

Report on RecSys 2016 Workshop on New Trends in Content-Based Recommender Systems

Toine Bogers
Aalborg University Copenhagen
Copenhagen, Denmark
toine@hum.aau.dk

Marijn Koolen
Huygens ING
Amsterdam, the Netherlands
marijn.koolen@uva.nl

Cataldo Musto Pasquale Lops Giovanni Semeraro
University of Bari “Aldo Moro”, Bari, Italy
{cataldo.musto | pasquale.lops | giovanni.semeraro}@uniba.it

Abstract

This article reports on the *CBRecSys 2016* workshop, the third edition of the workshop on *New Trends in Content-based Recommender Systems*, co-located with RecSys 2016 in Boston, MA. Content-based recommendation has been applied successfully in many different domains, but it has not seen the same level of attention as collaborative filtering techniques have. Nevertheless, there are many recommendation domains and applications where content and metadata play a key role, either in addition to or instead of ratings and implicit usage data. The CBRecSys workshop series provides a dedicated venue for work dedicated to all aspects of content-based recommender systems.

I Motivation and Goals

Content-based recommendation has been applied successfully in a variety of domains [8], yet it has not seen the same level of attention as collaborative filtering techniques have. Competitions like the Netflix Prize¹, CAMRA², and the Yahoo! Music KDD Cup 2011 [10] spurred on advances in collaborative filtering and how to utilize ratings and usage data from a community of users. However, there exist many recommendation domains and applications where metadata and item content play a key role, either in addition to or instead of ratings and implicit usage data. For some domains the relationship between content and usage data has seen thorough investigation already (e.g. movies [15]). For many other domains, however, such as books, news, scientific articles, and Web pages we still do not know *if* and *how* these data sources should be combined to provide the

¹<http://www.netflixprize.com/>

²<http://www.dai-labor.de/camra2010/>

best recommendation performance. On top of that, in certain contexts of search, content-based recommendation plays a role as well. Personalised IR is strongly related to recommendation, but recent research in book search also identified search scenarios that combine aspects of retrieval and recommendation [11].

The *CBRecSys* workshop series aims to address this by providing a venue for papers dedicated to all aspects of and new trends in content-based recommender systems. This includes both recommendation in domains where textual content is abundant (e.g., books, news, scientific articles, jobs, educational resources, and Web pages) as well as dedicated comparisons and combinations of content-based techniques with collaborative filtering approaches. After two successful editions organized at RecSys 2014 in Silicon Valley [6, 7] and RecSys 2015 in Vienna [4, 5], the third edition of the workshop, *CBRecSys 2016*, was co-located with RecSys 2016 in Boston, MA.

2 Workshop Focus & Format

Authors were encouraged to submit papers on a variety of topics related to content-based recommendation, including:

- Opinion mining for text/book recommendation
- Semantic recommendation
- Content-based recommendation to alleviate cold-start problems
- Deep learning for content representation
- Serendipity and diversity
- Cross-domain recommendation

In particular, papers submitted to the the workshop focused on the following topics. As in previous years, we received several submissions that examined how semantic information—usually in the form of Linked Open Data resources—can be integrated in recommendation algorithms. Another popular approach in 2016, related to semantic recommendation, was the use of text mining techniques on textual resources, such as reviews, to improve the quality of the recommendations. Deep learning, currently one of the most popular machine learning techniques, also made its way to *CBRecSys*. Finally, one of the accepted submissions examined the problem of content mining for multimedia content. See Section 2.2 for more details about the accepted submissions.

A total of 14 papers were submitted, of which 9 were accepted for an oral presentation (for an acceptance rate of 64%). All papers were reviewed by at least 3 reviewers from the program committee, consisting of international experts in the field. In terms of attendance, *CBRecSys 2016* proved to be a popular workshop filling up the maximum room capacity of 30 with around another 30 attendees on the waiting list.

2.1 Keynotes

The third edition of CBRecSys featured two excellent keynote presentations. In the opening keynote, Barry Smyth (University College Dublin) discussed the role that user-generated reviews could play in a new generation of recommender systems. Barry presented an overview of the work he and his group have done on leveraging opinion mining techniques to extract features and sentiment from unstructured review text, and subsequently how to incorporate this information in recommendation ranking and explanation.

The closing keynote was given by Bamshad Mobasher (DePaul University), who contrasted content and context by focusing on the role of semantic and social knowledge in context-aware recommendation. He presented three different approaches to this: (1) an approach where user profiles are represented as mixtures of the latent topics, allowing for a unified model of users, items, and the metadata associated with contexts; (2) an approach where contextual information is obtained by mining social annotations or other textual features associated with a user's current situation and used in combination with user preference histories to compute a utility function over the set of items; and (3) an approach that emphasizes the role of a domain ontology in the form of a concept hierarchy as an integral part of a user's evolving contextual profile.

2.2 Accepted Papers

This section briefly summarizes the eight accepted papers presented at the workshop. In their work, Lenhart and Herzog [14] addressed the problem of recommending sports news, where users have strong emotional attachments to particular sports, teams and players. They propose a weighted hybrid content-based and collaborative filtering system, with user profiles based on reading behaviour as implicit feedback. The content-based model only takes recent articles into account, to avoid stale news. A collaborative filtering component is added that which allows recommending trending news and older articles that have recently been read. They implemented a recommender dashboard in a large news site and used online evaluation on 1% of users. The content-based approach leads to high accuracy but low diversity. The addition of collaborative filtering increases diversity through recommending trending news and older (yet still relevant) articles.

Lak et al. [13] investigated different text-based content analysis techniques for ranking news. One year of historical data is used to extract text features and combine them with author reputation, article freshness and genre. Both SVM and Random Forests algorithms are used to train classifiers on three types of text features: keyword popularity, TF-IDF and word2vec. Performance drops as older articles are included for all three types of features, but least for word2vec. Overall TF-IDF gives the best performance.

Deldjoo et al. [9] proposed to enrich content-based movie recommendation through automatic extraction of low-level visual features from shots and key frames. The MovieLens dataset is extended with movie trailers from YouTube. The extracted features include dynamic features such as object and camera motion and static features such as color variance and lighting key. The use of visual features improves recommendation performance with respect to baselines that use genre and popularity information.

Bauman et al. [3] presented a new method for recommending not only items but also specific aspects of items, in the domains of restaurant, hotel beauty/spa recommendations. The method extracts aspect-based sentiments from analysis in user reviews and predicts aspect-based sentiment

of users. The system is trained and tested on Yelp data, with evaluation showing improved rating prediction when the system can predict the sentiment of aspects that users mention in their reviews.

Rosati et al. [16] used RDF graph embeddings (rdf2vec) based on short random walks on local sub-graphs in DBpedia and Wikidata to obtain entity sequences. They experiment with different link directions (incoming, outgoing) and types (e.g. RDF type, subject) and test their methods on ratings data for movies and books. Evaluation shows that for the dense movie rating data set, collaborative filtering outperforms content-based approaches. But for book rating data from LibraryThing, which is more sparse, their content-based approach is better. This suggests recommender systems can effectively exploit content-based strategies for domains and user populations where overlapping ratings are rare.

Vagliano et al. [20] proposed a dynamic hybrid algorithm that traverses resource and category links in DBpedia to recommend movies. They conducted a user study using a mobile application and compared the proposed approach against other Linked-Data recommender systems, and found that their approach hits a sweet spot of good accuracy and good novelty, where the recommended list of movies contains a good mix of known (but not watched) and unknown movies.

Ahn et al. [1] proposed a rank aggregation model for quote recommendation exploiting complementary strengths of different algorithms for different aspects: Random Forests for word discrimination, Convolutional Neural Networks for local semantic features and Recurrent Neural Networks to exploit word and sentence ordering. Recommended quotes are chosen from three quote data sets containing half a million quotes, with mentions and contextual data from Twitter, blogs and Gutenberg books. Evaluation shows the three algorithms are highly complementary and together significantly outperform state-of-the-art baselines.

Alzoghbi et al. [2] used learning-to-rank over preference pairs for research paper recommendation. Their approach uses only publicly available metadata and is trained on preference pairs, with papers marked as interesting taken as positive examples and other papers from the same conference as negative examples.

Korpusik et al. [12] used recurrent neural networks to predict which Twitter users are potential buyers of a certain product, to personalise product recommendation. The authors identified regular expression patterns associated with buying or desiring a product and trained recurrent neural networks on tweets both with and without sequence information, with the best performance achieved with a feed-forward network exploiting sequential information.

2.3 Discussion Session about Current and Future Challenges

The workshop was concluded by a discussion session on current trends and future challenges in content-based recommender systems. One of the most important challenges that was brought up by several of the participants was the difficulty of comparing the results of different recommendation algorithms across data sets as reported in research. Currently, it is difficult to know whether newly-proposed algorithms actually outperform other state-of-the-art approaches, because of differences in the data sets used, filtering performed on these data sets, use of evaluation metrics, and other elements of the experimental setup. An additional problem is that details about the experimental setup or evaluation regimen are often missing or lacking in detail, which makes it even harder to compare a new algorithm to other state-of-the-art approaches.

Naturally, this issue is important not just for content-based recommendation, but for recom-

mender systems research in general, and was signaled already in 2014 by Said & Bellogín [17]. They also proposed RIVAL, a toolkit for transparent and objective benchmarking of recommender systems software. This would enable making robust and comparable assessments of their recommendation quality [18]. Another possible (and partial) solution to this problem could be to set up an recommender systems evaluation infrastructure that stores the results of recommendation runs, as well as the data sets and details about the experimental setup(s). Similar to the PROMISE evaluation infrastructure³ for information retrieval research, this would greatly streamline the process of directly comparing new algorithms to older state-of-the-art algorithms.

Another challenge brought up was the need for data sets that enable the study of feature preferences by users: which aspects, properties, or features of a product do they prefer as opposed to only studying item-level preferences. More research is also needed on how to elicit and collect such feature preference information. A better understanding of feature preferences could lead to higher-quality recommendations. It would also enable the development of approaches to multi-criteria feature selection that take other dimensions besides accuracy into account, such as the potential ‘explainability’ of recommendation.

Finally, the topic of explaining content-based recommendation evoked additional discussion, especially in light of the EU’s General Data Protection Regulation (GDPR) law⁴ which will come into effect in 2018. Specific provisions of the law may be interpreted in such a way that EU citizens might soon have a way to demand explanations of the decisions that algorithms make about them. In addition to a possible legal requirement, generating explanations for recommendations can also improve the transparency, scrutability, trust, effectiveness, persuasiveness, efficiency, and satisfaction of a recommender system [19]. Participants quickly agreed that the relevant importance of this depends on the application domain, which was followed by a lively discussion about the relative importance of explanations for the different domains.

3 Future Workshop

The continued popularity of the CBRecSys workshops in terms of both submissions and participation clearly shows a consistent interest in research on content-based recommendation. We are therefore planning a fourth and final edition of the CBRecSys workshop at RecSys 2017 in Como, Italy.

4 Acknowledgments

We would like to thank ACM and RecSys for hosting this workshop, the RecSys workshop chairs Elizabeth Daly and Dietmar Jannach. We would also like to thank the program committee: Jon Atle Gulla, Shlomo Berkovsky, Ludovico Boratto, Robin Burke, Iván Cantador, Federica Cena, Paolo Cremonesi, Marco de Gemmis, Ernesto William De Luca, Tommaso Di Noia, Peter Dolog, Fabio Gaspiretti, Christina Gena, Frank Hopfgartner, Juan F. Huete, Jaap Kamps, Silvia Likavec, Babak Loni, Fedelucio Narducci, Casper Petersen, Shaghayegh Sahebi, Alan Said, Marko Tkalčič, and

³See <http://promise-noe.eu> for more information.

⁴Available at <http://ec.europa.eu/justice/data-protection/>.

Bei Yu. Final thanks are due to the paper authors, our invited speakers Barry Smyth and Bamshad Mobasher, and the participants for an interesting and lively workshop.

The workshop material (list of accepted papers, invited talk, and the workshop schedule) can be found on the CBRecSys 2016 workshop website at <http://cbrecsys2016.aau.dk/>. The proceedings were published as a CEUR Workshop Proceedings volume, available at <http://ceur-ws.org/Vol-1673/>.

References

- [1] Y. Ahn, H. Lee, H. Jeon, S. Ha, and S.-G. Lee. Quote Recommendation for Dialogs and Writings. In *CBRecSys '16: Proceedings of the 3rd Workshop on New Trends on Content-Based Recommender Systems*, pages 39–42, 2016.
- [2] A. Alzoghbi, V. A. A. Ayala, P. M. Fischer, and G. Lausen. Learning-to-Rank in Research Paper CBF recommendation: Leveraging Irrelevant Papers. In *CBRecSys '16: Proceedings of the 3rd Workshop on New Trends on Content-Based Recommender Systems*, pages 43–46, 2016.
- [3] K. Bauman, B. Liu, and A. Tuzhilin. Recommending Items with Conditions Enhancing User Experiences Based on Sentiment Analysis of Reviews. In *CBRecSys '16: Proceedings of the 3rd Workshop on New Trends on Content-Based Recommender Systems*, pages 19–22, 2016.
- [4] T. Bogers and M. Koolen. Report on RecSys 2015: Workshop on New Trends in Content-Based Recommender Systems. *SIGIR Forum*, 49(2):141–146, 2015.
- [5] T. Bogers and M. Koolen. Second Workshop on New Trends in Content-based Recommender Systems (CBRecSys 2015). In *RecSys '15: Proceedings of the Ninth ACM Conference on Recommender Systems*, pages 339–340, 2015.
- [6] T. Bogers, M. Koolen, and I. Cantador. Workshop on New Trends in Content-based Recommender Systems (CBRecSys 2014). In *RecSys '14: Proceedings of the Eighth ACM Conference on Recommender Systems*, pages 379–380, 2014.
- [7] T. Bogers, M. Koolen, and I. Cantador. Report on RecSys 2014: Workshop on New Trends in Content-Based Recommender Systems. *SIGIR Forum*, 49(1):20–26, 2015.
- [8] M. de Gemmis, P. Lops, C. Musto, F. Narducci, and G. Semeraro. Semantics-Aware Content-Based Recommender Systems. In F. Ricci, L. Rokach, and B. Shapira, editors, *Recommender Systems Handbook*, pages 119–159. Springer, 2015.
- [9] Y. Deldjoo, M. Elahi, and P. Cremonesi. Using Visual Features and Latent Factors for Movie Recommendation. In *CBRecSys '16: Proceedings of the 3rd Workshop on New Trends on Content-Based Recommender Systems*, pages 15–18, 2016.
- [10] G. Dror, N. Koenigstein, Y. Koren, and M. Weimer. The Yahoo! Music Dataset and KDD-Cup '11. In *JMLR Workshop and Conference Proceedings*, volume 18 of *Proceedings of KDD Cup 2011*, pages 3–18, 2012.

-
- [11] M. Koolen, T. Bogers, A. van den Bosch, and J. Kamps. Looking for books in social media: An analysis of complex search requests. In *Advances in Information Retrieval: 37th European Conference on IR Research (ECIR 2015)*, LNCS. Springer, 2015.
- [12] M. Korpusik, S. Sakaki, F. Chen, and Y.-Y. Chen. Recurrent Neural Networks for Customer Purchase Prediction on Twitter. In *CBRecSys '16: Proceedings of the 3rd Workshop on New Trends on Content-Based Recommender Systems*, pages 47–50, 2016.
- [13] P. Lak, C. Babaoglu, A. B. Bener, and P. Pralat. News Article Position Recommendation Based on the Analysis of Article's Content – Time Matters. In *CBRecSys '16: Proceedings of the 3rd Workshop on New Trends on Content-Based Recommender Systems*, pages 11–14, 2016.
- [14] P. Lenhart and D. Herzog. Combining Content-based and Collaborative Filtering for Personalized Sports News Recommendations. In *CBRecSys '16: Proceedings of the 3rd Workshop on New Trends on Content-Based Recommender Systems*, pages 3–10, 2016.
- [15] I. Pilászy and D. Tikk. Recommending New Movies: Even a Few Ratings Are More Valuable Than Metadata. In *RecSys '09: Proceedings of the Third ACM Conference on Recommender Systems*, pages 93–100. ACM, 2009.
- [16] J. Rosati, P. Ristoski, T. D. Noia, R. D. Leone, and H. Paulheim. RDF Graph Embeddings for Content-based Recommender Systems. In *CBRecSys '16: Proceedings of the 3rd Workshop on New Trends on Content-Based Recommender Systems*, pages 23–30, 2016.
- [17] A. Said and A. Bellogín. Comparative Recommender System Evaluation: Benchmarking Recommendation Frameworks. In *RecSys '14: Proceedings of the 8th ACM Conference on Recommender Systems*, pages 129–136, 2014.
- [18] A. Said and A. Bellogín. RIVAL: A Toolkit to Foster Reproducibility in Recommender System Evaluation. In *RecSys '14: Proceedings of the 8th ACM Conference on Recommender Systems*, pages 371–372, 2014.
- [19] N. Tintarev and J. Masthoff. Designing and Evaluating Explanations for Recommender Systems. In F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, editors, *Recommender Systems Handbook*, pages 479–510. Springer, 2011.
- [20] I. Vagliano, C. Figueroa, O. R. Rocha, M. Torchiano, C. Faron-Zucker, and M. Morisio. ReDYAl: A Dynamic Recommendation Algorithm based on Linked Data. In *CBRecSys '16: Proceedings of the 3rd Workshop on New Trends on Content-Based Recommender Systems*, pages 31–38, 2016.