A Framework for Evaluation of Information Filtering Techniques in an Adaptive Recommender System

John O'Donovan and John Dunnion

Department of Computer Science,
University College Dublin,
Ireland.
{john.odonovan,john.dunnion}@ucd.ie

Abstract. This paper proposes that there is a substantial relative difference in the performance of information-filtering algorithms as they are applied to different datasets, and that these performance differences can be leveraged to form the basis of an Adaptive Information Filtering System. We classify five different datasets based on metrics such as sparsity, user-item ratio etc, and develop a regression function over these metrics in order to predict suitability of a particular recommendation algorithm to a new dataset, using only the aforementioned metrics. Our results show that the predicted best algorithm does perform better for the new dataset.

1 Introduction

Collaborative Filtering (CF) is a broad term for the process of recommending items to users based on similarities in user taste [10][3]. An increasing number of online stores provide collaborative recommender systems on their sites, e.g. E-Bay, Amazon.com etc. CF techniques tend to have the following advantages over others: They do not require items to be machine-analysable (as explained in [6]), they can arrive at serendipitous recommendations [12]. They also require little knowledge-engineering overhead. CF techniques are also subject to two serious restrictions. Sparsity Restriction: In any given case, it is unlikely that two users have co-rated many of the items in the system. Accurate similarity measurements depend on rich user profiles with high overlap. Latency Restriction: This affects new or unique items. These items will not be recommended by a system until they are included in a sufficient number of user profiles, as outlined in [12]. Similarity can be computed for CF by several well-known techniques, auch as Cosine Similarity, Spearman's or Pearson's Correlation[8]. For all of our similarity calculations we employ Pearsons, as it is the most widely used and allows for better comparison with other systems.

Different implementations of collaborative filtering will be affected to varying degrees by the problems mentioned above. Their performance will change based on the dataset that they operate on, and the information they harness to compile a similarity model. For example, for a situation where the set of items to be recommended is relatively small and static, and there are a large number of users, it would be advisable to employ an item-based approach[11][5] to CF,