

Implicit User Profiling in News Recommender Systems

Jon Atle Gulla¹, Arne Dag Fidjestøl¹, Xiaomeng Su² and Humberto Castejon²

¹*Dep. of Computer and Information Science, Norwegian University of Science and Technology, Trondheim, Norway*

²*Telenor Group, Trondheim, Norway*

{jag, adf}@idi.ntnu.no, {xioameng.su, humberto.castejon}@telenor.com

Keywords: Recommender Systems, Personalization, User Profiling, Mobile News, Big Data, Information Retrieval.

Abstract: User profiling is an important part of content-based and hybrid recommender systems. These profiles model users' interests and preferences and are used to assess an item's relevance to a particular user. In the news domain it is difficult to extract explicit signals from the users about their interests, and user profiling depends on in-depth analyses of users' reading habits. This is a challenging task, as news articles have short life spans, are unstructured, and make use of unclear and rapidly changing terminologies. This paper discusses an approach for constructing detailed user profiles on the basis of detailed observations of users' interaction with a mobile news app. The profiles address both news categories and news entities, distinguish between long-term interests and running context, and are currently used in a real iOS mobile news recommender system that recommends news from 89 Norwegian newspapers.

1 INTRODUCTION

With the growing popularity of smartphones and tablets an abundance of news apps have been introduced over the last few years as alternatives to online web news sites. News aggregator apps like News360, Flipboard, Pulse and Feedly are not linked to one particular media house, but make use of publically available news stories as soon as they are published on the Internet. These mobile applications allow users to read neatly presented news stories from numerous media houses using interfaces that require only very limited interaction with the system. The features of these apps vary, though information filtering and user friendliness are key features of many of these applications (Haugen, 2013).

Advanced recommendation technologies are still rare in commercial news apps, but some of them now include simplified recommendation features, and several research prototypes experiment with new and promising recommendation approaches that analyze both the content of news articles and the users' social networks. This is not very surprising, as mobile news may benefit from these recommendation technologies for a number of reasons: (i) the constant flow of news easily leads to information overload problems for the reader, (ii) the small screen prevents the apps from showing

several stories or proper news overviews, and (iii) the lack of an efficient keyboard hampers the reader's interaction with the news apps.

However, an efficient recommender system needs some understanding of the users' preferences and interests. Some news apps do not keep any particular information about the user and there is no user-tailoring of news or recommendations offered. Other apps allow the user to select a category or post a query. The category or query associated with the user is then the simplest form of user profile found in mobile news apps. For full-fledged recommender systems, though, a more comprehensive representation of user interests and preferences is needed to provide personalized news services. If news stories read in the past is a good indication of the user's current preferences, we need to analyze these stories and find ways of capturing their content in her user profile. If stories read by friends or similar people are assumed to be relevant, we need ways of relating the user's profile to other users' profiles and find exactly those other users that have read news stories that are most likely to be of interest to her. There are many aspects of user profiles, and the choice of recommendation technique also influences the structure, content and maintenance of these profiles.

The NTNU SmartMedia project is experimenting with advanced recommendation techniques in a

news app for Norwegian newspapers. Users are in this app associated with devices, and the mobile news recommender engine ranks news using a combination of freshness, geographical proximity and match with users' preferences. Central in the work is the definition of user profiles that reflect and interpret user behavior and support both content-based and collaborative filtering.

The purpose of this study is to investigate to what extent a minimalistic swipe-based mobile news interface can be used to infer user preferences without any explicit signals from the users. The paper itself is structured as follows: After the introduction in Section 1, we discuss the particular challenges and needs in news recommender systems in Section 2. In Section 3 we briefly present earlier work on user profiling in recommender systems. Section 4 introduces the SmartMedia news recommendation project and its overall method for extracting user profiles from end user's interaction with the mobile app. After discussing the way this profiling technique can be configured to serve different purposes in Section 5, we conclude the paper in Section 6.

2 NEWS RECOMMENDATION

A news recommender system filters incoming news and presents to each individual user a ranked list of news that it deems most relevant to the particular user at hand. Since these are articles that the users have not already read and evaluated, we need to find a way of guessing users' interests based on their previously observed behavior.

In formal terms, the task in news recommendation is to estimate and rank the evaluations of news articles unknown to a user (Borges & Lorena, 2010; Jannach et. al, 2010). Assume a set of users U and a set of news articles A , the recommender system needs a utility function s that defines the evaluation v of an article a for a user u :

$$s: U \times A \rightarrow V,$$

in which V is a completely ordered set formed by non-negative values within an interval, e.g. 0 to 1 or 0 to 100. The system will recommend an article a' that maximizes the utility function for the user:

$$a' = \arg \max_{a \in A} s(u, a)$$

For those articles that have not been evaluated by the user, the system needs to estimate their evaluations from relatedness with other articles that have in fact

been read. Techniques like decision trees, Bayesian classifiers, support vector machines, singular value decomposition, clustering and various similarity scores have been used as part of this estimation process.

In content-based filtering, which is the primary concern of this paper, the estimation is all based on the content of previously read news articles. The assumption is that users read articles on topics they find interesting, and the users' interests do not change substantially from one day to another. If a particular user preferred to read about politics one day, chances are good that she will also be interested in political news the next day. Moreover, the degree of interest in a topic may be reflected in the frequency with which the user visits news stories of this topic or category.

2.1 News Challenges

Recommender systems for products and services normally deal with items that are rather stable both in terms of number and descriptions. The recommendation task is to compare a fixed number of items that all have structured and well understood properties. There is no relevant temporal dimension, and there is no problem separating one product or service from another.

The news domain is intrinsically more dynamic and unpredictable to deal with. There are new stories coming up all the time. Some of these uncover new events, while others just report on the progress of an already published event. There may be several conflicting stories of the same event, but there may also be several events discussed in the same story.

For news recommendation there are three particular challenges that need to be addressed:

First, the news domain is characterized by fluctuating and unclear vocabularies and ever changing news topics. Rather than ranking a fixed number of products, the system need to rank a dynamically growing number of events that may have very little in common with what the user read yesterday. The vocabularies change as new stories emerge, making it difficult to detect how stories are related and thereby estimate their relevance to the user.

Second, the short life span of news stories renders most news stories irrelevant, even though they seem consistent with the user's interests. For a developing story, it makes sense to recommend the latest article, unless there are substantial differences in quality and depth that make an older article more

informative. On the other hand, if collaborative filtering is used, it may be difficult to pick up the latest article before anyone else in your community has read and rated it.

Third, most media houses consider serendipity to be important to hold on to its readers. Serendipity refers to the introduction of news that may not have been selected or preferred by the user, but may still be interesting because they are surprising, alarming, important, have a particular journalistic quality, or simply add variation into the news stream. Serendipitous news help media houses expand users interests in their products, since they persuade users to read stories that are alien to them and that would never be read if the user only received stories consistent with their reading history.

While other recommender systems concentrate on the utility function itself, news recommender systems need also to consider content analysis techniques to extract and structure the topics pertaining to the news. They need to deal with ambiguities, uncertainties and inconsistencies that affect the way articles should be formally represented and compared to users' interest profiles.

3 RELATED WORK

A user profile is a representation of user's interests and preferences that is used to verify to what extent news stories are relevant to a particular user. The profiles are built for each individual user, are regularly updated, and describe topics, news categories and relevant features of the users.

In principle, there are two types of user profiles, *profiles based on implicit feedback* and *profiles based on explicit feedback* (Gauch et al, 2007; Lops et. al, 2011). Implicit user profiles are automatically extracted by the recommender systems themselves and may or may not be a correct representation of users' interests. In general, implicit feedback methods assign relevance scores to user actions on news articles like saving, sharing, bookmarking, etc. Explicit profiles are entered or approved by the users in question, but tend to be slightly less detailed than the implicit ones. The user-selected categories in Zite and Flipboard form very simple explicit user profiles, while the thumbs up/down approach in Zite and News360 provide explicit signals for what topics should be included in the profiles. However, adopting numeric or symbolic scales increases the users' cognitive load and may not be adequate for capturing emotions or attitudes towards the news. For mobile news recommendation it seems difficult

to require that the user enter and regularly update extensive representations of her interests, though, and more advanced techniques for profile construction on the basis of implicit feedback are needed.

A system combining explicit feedback and automatic learning is described in Singh et. al (2006). After building an initial interest category hierarchy on the basis of explicit feedback on a number of articles, the system analyzes user feedback from ongoing news sessions and automatically adds new leaf categories or update existing ones in the interest hierarchy. A similar approach is taken by Kim & Chan (2003), though they depend less on explicit user feedback.

Billsus and Pazzani (2000) have developed an approach in DailyLearner for interpreting implicit user feedback on news articles. A user click on the headline of an article is taken as a signal of interest, and an initial score of 0.8 is set. If the user is requesting more pages of the story, the score will be gradually increased until it reaches a maximum of 1.0 when all pages have been consulted. Similarly, a skipped article is assumed to be of no interest and is given a negative score that is subtracted from the system's prediction score for the article. All these scores are combined into a user profile that lists weighted informative words typical for the user's interests and preferences.

Liu, Dolan and Pedersen (2010) build user profile vectors that express users' interests in specific news categories over time. For each user they record the distribution of clicks and associate these click rates with categories on a monthly basis. This allows them to assess for every user the proportion of time spent on reading news from each category as well as to reflect on the development of her interests from one month to another.

In Cantador et. al (2008) they use semantic expansion in combination with a standard user preference algorithm for content-based recommendation. They observe that automatically learned profiles tend to be dominated by the main characteristics of user's preferences, preventing the recommendation engine from recommending news that are related albeit not directly addressed by the profile. Their solution is to use an ontology to include additional weighted concepts in the profile that are related to the original learned concepts or terms. Then they combine this expanded profile with a running context, which is a weighted set of concepts from the user's latest interactions with the system – to produce a contextualized version of the user's preferences that filters out topics that are out

of focus.

Implicit signals of user interests require that the user interface is designed in such a way that different levels of interests in a news article are reflected in the user's interaction with the system. This is a challenge in itself, as most news apps try to minimize the required interaction between system and reader. More interaction is good for analysing user interests, though it may lead to a less attractive user experience.

4 THE SMARTMEDIA PROJECT

The SmartMedia project at NTNU was initiated in 2011 as part of the university's collaboration with the Norwegian media industry and major Telecom companies. The project targets recommendation technologies for the mobile news domain and addresses content-based and collaborative techniques in combination with semantic and linguistic theories (in Gulla et. al (2014)).

Another important issue is the use of platforms, frameworks and methods for designing efficient cross-platform mobile user interfaces. Figure 1 shows the current mobile news recommender systems for the iOS platform. Details of this implementation can be found in Tavakolifard et. al (2013). A new version of the system in HTML5/JavaScript is under development and will introduce new features for hybrid news recommendation.




Figure 1: SmartMedia's news recommender app.

The current iOS news app indexes about 1,500 news articles per day. The articles are sourced from 89

Norwegian newspapers and are mostly in one of Norway's two official languages. During the indexing process, we extract and index the article's category (e.g. sports or lifestyle), its location (e.g. Oslo), and relevant keywords characterizing the article's content. The keywords are partly named entities (e.g. Lionel Messi) from a NER (named entity recognition) component, partly important nominals that refer to concepts discussed in the article. The underlying news index is in SolR, an open source search platform from the Apache Lucene project, and is built up as a traditional large-scale search engine index (see for example Gulla et. al (2002) and Solskinnsbakk & Gulla (2010)).

5 USER PROFILING

In SmartMedia's user profiling component a set of 10 user acts have been identified as implicit signals of user interest. These are *Opened article view*, *Article view time*, *Preview time*, *Clicked category*, *Shared twitter*, *Shared facebook*, *Shared mail*, *Starred article*, *Viewd map*, and *Viewed similar article*. We monitor the time spent reading a news story or looking at its preview, assuming that large view dwell times indicate a stronger interest in the topic.

The construction of a user profile for a particular user follows a three-step process:

1. *Extract interests from user acts*
2. *Build running context from all user acts*
3. *Combine running context and long-term interests into new user profile*

In the following, we go into detail of each of these three steps.

5.1 User Acts

Whenever the user is interacting with the news app, her actions are logged and stored with all relevant data associated with the act. This includes not only the user ID and the act itself, but also the time of the act, the location of the news reported in the article, the news category, as well as entities and important key phrases describing the content of the article.

Figure 2 shows how a particular user act is represented in the system. As seen from the timestamp, this is the act of reading a particular news story at 16:41 on 2nd June 2013. The user spent 1.4 seconds reading it, and the article itself was a news article about an Indian man that resisted an arrest in Kristiansand in Norway. The content is

revealed from the keywords (named entities and important nominal phrases) that are listed under the ‘tags’ label.

As different types of user acts may signal different levels of user interest, each type of act t is associated with a constant α_t . For timed user acts, like reading a preview or the full version of a news story, we only want to consider this reading if it is a significant act, i.e, she spends more time than usual on the story.

```
{
  "_id" : "241B50BE-DF5-4AAB-A12D-98D4A4606028" ,
  "articleId" : "318218311" ,
  "userId" : "bf4d2b7adec01da0ddc8c3317088bcdc6" ,
  "eventType" : "TIME_SPENT_ARTICLE_VIEW" ,
  "timestamp" : { "$date" : "2013-06-02T16:41:15.511Z" } ,
  "geoLocation" :
    { "name" : "" ,
      "type" : "" ,
      "longitude" : 8.00354 ,
      "latitude" : 58.138821 } ,
  "properties" : { "duration" : "1.427272" } ,
  "tags" :
    [ "agder politidistrikt" ,
      "havnet" ,
      "satt" ,
      "rebelsk mannen" ,
      "operasjonsleder" ,
      "kristiansund slo" ,
      "politiet" ,
      "kristiansand skallet" ,
      "Maharashtra" ,
      "India" ,
      "Egersund" ,
      "Rogaland" ,
      "Norway" ] ,
  "categories" : [ "NEWS" ]
}
```

Figure 2: User act logging.

At the content level, we now assume that the interests of the user are reflected by the important keywords found in the article she has read, modified by the importance of the user act. We model this as a *user act interest vector* $\vec{E}_i = (e_1, \dots, e_k)$ for act i , where e_j is defined as

$$e_j = \begin{cases} 0 & \text{if act } i \text{ is insignificant timed act} \\ \alpha_t * f_j & \text{otherwise} \end{cases}$$

α_t denotes the importance of user acts of type t , while f_j is the frequency of keyword j . In the current implementation α_t is set to 1 for all types t , though experiments show that some of these acts are more significant than others.

A similar analysis is used for categories, giving us also a user act category vector.

5.2 Running Context

A running context for a user is built from all user acts of that user that have happened after the last time a full user profile was generated. It describes

the overall topics of what the user has been clicking on or reading lately without reflecting what she might have been interested in at earlier occasions. The running context gives us the user’s current news focus.

Formally, the *running content context* \vec{V}_R is defined as follows:

$$\vec{V}_R = \sum_{i \in T} E_i$$

where $i \in T$ are the user acts in the time span T after the last time her profile was generated.

A similar running category context is also defined.

Implementationally, the mobile news app records every gesture from the user and maintains an updated running context at all times. Whenever the user decides to reset his context or consult her stored profile, the running context is compared with her old user profile on the server side. When the user ends her session, the running context is incorporated into her long-term user profile, so that she will start with an updated user profile and an empty running context next time.

5.3 User Profile

Our user profiles combine long-term interests with users’ focus on current news stream. The idea is that the user may be interested in following news events that are unfolding, even though they are not necessary perfectly in line with her general interests and preferences.

For each user the system stores an old user profile \vec{V}_0 . The new user profile \vec{V}_U is calculated from the old profile and the running context as follows:

$$\vec{V}_U = c\vec{V}_0 + (1 - c)\vec{V}_R$$

If c is set to 1, the user profile does not take into account the running context and will never change. A value of 0 makes the old profile irrelevant, and only the running context is used to recommend news stories to the user. I practice the constant c should be carefully selected to both address the users’ need for relevant news and the content providers’ need to promote the latest news.

As before there is a similar calculation for constructing the user profile at category level.

It is important to note that we also assume that long-term interests fade unless they are renewed when similar topics come up in the news stream. Since c will always be below 1, the features of old

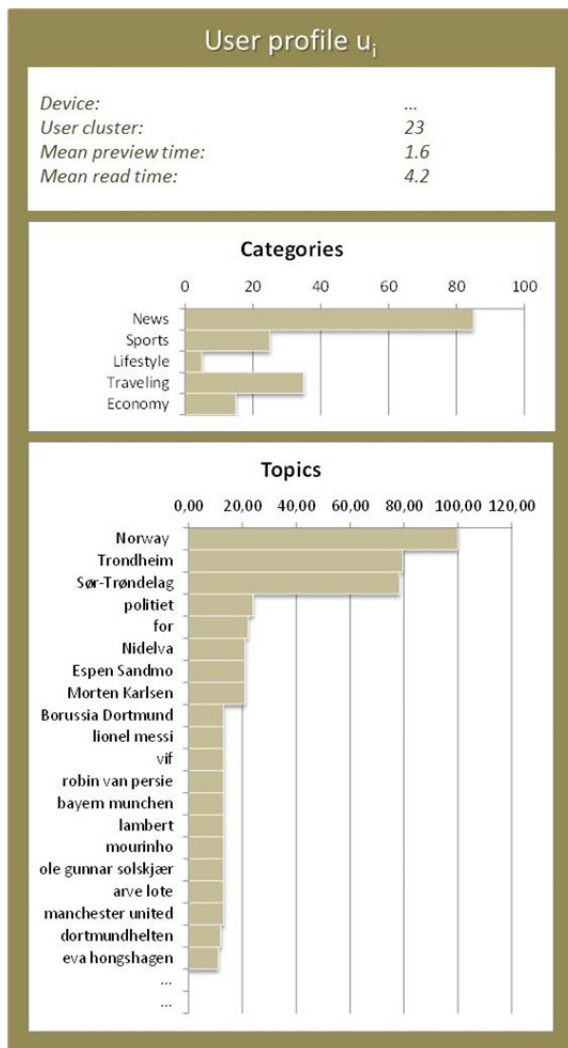


Figure 3: Generated user profile.

user profiles gradually disappear as they are updated with running contexts. The lower the c value is, the faster the system will forget the user’s old profile.

Figure 3 shows a generated user profile for a user that has just started using the system. Both the category part and the content part of the profile are shown. The category part tells us that this user is mostly interested in news, though she has also enjoyed some stories about traveling and sports. At the content level, we notice that the user has preferred local news, as *Norway*, *Trondheim* and *Sør-Trøndelag* are all location names relevant to his home place. The rest of the profile exhibits some sports personalities as well as culture persons in the Norwegian media. At the top of the profile, there are two variables, *mean preview time* and *mean read time*, that are used to determine if a timed user act is significant or not. If the user spends less than 4.2

seconds reading a full news story, the system will not consider this reading act relevant to the calculation of her user profile.

6 DISCUSSION

The SmartMedia news recommender app sources news from all major newspapers in Norway. Some of these news come with meta data like category and tags, though most news stories are just unstructured of texts that need to be analyzed to locate the news, categorize the stories and extract informative keywords that describe their content. Since these keywords are critical for both the construction of profiles and subsequent recommendation process, their quality and availability are important to the recommendation process.

With the exception of finance articles, news stories in Norway are on average 220 words (tokens) long. Finance news tend to be more than 1,000 words long and written in a more factual and objective manner.

On the average our news articles contain the following named entities:

Named entity	Occurrences per story
<i>Locations</i>	1.6
<i>Persons</i>	2.3
<i>Organizations</i>	2.3
<i>Roles</i>	0.8

Since there are so few named entities in news stories, our user profiles also include important nominal phrases that constitute more concepts than references to particular entities.

The evaluation of user profiles has so far focused on the contribution of each type of user acts on the quality of the profile (Nilsen, 2013). In one experiment we used sensitivity analyses to assess the influence of each type of user acts, and to what extent they are consistent with each other. The results show that the *Opened article view* act in many cases provided a good average of the other acts taken together, though the other acts were needed to deal with particularly interesting or uninteresting news stories. Another experiment analyzed the impact on user profile size. Not surprisingly, the *Preview time* act accounted for 82% of the profile vectors’ contents after some time of regular use. This may be an indication that our threshold for taking this act into account is set to low or has to reflect a more dynamic rule than just accepting any reading act longer than the average one. Quite likely the lengths of the news stories

need to be a factor when considering the reading time.

Experiments with the news recommender system shows that there are three issues that need to be carefully considered when user profiles are constructed:

- **Granularity.** In the current implementation we consider category interests and content interests equally important. This is a choice that seems reasonable to most users, though there are many users that would prefer a stronger focus on either the category as a whole (“i like all sports”) or on its particular topics (“I don’t care about football, but like Messi and Barcelona”).
- **Persistence.** The balance between stable long-term user interests and short-term news context is delicate. If long-term interests are preferred, the user risks that there is no relevant news available and she will leave the app. Emphasizing the running context too much is also unfortunate, as this will make the app appear similar from one user to another.
- **Serendipity.** There is also a balance between profile-relevant news stories and serendipitous news. Most content providers like to include also news that are not directly within a user’s profile to trigger new interests and widen her perspective (Das et. al, 2007). This is in our case achieved using content-based filtering for profile-driven news and collaborative filtering for serendipitous news, and the weighting of the techniques decide the balance between these two kinds of news.

In sum, it is difficult to identify an ideal user profile in news recommender systems. Users have different priorities than content providers, and the various weighting schemes that may be used to address the granularity, persistence and serendipity of news typically depend on whether they reflect the users’ or the content providers’ perspective.

It is also worth noting that freshness and location seem very important to users, even if the stories themselves are not necessarily the best match with their user profiles. In this respect news recommendation differs fundamentally from product recommendation on online shopping sites, where users have a clear and stable interest that calls for all relevant products independently of age or origin.

Our profile construction approach is in many ways similar to what is used in SCENE (Li et al., 2011). However, whereas we combine content phrases and entities in one vector, they split it up and also consider more complex phrases as part of their content description. A more fundamental difference,

though, is the granularity of the user click analysis itself. As opposed to SCENE and most other news apps that look at news stories as either read or not read, we consider in more detail what the user is doing with the stories and use these smaller acts as cues for refining the users’ real interests and preferences. This technique is to some extent also exploited in DailyLearner (Billisus and Pazzani, 2000), though the SmartMedia app takes every user gesture as an indication of user satisfaction or dissatisfaction.

7 CONCLUSIONS

The Smartmedia mobile news recommender system demonstrates both the value and serious challenges of implicit user profiling. Our mobile news recommender system uses extensive user profiles to recommend news stories that deal with the same topics as the profiles. However, the system is highly configurable, with a number of parameters that seriously affect the news stories recommended to the users. As there is no obvious best configuration of the system, since users and content providers have different agendas, only time will tell what is an acceptable weighting scheme for both readers and media houses.

The news app has been available online since the summer of 2013, and we are now in the process of building a new mobile user interface based on feedback from the early months. We are also – together with the media industry – collecting a large training set with real news and real users over several months, which we intend to use to refine our weighting schemes for balancing different recommendation strategies.

Research-wise we are gradually expanding the recommender system with semantic features for modeling news events and the entities involved in these events. This involves the definition of news taxonomies not very different from what is used in semantic search solutions (see e.g. Brasethvik & Gulla, 2002). With a new log-in feature we also intend to use social media sites like Twitter to expand and deepen the understanding of users’ interests and preferences (e.g. O’Banion, 2012).

ACKNOWLEDGEMENTS

This research was supported by Telenor Group as part of their collaboration with the Department of

Computer and Information Science at the Norwegian University of Science and Technology in Trondheim.

REFERENCES

- Billsus, D. and Pazzani, M. J., 2000. User Modeling for Adaptive News Access. *User Modeling and User-Adapted Interaction*, 10, pp. 147-180.
- Borges, H. L. and Lorena, A. C., 2010. A Survey of Recommender Systems for News Data. In Szczerbicki & Nguyen (eds.), *Smart Information and Knowledge Management*, SCI 260, pp. 129-151. Springer.
- Brasethvik, T. and Gulla, J. A., 2002. A conceptual modeling approach to semantic document retrieval. In Proceedings of the 14th international Conference on *Advanced Information Systems Engineering (CAiSE'02)*, pp. 167-182. Springer.
- Cantador, I. Bellogin, A. and Castells, P., 2008. Ontology-Based Personalised and Context-Aware Recommendations of News Items. In Proceedings of the 7th International Conference on Web Intelligence, pp. 562-565. IEEE.
- Das, A. S. Datar, M. Garg, A. and Rajaram, S., 2007. Google news personalization: scalable online collaborative filtering. In *Proceedings of the 16th international conference on World Wide Web*, pp. 271-280. ACM.
- Gauch, S., Speretta, M., Chandramouli, A., and Micarelli, A., 2007. User profiles for personalized information access. *The adaptive web*, pp. 54-89. Springer.
- Gulla, J. A. Auran, P. G. and Risvik, K. M., 2002. Linguistic Techniques in Large-Scale Search Engines. In *Proceedings of the 6th International Conference on Applications of Natural Language to Information Systems (NLDB'02)*, pp. 218-222.
- Gulla, J. A., Ingvaldsen, J. E., Fidjestøl, A. D., Nilsen, J. E., Haugen, K. R., Su, X., 2014. Learning User Profiles in Mobile News Recommendation. Accepted for publication in *Journal of Print and Media Technology Research*.
- Haugen, K. R., 2013. *Mobile News: Design, User Experience and Recommendation*. MSc thesis. NTNU, Trondheim.
- Jannach, D. Zanker, M. Felfernig, A. and Friedrich, G., 2010. *Recommender Systems: An Introduction*. Cambridge University Press.
- Kim, H. R. and Chan, P. K., 2003. Learning implicit user interest hierarchy for context in personalization. In *Proceedings of the 8th international conference on Intelligent user interfaces*, pp. 101-108. ACM.
- Li, L., Wang, D., Li, T., Know, D., and Padmanabhan, B., 2011. SCENE: a scalable two-stage personalized news recommendation system. In *Proceedings of SIGIR'11*, pp. 125-134. ACM.
- Liu, J. Dolan, P. and Pedersen, E. R., 2010. Personalized news recommendation based on click behavior. In *Proceedings of the 15th international conference on intelligent user interfaces*, pp. 31-40. ACM.
- Lops, P., de Gemmis, M. and Semeraro, G., 2011. Content-based Recommender Systems: State of the Art and Trends. In Ricci, Rokach, Shapira and Kantor (Eds.), *Recommender Systems Handbook*, Chapter 3, pp. 73-106. Springer.
- Nilsen, J. E., 2013. Large-Scale User Click Analysis in News Recommendation. MSc thesis, NTNU, Trondheim.
- O'Banion, S. Birnbaum, L. and Hammond, K., 2012. Social media-driven news personalization. In *Proceedings of the 4th ACM RecSys workshop on Recommender systems and the social web*. pp. 45-52. ACM.
- Rajaraman, A. and Ullman, J. D., 2011. *Mining of Massive Datasets*. Cambridge University Press.
- Singh, S., Shepherd, M., Duffy, J. and Watters, C., 2006. An Adaptive User Profile for Filtering News Based on a User Interest Hierarchy. In Proceedings of the American Society for Information Science and Technology, Volume 43, Issue 1, pp. 1-21, 2006.
- Solskinnbakk, G. and Gulla, J. A., 2010. Combining ontological profiles with context in information retrieval. *Data & Knowledge Engineering*, 69(3), pp. 251-260.
- Tavakolifard, M. Gulla, J. A. Almeroth, K. C. Ingvaldsen, J. E. Nygreen, G. and Berg, E., 2013. Tailored News in the Palm of your HAND: A Multi-Perspective Transparent Approach to News Recommendation. In *Proceedings of 22nd International World Wide Web Conference (WWW'13), Companion Volume*, pp. 305-308, May, Rio de Janeiro.