

# Predicting Recommender System Performance from an Information Retrieval Perspective

Alejandro Bellogín, Pablo Castells, Iván Cantador  
Universidad Autónoma de Madrid, Departamento de Ingeniería Informática

## ABSTRACT

Personalised recommender systems aim to help users retrieve relevant information from large collections, by automatically finding and suggesting products of likely interest based on observed evidence of the users' preferences. These preferences are difficult to guess, and therefore recommender systems have a considerable variance in their success ratio in estimating the user's tastes and interests. In such a scenario, self-predicting the chances that a recommendation may be accurate before actually submitting it to a user becomes an interesting capability from many perspectives. Performance prediction has been studied in the context of Information Retrieval [3], but there is little if any prior research of this problem in the recommendation domain.

In our work we investigate the definition and formalisation of performance prediction methods for recommender systems [1]. Specifically, we propose and analyse an array of methods based on adaptations of performance predictors from Information Retrieval - mainly the query clarity predictor, which captures the query ambiguity with respect to a given document collection [2]. We define several language models according to various probability spaces (items, ratings, and combinations thereof) to capture different aspects of the users and items involved in the recommendation task. Through a series of empirical experiments on several datasets, we have found strong correlations between the proposed predictors and the performance measured, confirming the predictive power of our methods. Furthermore, we have also found a correspondence between such predictive power and performance enhancements in two applications for recommendation: dynamic recommender ensembles and neighbour selection and weighting.

## BODY

*We develop new recommendation effectiveness prediction methods based on query clarity. Recommendation performance is, in fact, predictable.*

## REFERENCES

- [1] A. Bellogín. *Recommender System Performance Evaluation and Prediction: An Information Retrieval Perspective*. PhD thesis, Univ. Autónoma de Madrid, Spain, Nov. 2012.
- [2] S. Cronen-Townsend, Y. Zhou, and W. B. Croft. Predicting query performance. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '02, pages 299–306, New York, NY, USA, 2002. ACM.
- [3] C. Hauff. *Predicting the Effectiveness of Queries and Retrieval Systems*. PhD thesis, Univ. of Twente, Enschede, Jan. 2010.

*Volume 2 of Tiny Transactions on Computer Science*

This content is released under the Creative Commons Attribution-NonCommercial ShareAlike License. Permission to make digital or hard copies of all or part of this work is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. CC BY-NC-SA 3.0: <http://creativecommons.org/licenses/by-nc-sa/3.0/>.