Distributed Recommender Systems for Internet Commerce

Badrul M. Sarwar
University of Minnesota, USA

Joseph A. Konstan
University of Minnesota, USA

John T. Riedl
University of Minnesota, USA

INTRODUCTION

Recommender systems (RSs) present an alternative information-evaluation approach based on the judgements of human beings (Resnick & Varian, 1997). It attempts to automate the word-of-mouth recommendations that we regularly receive from family, friends, and colleagues. In essence, it allows everyone to serve as a critic. This inclusiveness circumvents the scalability problems of individual critics—with millions of readers it becomes possible to review millions of books. At the same time it raises the question of how to reconcile the many and varied opinions of a large community of ordinary people. Recommender systems address this question through the use of different algorithms: *nearest-neighbor algorithms* (Resnick, Iacovou, Suchak, Bergstrom, & Riedl, 1994; Shardanand et al., 1994), *item-based algorithms* (Sarwar, Karypis, Konstan, & Riedl, 2001), *clustering algorithms* (Ungar & Foster, 1998), and *probabilistic* and *rule-based learning algorithms* (Breese, Heckerman, & Kadie, 1998), to name but a few. The nearest-neighbor-algorithm-based recommender systems, which are often referred to as *collaborative filtering* (CF) *systems* in research literature (Maltz & Ehrlich, 1995), are the most widely used recommender systems in practice. A typical CF-based recommender system maintains a database containing the *ratings* that each customer has given to each product that customer has evaluated. For each customer in the system, the recommendation engine computes a *neighborhood* of other customers with similar opinions. To evaluate other products for this customer, the system forms a normalized and weighted average of the opinions of the customer's neighbors.

The emergence of the Internet and its far-reaching deployment is changing the way commerce is done. Economists and commerce experts are now suggesting companies to shift from the old world of mass production characterized by "...standardized products, homogenous markets, and long product life and development cycles..." to the new world where "...variety and customization supplant standardized products" (Pine, 1993). In his famous book *Mass Customization*, Joe Pine also suggests that building one product is simply not adequate anymore. Companies need to be able to develop multiple products that meet the multiple needs of multiple customers. The movement toward Internet commerce has allowed companies to provide customers with more options. However, in expanding to this new level of customization, businesses increase the amount of information that customers must process before they are able to select which items meet their needs. Traditional data-analysis techniques are often not sufficient to process this huge amount of data in real time as needed by the Internet sites. Recommender systems, by providing a "personalized" interface to each customer, can potentially automate personalization on the Internet. Personalization to this extent is one way to realize Pine's ideas of the "new world order" of Internet commerce.

As discussed by Schafer, Konstan, and Riedl (1999), recommender systems can help Internet commerce sites boost their business in several ways. First, by providing personalized recommendations on various products, they help convert browsers into buyers. Visitors to a Web site often look over the site without ever purchasing anything. Recommender systems can help customers find products they wish to purchase and can potentially increase sales. Second, recommender systems improve cross-sell by suggesting additional products for the customer to purchase. If the recommendations are good, the average order size should increase. For instance, a site might recommend additional products in the checkout process based on those products already in the shopping cart. Third, recommender systems help capture customer loyalty. Internet commerce is getting competitive day by day. Freed from large capital investment and recurring costs for physical storefronts, an unprecedented number

of businesses are using the Internet to market and sell goods, potentially creating a vicious price war. As a consequence, gaining customer loyalty becomes an essential strategy for businesses to survive on the Internet nowadays (Reichheld, 1993; Reichheld & Sesser, 1990). Recommender systems improve loyalty by creating a value-added relationship between the site and the customer. Sites invest in learning about their users, use recommender systems to operationalize that learning, and present custom interfaces that match customer needs. Customers repay these sites by returning to the ones that best match their needs.

RECOMMENDATION INTERFACES

There is more than one way to display recommendations to a customer. The method selected may well depend on how the Internet commerce site wants the customer to use the recommendation. In the following we will examine several recommendation interfaces and how each assists the site in making money. While some of these methods have their roots in traditional commerce, each of them draws upon the strengths of ubiquitous Internet to provide more powerful recommendations. We present these interfaces in Table 1 (Sarwar, Konstan, & Riedl, 2001; Schafer et al., 1999).

DISTRIBUTED RECOMMENDER SYSTEMS

In the past, participating in commerce meant that the consumer had to travel to the location of the store from which he or she wanted to purchase a product. Today, participating in commerce may be as easy as moving a mouse and typing a few keystrokes. In the future, participating in commerce will become even easier. With the introduction of new wireless devices that enable commerce on the palmtop, consumers will be able to shop from wherever they happen to be. The result of technical improvements like wireless Web browsers is that consumers will come to expect the same shopping experience when travelling as they currently receive when directly connected to the Internet. For instance, a customer who uses a restaurant recommendation service in his or her hometown will expect to be able to use the same restaurant recommendation service while traveling. In fact, the recommendations will be even more valuable on the road since the customer will know less about what is available.

Creating good recommendations for a traveling consumer is challenging, though, especially for the leading recommendation technology, collaborative filtering. Many Internet commerce sites, including several of the largest ones, are now using CF recommender systems as part of their personalization effort. As the largest Internet commerce sites attempt to use these systems, however, they are discovering the drawbacks of today's centralized systems. While a centralized architecture may be useful for smaller applications, there are several key drawbacks to centralization as shown in Table 2. Together, these issues point toward distributed architectures for collaborative filtering (Sarwar, Konstan et al., 2001).

There are four fundamemtal components that are to be considered for designing a framework for distributed recommender system:

1) *Products*: Are the products of local or global interest?

Table 1. Recommendation interfaces for Internet commerce

Browsing: Recommended browsing helps the users narrow down their choices and makes them feel more confident about their decision buying decision by providing organized access to the recommendations.
Similar Item: Recommender systems display items based on the item in which a customer has shown interest. In so doing, sites increase customers' exposure to their product line and are ideally able to sell more items per order.
E-mail: Recommendations can also be delivered directly to customers through e-mail. This extension of the traditional direct-mail campaign is also expected to generate more sales.
Text Comments: Sites can provide customers with recommendations based directly on the text comments of other customers.
Average Rating: Recommender systems show the average ratings for particular items.
Top *N*: Once each site has learned details about a customer's likes and dislikes, each is able to provide the customer with a personalized list of the top-*N* unrated items for that customer. It helps sites convert browsers into buyers as well as helps customers in making a decision about a product that they originally held in doubt.