

Information Retrieval and User-Centric Recommender System Evaluation

Alan Said[†], Alejandro Bellogín[†], Arjen De Vries[†], Benjamin Kille^{*}

CWI, The Netherlands[†], TU Berlin, Germany^{*}
{alan, alejandro.bellogin, arjen.de.vries}@cwi.nl[†], kille@dai-lab.de^{*}

Abstract. Traditional recommender system evaluation focuses on raising the accuracy, or lowering the rating prediction error of the recommendation algorithm. Recently, however, discrepancies between commonly used metrics (e.g. precision, recall, root-mean-square error) and the experienced quality from the users' have been brought to light. This project aims to address these discrepancies by attempting to develop novel means of recommender systems evaluation which encompasses qualities identified through traditional evaluation metrics and user-centric factors, e.g. diversity, serendipity, novelty, etc., as well as bringing further insights in the topic by analyzing and translating the problem of evaluation from an Information Retrieval perspective.

Keywords: Recommender Systems, Evaluation, Information Retrieval

1 Introduction

The project is framed in the Recommender Systems (RS) field. The aim of RSs is to assist users in finding their way through huge databases and catalogues, by filtering and suggesting relevant items, taking into account the users' preferences (i.e., tastes, priorities, etc.).

2 Novel Methods for Recommender System Evaluation

Over the last two decades, a vast amount of research in RS has lead to great progress in terms of prediction accuracy [1]. Today, the majority of the work on RS is based on top-n recommendation or rating prediction; the former requires bi/unary interaction data between users and items, whereas the latter requires a dataset with ratings. This type of evaluation is also common in information retrieval (IR) systems [3].

2.1 User-centric Evaluation

Both top-n and rating prediction-based evaluation build on several assumptions which could potentially have negative effects on recommendation algorithms tuned solely on these evaluation metrics [4, 1, 5, 6]. These are:

- there is an absolute ground truth which the RS should attempt to identify,
- users are primarily interested in the items which have received the highest ratings,
- higher top-n accuracy or lower rating error levels translate to a higher perceived usefulness from the users.

Recent work has, however, shown that these assumptions are not always true in the RS context, e.g. [7–9]. In specific cases, the assumptions are detrimental to the users’ perceived quality.

The main focus of this sub-project is to analyze the discrepancies between *offline* and *online* evaluation, and to gain insights into the subjective qualities of various recommendation algorithms and their specific qualities. With this in mind, there are several goals we will strive to achieve:

- Analyze the correlation of IR-related evaluation metrics and user-centric concepts such as diversity, novelty and other non-quantifiable RS aspects.
- Identify whether there exists a correlation between the properties of items that are regarded as false recommendations in traditional (offline) IR and RS evaluation settings and true (high quality) recommendations in user-centric (online) evaluation, taking factors such as serendipity and usefulness into consideration.
- Evaluate whether metrics from research areas outside of IR and RS (e.g. signal processing, economics) can estimate sought for qualities better than the currently used RS and IR metrics.
- Improve the understanding of which evaluation metrics should be applied to RSs in different contexts, using different algorithms, data sets, and other system-specific features.

For these purposes, we will use a variety of data sets from the multimedia domain, ranging from publicly available, e.g. Movielens, Last.fm, as well as proprietary from other related services, e.g. Filmtipset, Moviepilot, etc. The variety of data will ensure the general applicability of the research results. Additional data will be collected through user studies and surveys.

2.2 Information Retrieval-based Evaluation

RS are usually considered as a special case of IR systems, specifically, one where no query is given and the information to be retrieved has to be inferred from previous user experiences. For this reason, some of the models and theories developed in IR have already been translated to RS, such as the Vector Space Model and the Probability Ranking Principle [10].

There are, however, several gaps in the understanding of RS as personalized IR systems, such as the need of formal methods to introduce implicit and contextual feedback in the recommendations and the lack of a proper evaluation framework.

A strong link between IR and RS has already been shown in our previous research (see [11, 12]), where we adapted to RSs different techniques proposed

in IR to predict the performance of a system. A natural next step is to explore how the evaluation in RS may benefit from extending this analogy between IR and recommendation, and applying more retrieval methodologies to recommendation.

More specifically, in this sub-project, we aim to exploit IR concepts, algorithms, and methodologies for recommendation. RS has a well known tradition in integrating contextual information which could be, in turn, transferred from RS to IR and investigate how such methods could be integrated in contextual IR and whether some benefits could be found by creating such links.

With this goal in mind, there are several objectives we aim to achieve:

- Analyze how evaluation in recommendation should be performed to obtain a general methodology that would result in interpretable and comparable results for the community, mainly by adapting and integrating models and metrics from IR.
- Identify if there is any correlation between the metrics typically used in offline experiments (e.g., precision) and those pervasive in real applications, more useful from a business point of view (such as the click through and conversion rates).
- Bring the models and theories used in context-aware recommendation to contextual IR and vice versa, in such a way that further interactions between these two areas could be found.
- Exploit implicit information for RS in novel ways based on current research from the use of search logs and other implicit sources of information in IR.

3 Current Results

At the moment of writing, we have obtained positive results regarding some of the aspects of this research project. Specifically, a workshop on reproducibility and replication in recommender evaluation has been accepted at the 2013 ACM RecSys conference. We are also working on revisiting several classic recommendation techniques in order to evaluate them under a common framework based on an IR-inspired recommendation method previously proposed in [14]. A novel evaluation protocol has also been researched specifically for RS. Additionally, we are researching different features that could help in the understanding of why some similarity functions perform better when applied within a recommendation strategy. Finally, we have analyzed different sources of implicit and explicit popularity scores in order to exploit them in a recommendation context; this study was recently accepted for publication as [15].

4 Acknowledgements

This work was carried out during the tenure of an ERCIM “Alain Bensoussan” Fellowship Programme. The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no.246016.

References

1. Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T.: Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst.* **22**(1) (January 2004) 5–53
2. Amatriain, X., Basilico, J.: Netflix recommendations: Beyond the 5 stars (part 1) – the netflix tech blog. <http://techblog.netflix.com/2012/04/netflix-recommendations-beyond-5-stars.html> (retrieved May 12, 2012) (April 2012)
3. Baeza-Yates, R.A., Ribeiro-Neto, B.: *Modern Information Retrieval*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA (1999)
4. Hill, W., Stead, L., Rosenstein, M., Furnas, G.: Recommending and evaluating choices in a virtual community of use. In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. CHI '95, New York, NY, USA, ACM Press/Addison-Wesley Publishing Co. (1995) 194–201
5. Amatriain, X., Pujol, J.M., Oliver, N.: I like it... i like it not: Evaluating user ratings noise in recommender systems. In: *Proceedings of the 17th International Conference on User Modeling, Adaptation, and Personalization: formerly UM and AH. UMAP '09*, Berlin, Heidelberg, Springer-Verlag (2009) 247–258
6. Cremonesi, P., Garzotto, F., Negro, S., Papadopoulos, A., Turrin, R.: Comparative evaluation of recommender system quality. In: *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems*. CHI EA '11, New York, NY, USA, ACM (2011) 1927–1932
7. Said, A., Fields, B., Jain, B.J.: User-centric evaluation of a k-furthest neighbor collaborative filtering recommender algorithm. In: *Proceedings of the ACM 2013 conference on Computer Supported Cooperative Work*, New York, NY, USA, ACM (2013)
8. Said, A., Jain, B., Narr, S., Plumbaum, T.: Users and noise: The magic barrier of recommender systems. In *Masthoff, J., Mobasher, B., Desmarais, M., Nkambou, R., eds.: User Modeling, Adaptation, and Personalization*. Volume 7379 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg (2012) 237–248 10.1007/978-3-642-31454-4_20.
9. Said, A., Jain, B.J., Narr, S., Plumbaum, T., Albayrak, S., Scheel, C.: Estimating the magic barrier of recommender systems: a user study. In: *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*. SIGIR '12, New York, NY, USA, ACM (2012) 1061–1062
10. Wang, J., Robertson, S., de Vries, A., Reinders, M.: Probabilistic relevance ranking for collaborative filtering. *Information Retrieval* **11**(6) (December 2008) 477–497
11. Bellogín, A.: *Recommender System Performance Evaluation and Prediction: an Information Retrieval Perspective*. PhD thesis, Universidad Autónoma de Madrid, Spain (November 2012)
12. Bellogín, A.: Performance Prediction in Recommender Systems. In *Konstan, J., Conejo, R., Marzo, J., Oliver, N., eds.: User Modeling, Adaption and Personalization*. Volume 6787 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, Berlin, Heidelberg (2011) 401–404
13. Said, A.: *Researching Novel Methods for Recommender System Evaluation*. In: *UMAP (extended proceedings)*, Rome (2013)
14. Bellogín, A., Wang, J., Castells, P.: Bridging memory-based collaborative filtering and text retrieval. (2012) 1–28
15. Bellogín, A., de Vries, A., He, J.: Artist popularity: do web and social music services agree? In: *Int. Conf. on Weblogs and Social Media (ICWSM)*, Boston (2013)