Evaluating the Relative Performance of Collaborative Filtering Recommender Systems

Humberto Jesús Corona Pampín[*], Houssem Jerbi[**], and Michael P. O'Mahony[**]

[*]Fashion Insights Centre, Zalando Ireland
[**]Insight Centre for Data Analytics, University College Dublin

Summary

Past work on the evaluation of recommender systems indicates that collaborative filtering algorithms are accurate and suitable for the top-N recommendation task. Further, the importance of performance beyond accuracy has been recognised in the literature. Here, we present an evaluation framework based on a set of accuracy and beyond accuracy metrics, including a novel metric that captures the uniqueness of a recommendation list. We perform an in-depth evaluation of three well-known collaborative filtering algorithms using three datasets. The results show that the user-based and item-based collaborative filtering algorithms have a high inverse correlation between popularity and diversity and recommend a common set of items at large neighbourhood sizes. The study also finds that the matrix factorisation approach leads to more accurate and diverse recommendations, while being less biased toward popularity [1][1].

References

1. Pampín, H.J.C., Jerbi, H., O'Mahony, M.P.: Evaluating the relative performance of collaborative filtering recommender systems. Journal of Universal Computer Science 21(13), 1849–1868 (2015)