
User Modeling and User-Adapted Interaction **12:** 331–370, 2002.

Hybrid Recommender Systems: Survey and Experiments[1]

ROBIN BURKE
Department of Information Systems and Decision Sciences, California State University, Fullerton, CA 92834, USA

(Received 23 January 2000; accepted in revised form 24 September 2001)

Abstract. Recommender systems represent user preferences for the purpose of suggesting items to purchase or examine. They have become fundamental applications in electronic commerce and information access, providing suggestions that effectively prune large information spaces so that users are directed toward those items that best meet their needs and preferences. A variety of techniques have been proposed for performing recommendation, including content-based, collaborative, knowledge-based and other techniques. To improve performance, these methods have sometimes been combined in hybrid recommenders. This paper surveys the landscape of actual and possible hybrid recommenders, and introduces a novel hybrid, EntreeC, a system that combines knowledge-based recommendation and collaborative filtering to recommend restaurants. Further, we show that semantic ratings obtained from the knowledge-based part of the system enhance the effectiveness of collaborative filtering.

Key words: case-based reasoning, collaborative filtering, electronic commerce, recommender systems

1. Introduction

Recommender systems were originally defined as ones in which 'people provide recommendations as inputs, which the system then aggregates and directs to appropriate recipients' (Resnick & Varian, 1997). The term now has a broader connotation, describing any system that produces individualized recommendations as output or has the effect of guiding the user in a personalized way to interesting or useful objects in a large space of possible options. Such systems have an obvious appeal in an environment where the amount of on-line information vastly outstrips any individual's capability to survey it. Recommender systems are now an integral part of some e-commerce sites such as Amazon.com and CDNow (Schafer, Konstan & Riedl, 1999).

It is the criteria of 'individualized' and 'interesting and useful' that separate the recommender system from information retrieval systems or search engines. The semantics of a search engine are 'matching': the system is supposed to return all those items that match the query ranked by degree of match. Techniques such

[1]The managing editor for this paper was Ingrid Zukerman

as relevance feedback enable a search engine to refine its representation of the user's query, and represent a simple form of recommendation. The next-generation search engine Google[1] blurs this distinction, incorporating 'authoritativeness' criteria into its ranking (defined recursively as the sum of the authoritativeness of pages linking to a given page) in order to return more useful results (Brin & Page, 1998).

One common thread in recommender systems research is the need to combine recommendation techniques to achieve peak performance. All of the known recommendation techniques have strengths and weaknesses, and many researchers have chosen to combine techniques in different ways. This article surveys the different recommendation techniques being researched — analyzing them in terms of the data that supports the recommendations and the algorithms that operate on that data — and examines the range of hybridization techniques that have been proposed. This analysis points to a number of possible hybrids that have yet to be explored. Finally, we discuss how adding a hybrid with collaborative filtering improved the performance of our knowledge-based recommender system Entree. In addition, we show that semantic ratings made available by the knowledge-based portion of the system provide an additional boost to the hybrid's performance.

1.1. RECOMMENDATION TECHNIQUES

Recommendation techniques have a number of possible classifications (Resnick & Varian, 1997; Schafer, Konstan & Riedl, 1999; Terveen & Hill, 2001). Of interest in this discussion is not the type of interface or the properties of the user's interaction with the recommender, but rather the sources of data on which recommendation is based and the use to which that data is put. Specifically, recommender systems have (i) background data, the information that the system has before the recommendation process begins, (ii) input data, the information that user must communicate to the system in order to generate a recommendation, and (iii) an algorithm that combines background and input data to arrive at its suggestions. On this basis, we can distinguish five different recommendation techniques as shown in Table I. Assume that **I** is the set of items over which recommendations might be made, **U** is the set of users whose preferences are known, **u** is the user for whom recommendations need to be generated, and **i** is some item for which we would like to predict **u**'s preference.

Collaborative recommendation is probably the most familiar, most widely implemented and most mature of the technologies. Collaborative recommender systems aggregate ratings or recommendations of objects, recognize commonalities between users on the basis of their ratings, and generate new recommendations based on inter-user comparisons. A typical user profile in a collaborative system consists of a vector of items and their ratings, continuously augmented as the user interacts with the system over time. Some systems used time-based discounting of ratings to account for drift in user interests (Billsus & Pazzani, 2000; Schwab et al., 2001). In some cases, ratings may be binary (like/dislike) or real-valued indicating degree

[1]URL http://www.google.com