, S. (2007). Understanding collaborative filtering parameters for personalized recommendations in e-commerce. *Electronic Commerce Research*, 7(3-4):293–314.

Li, L., Wang, D., Li, T., Knox, D., and Padmanabhan, B. (2011). Scene: A scalable two-stage personalized news recommendation system. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '11, pages 125–134, New York, NY, USA. ACM.

Papagelis, M., Plexousakis, D., and Kutsuras, T. (2005). Alleviating the sparsity problem of collaborative filtering using trust inferences. In Herrmann, P., Issarny, V., and Shiu, S., editors, *Trust Management*, volume 3477 of *Lecture Notes in Computer Science*, pages 224–239. Springer Berlin Heidelberg.

Reuters Institute for the Study of Journalism (2014). Digital News Report. Available at http://www.digitalnewsreport.org/.

Ricci, F. (2010). Mobile recommender systems. *Information Technology & Tourism*, 12(3):205–231.

Ricci, F. (2012). Contextualizing recommendations. In *ACM RecSys Workshop on Context-Aware Recommender Systems (CARS '12), In conjunction with the 6th ACM Conference on Recommender Systems (RecSys '12)*. ACM.

Said, A., Bellogín, A., and de Vries, A. (2013). News recommendation in the wild: Cwis recommendation algorithms in the nrs challenge. In *Proceedings of the 2013 International News Recommender Systems Workshop and Challenge. NRS*, volume 13.

Shani, G. and Gunawardana, A. (2013). Tutorial on application-oriented evaluation of recommendation systems. *AI Communications*, 26(2):225–236.

Telematica Instituut / Novay (2009). Duine Framework. Available at http://duineframework.org/.

Thomson Reuters (2008-2013). Open Calais. Available at http://www.opencalais.com/.

Weiss, A. S. (2013). Exploring news apps and location-based services on the smartphone. *Journalism & Mass Communication Quarterly*, 90(3):435–456.

Yu, Z., Zhou, X., Zhang, D., Chin, C.-Y., Wang, X., and men, J. (2006). Supporting context-aware media recommendations for smart phones. *Pervasive Computing, IEEE*, 5(3):68–75.