

7th International Workshop on News Recommendation and Analytics (INRA 2019)

in conjunction with 13th ACM Conference on Recommender Systems (RecSys 2019), 16-20 September, Copenhagen, Denmark

Özlem Özgöbek

Norwegian University of Science and Technology
Trondheim, Norway
ozlem.ozgobek@ntnu.no

Jon Atle Gulla

Norwegian University of Science and Technology
Trondheim, Norway
jon.atle.gulla@ntnu.no

Benjamin Kille

Institute of Technology Berlin
Berlin, Germany
benjamin.kille@tu-berlin.de

Andreas Lommatzsch

Institute of Technology Berlin
Berlin, Germany
andreas.lommatzsch@dai-labor.de

1 INTRODUCTION

The 7th International Workshop on News Recommendation and Analytics (INRA 2019)¹ is held in conjunction with in conjunction with 13th ACM Conference on Recommender Systems (RecSys 2019), 16-20 September, Copenhagen, Denmark. This workshop aims to bring researchers, media companies, and practitioners together, in order to exchange ideas about how to create and maintain a trusted and sustainable environment for digital news production and consumption. This version of INRA workshop series includes a keynote speaker and 10 peer reviewed papers, where each paper have been reviewed by at least two program committee members. INRA 2019 have received 16 submissions and has an acceptance rate of 62.5%.

In INRA 2019, thinking of creating a more interactive workshop setting, we have introduced a poster session. All the accepted papers' authors have given the chance to display their works as a poster during the workshop. More than half of the authors had a poster and we have observed interactions between the authors and the participants during the half an hour break. During the call for papers of INRA 2019, we have provided the researchers access to several data sets and an evaluation platform for news recommendations, in case they would like to test their systems by using them [6]. Unfortunately, we have not received any submissions using these data sets and platforms.

In this year's edition, we mainly focus on three categories: News recommendation, news analytics, and ethical aspects of news recommendation. More information can be found in [6].

2 WORKSHOP DETAILS

2.1 Keynote Speech

Democracy, Diversity and Design - Sharing experiences from an interdisciplinary project

Dr. Natali Helberger, University of Amsterdam

Abstract: News-recommender systems, which automatically select the content of newsletters, personalized news apps or social-media news feeds are playing an increasingly critical role in helping users to filter and sort information. And as such are fulfilling a crucial role in democratic society. Data analytics and recommender systems are going to be more and more pivotal in deciding what kind of news the public does and does not see. Depending on their design, recommenders can either unlock the diversity of online information for their or lock them into so-called filter bubbles. The challenge for the development of diversity-sensitive recommenders is defining what diversity in recommendations actually means. Often conceptualised as a measure of variance or even serendipity, diversity is an inherently normative concept, deeply rooted in democratic theory and our ideas of what it means to live in a democratic society. Funded by the SIDN fonds, a team of legal scholars, communication and computer scientists from the University of Amsterdam and RTL have worked on a project that translates insights from democratic theory into concrete metrics that can help to assess the performance of news recommenders. Condensing a concept that is as vague and colourful as diversity into a number of concrete metrics is not a trivial task. In my keynote I would like to present some of our work, and draw some lessons for future work on 'diversity by design'.

2.2 Accepted Papers

- **Public Service Media, Diversity and Algorithmic Recommendation: Tensions between Editorial Principles and Algorithms in European**

¹<http://research.idi.ntnu.no/inra/2019>

PSM Organizations, Jannick Kirk Sørensen

Abstract: Public Service Media (PSM) websites are an interesting case for the implementation of recommender systems for media personalization, as the PSM organizations need to balance the optimization of exposure with traditional but ill-defined PSM policy goals such as fairness, viewpoint diversity and transparency. Furthermore, the mathematical logic of recommender system needs to be adapted to the legacy broadcasting scheduling and publishing strategies and procedures. Finally, as the PSM organizations step into new territories, domestication and adaption of the recommender system technologies must take place while PSM organizations try to embrace the new knowledge and new professions associated with recommender systems. Based on 25 in-depth interviews conducted from December 2016 to April 2019, this paper presents a cross European analysis of the implementation of recommender systems in nine European public service media organizations from eight countries. The findings indicate that PSM organizations, although viewing personalisation as competitive necessity, approach recommendation systems with hesitation in order to maintain core PSM-values in the online environment. Furthermore, although the collaborative filtering chosen by the PSM organizations indicate a user-centered approach, curation systems on top of recommender systems re-install a broadcaster-centric approach.

- **Semi-supervised sentiment analysis for under-resourced languages with a sentiment lexicon, Peng Liu, Cristina Marco and Jon Atle Gulla**

Abstract: This paper presents the results of using semi-supervised sentiment analysis on an under-resourced language such as Norwegian. To perform these experiments, two external resources have been used: an available training corpus containing Norwegian reviews from major newspaper sources (NoRec), and a newly created general sentiment lexicon for Norwegian. The results of our experiments show that the performance improves significantly when the sentiment lexicon is used. Besides, the best results are obtained using Support Vector Machines (SVM) as the machine learning algorithm used for training with an AUC score of around 92%. An alternative statistical measure was used for evaluation, Area Under ROC Curve (AUC), in order to deal with the highly imbalanced nature of the dataset.

- **On the Importance of News Content Representation in Hybrid Neural Session-based Recommender Systems, Gabriel De Souza P. Moreira, Dietmar Jannach and Adilson Marques Da Cunha**

Abstract: News recommender systems are designed to surface relevant information for online readers by personalizing their user experiences. A particular problem in that context is that online readers are often anonymous, which means that this personalization can only be based on the last few recorded interactions with

the user, a setting named session-based recommendation. Another particularity of the news domain is that constantly fresh articles are published, which should be immediately considered for recommendation. To deal with this item cold-start problem, it is important to consider the actual content of items when recommending. Hybrid approaches are therefore often considered as the method of choice in such settings. In this work, we analyze the importance of considering content information in a hybrid neural news recommender system. We contrast content-aware and content-agnostic techniques and also explore the effects of using different content encodings. Experiments on two public datasets confirm the importance of adopting a hybrid approach. Furthermore, we show that the choice of the content encoding can have an impact on the resulting performance.

- **Defining a Meaningful Baseline for News Recommender Systems, Benjamin Kille and Andreas Lommatzsch**

Abstract: Evaluation protocols for news recommender systems typically involve comparing the performance of methods to a baseline. The difference in performance ought to tell us what benefit we can expect from using a more sophisticated method. Ultimately, there is a trade-off between performance and effort in implementing and maintaining a system. This work explores what baselines have been used, what criteria baselines must fulfil, and evaluates a variety of baselines in a news recommender evaluation setting with multiple publishers. We find that circular buffers and trend-based predictions score highly, need little effort to implement, and require no additional data. Besides, we observe variations among publishers, suggesting that not all baselines are equally competitive in different circumstances.

- **On-the-Fly News Recommendation Using Sequential Patterns, Mozghan Karimi, Boris Cule and Bart Goethals**

Abstract: The news recommendation problem poses a number of specific challenges that established recommendation techniques, successful in other settings, do not tackle adequately. For example, unlike in other domains, the relevance of news articles drops significantly over time, and the order in which users visit news articles matters greatly. Furthermore, in the context of breaking news, user interests can change rapidly, and there is a need to generate recommendations on-the-fly, taking into account recently published articles and the latest trends among users' preferences. To address these issues, we use a form of sequential pattern mining to generate up-to-date news recommendations on a click-by-click basis. In this approach, patterns are mined incrementally from the incoming clickstream so that new items and trends are considered. Our experimental

evaluation demonstrates that our method compares favorably with existing techniques and outperforms them on a variety of metrics.

- **Giveme5W1H: A Universal System for Extracting Main Events from News Articles, Felix Hamborg, Corinna Breitingner and Bela Gipp**

Abstract: Event extraction from news articles is a commonly required prerequisite for various tasks, such as article summarization, article clustering, and news aggregation. Due to the lack of universally applicable and publicly available methods tailored to news datasets, many researchers redundantly implement event extraction methods for their own projects. The journalistic 5W1H questions are capable of describing the main event of an article, i.e., by answering who did what, when, where, why, and how. We provide an in-depth description of an improved version of Giveme5W1H, a system that uses syntactic and domain-specific rules to automatically extract the relevant phrases from English news articles to provide answers to these 5W1H questions. Given the answers to these questions, the system determines an article's main event. In an expert evaluation with three assessors and 120 articles, we determined an overall precision of $p=0.73$, and $p=0.82$ for answering the first four W questions, which alone can sufficiently summarize the main event reported on in a news article. We recently made our system publicly available, and it remains the only universal open source 5W1H extractor capable of being applied to a wide range of use cases in news analysis.

- **Recommendation systems for news articles at the BBC, Maria Panteli, Alessandro Piscopo, Adam Harland, Jonathan Tutchner and Felix Mercer Moss**

Abstract: Personalised user experiences have improved engagement in many industry applications. When it comes to news recommendations, and especially for a public service broadcaster like the BBC, recommendation systems need to be in line with the editorial policy and the business values of the organisation. In this paper we describe how we develop recommendation systems for news articles at the BBC. We present three models and describe how they compare with baseline approaches such as random and popularity. We also discuss the metrics we use, the unique challenges we face and the considerations needed to ensure the recommendations we generate uphold the trust and quality standards of the BBC.

- **Trend-responsive user segmentation enabling traceable publishing insights. A case study of a real-world large-scale news recommendation system, Joanna Misztal-Radecka, Dominik Rusiecki, Michał Żmuda and Artur Bujak**

Abstract: The traditional offline approaches are no

longer sufficient for building modern recommender systems in domains such as online news services, mainly due to the high dynamics of environment changes and necessity to operate on a large scale with high data sparsity. The ability to balance exploration with exploitation makes the multi-armed bandits an efficient alternative to the conventional methods, and a robust user segmentation plays a crucial role in providing the context for such online recommendation algorithms. In this work, we present an unsupervised and trend-responsive method for segmenting users according to their semantic interests, which has been integrated with a real-world system for large-scale news recommendations. The results of an online A/B test show significant improvements compared to a global-optimization algorithm on several services with different characteristics. Based on the experimental results as well as the exploration of segments descriptions and trend dynamics, we propose extensions to this approach that address particular real-world challenges for different use-cases. Moreover, we describe a method of generating traceable publishing insights facilitating the creation of content that serves the diversity of all users needs.

- **Enriched Network Embeddings for News Recommendation, Janu Verma**

Abstract: News aggregators collect content from various sources and present them in one website or mobile application for easy access. A key challenge for the news applications is to help users discover relevant articles. Both the user experience and the key metrics depend on the high-quality personalized recommendations. However, building a news recommendation presents a set of challenges due to the large number of articles being published every hour, the surge and decline in the popularity of news, and critical nature of recency etc. In this paper, we present a graph-based news recommendation model which is deployed on a real-world news application. Our system is a hybrid of collaborative-filtering and the content-based filtering. We enrich the user-article interaction graph by adding new nodes corresponding to the named entities extracted from the contents of the articles. The random walk based graph embeddings are used to learn latent representation for users, articles and named entities in the same space. We evaluate the learned embeddings via a multi-class classification of news articles into high-level categories. We propose a recommendation system based on the binary classification problem which takes as input a combination of the user, item and entity embeddings and computes the probability of the user clicking on the article. We perform experiments to show the superiority of our model to the previous system.

- **Leveraging Emotion Features in News Recommendations, Nastaran Babanejad, Ameeta Agrawal, Heidar Davoudi, Aijun An and Manos Papageorgis**

Abstract: Online news reading has become very popular as the web provides access to news articles from millions of sources around the world. As a specific application domain, news recommender systems aim to give the most relevant news article recommendations to users according to their personal interests and preferences. Recently, a family of models has emerged that aims to improve recommendations by adapting to the contextual situation of users. These models provide the premise of being more accurate as they are tailored to satisfy the continuously changing needs of users. However, little attention has been paid to the emotional context and its potential on improving the accuracy of news recommendations. The main objective of this paper is to investigate whether, how and to what extent emotion features can improve recommendations. Towards that end, we derive a large number of emotion features that can be attributed to both items and users in the domain of news. Then, we devise state-of-the-art emotion-aware recommendation models by systematically leveraging these features. We conducted a thorough experimental evaluation on a real dataset coming from news domain. Our results demonstrate that the proposed models outperform state-of-the-art non-emotion-based recommendation models. Our study provides evidence of the usefulness of the emotion features at large, as well as the feasibility of our approach on incorporating them to existing models to improve recommendations.

2.3 Previous Workshops

7th International Workshop on News Recommendation and Analytics (INRA 2019) is based on the following previous workshops:

- International News Recommender Systems Workshop and Challenge (NRS) ² held in conjunction with the 7th ACM Recommender Systems Conference in 2013. This workshop had a minimal scope, which restricted the number of submissions and led to an acceptance rate of 75%.
- International Workshop on News Recommendation and Analytics (NRA) 2014 ³ held in conjunction with 22nd Conference on User Modelling, Adaptation and Personalization (UMAP) in 2014. In this workshop, we have expanded the scope with news analytics, which is closely linked with news recommendation. This expansion of the scope led to more submissions and a 50% acceptance rate. [3]
- 3rd International Workshop on News Recommendation and Analytics (INRA) 2015 ⁴ held in conjunction with ACM RecSys 2015 Conference in September 2015, Vienna, Austria. Acceptance rate is 66%. [4]

²<http://recsys.acm.org/recsys13/nrs>

³<http://research.idi.ntnu.no/nra2014>

⁴<http://research.idi.ntnu.no/inra/2015>

- 4th International Workshop on News Recommendation and Analytics (INRA) 2016 ⁵ held in conjunction with 24th Conference on User Modeling, Adaptation and Personalization (UMAP 2016), Halifax, Canada. Acceptance rate is 75%. [2]
- 5th International Workshop on News Recommendation and Analytics (INRA) 2017 ⁶ held in conjunction with IEEE/WIC/ACM International Conference on Web Intelligence (WI), 23-26 August 2017, Leipzig, Germany. Acceptance rate is 70% [1]
- 6th International Workshop on News Recommendation and Analytics (INRA 2018) ⁷ held in conjunction with CIKM 2018. 22-26 October 2018, Turin, Italy. Acceptance rate is 75%. [5]

3 ORGANIZATION

3.1 Workshop Chairs

- **Özlem Özgöbek**, Department of Computer and Information Science, Norwegian University of Science and Technology (NTNU), Norway
- **Benjamin Kille**, Institute of Technology Berlin, Germany
- **Jon Atle Gulla**, Department of Computer and Information Science, Norwegian University of Science and Technology (NTNU), Norway
- **Andreas Lommatzsch**, Institute of Technology Berlin, Germany

3.2 Program Committee Members

- Alejandro Bellogin, Universidad Autónoma de Madrid (UAM), Spain
- Andreas Lommatzsch, Technische Universität Berlin, Germany
- Asbjørn Følstad, SINTEF, Norway
- Benjamin Kille, Technische Universität Berlin, Germany
- Cristina Marco, Amazon Alexa, Turin, Italy
- Frank Hopfgartner, Information School of University of Sheffield, UK
- Lemei Zhang, Norwegian University of Science and Technology, Norway
- Mozghan Karimi, University of Antwerp, Belgium
- Özlem Özgöbek, Norwegian University of Science and Technology, Norway
- Peng Liu, Norwegian University of Science and Technology, Norway
- Shumpei Okura, Yahoo! Reserach Japan

REFERENCES

- [1] 2017. *WI '17: Proceedings of the International Conference on Web Intelligence, SESSION: INRA*. ACM, New York, NY, USA.
- [2] Jon Atle Gulla, Luc Martens, Özlem Özgöbek, Nafiseh Shabib, and Toon De Pessemier. 2016. 4nd International Workshop on

⁵<http://research.idi.ntnu.no/inra/2016>

⁶<http://research.idi.ntnu.no/inra/2017>

⁷<http://research.idi.ntnu.no/inra>

- NewsRecommendation and Analytics (INRA2016). http://ceur-ws.org/Vol-1618/INRA_preface.pdf
- [3] Jon Atle Gulla, Ville Ollikainen, Özlem Özgöbek, and Nafiseh Shabib. 2014. 2nd International Workshop on NewsRecommendation and Analytics (NRA2014). http://ceur-ws.org/Vol-1181/nra2014_preface.pdf
- [4] Jon Atle Gulla, Bei Yu, Özlem Özgöbek, and Nafiseh Shabib. 2015. 3rd International Workshop on News Recommendation and Analytics (INRA 2015). In *Proceedings of the 9th ACM Conference on Recommender Systems, RecSys 2015, Vienna, Austria, September 16-20, 2015*. 345–346. <https://dl.acm.org/citation.cfm?id=2798721>
- [5] Özlem Özgöbek, Benjamin Kille, Jon Atle Gulla, Martha Larson, and Andreas Lommatzsch. 2018. 6th International Workshop on News Recommendation and Analytics (INRA 2018). In *Proceedings of the CIKM 2018 Workshops co-located with 27th ACM International Conference on Information and Knowledge Management (CIKM 2018), Torino, Italy, October 22, 2018*. <http://ceur-ws.org/Vol-2482/paper10.pdf>
- [6] Özlem Özgöbek, Benjamin Kille, Jon Atle Gulla, and Andreas Lommatzsch. 2019. The 7th International Workshop on News Recommendation and Analytics (INRA 2019). In *Proceedings of the 13th ACM Conference on Recommender Systems (RecSys '19)*. ACM, New York, NY, USA, 558–559. <https://doi.org/10.1145/3298689.3346972>