

Recommender Systems for E-Commerce: Challenges and Opportunities

From: AAAI Technical Report WS-99-01. Compilation copyright © 1999, AAAI (www.aaai.org). All rights reserved.

Robert Driskill and John Riedl

Net Perceptions
7901 Flying Cloud Drive
Eden Prairie, Minnesota 55344
{rdriskill,riedl}@netperceptions.com

Abstract

Recommender systems are an AI technology that has become an essential part of business for many E-commerce sites. They serve many types of E-commerce applications, from direct product recommendation for an individual to helping someone find a gift for a third party. In this paper, we provide a brief overview of how recommender systems are being used in E-commerce today, and analyze four key challenges for recommender systems in the future: hybrid data, predictable recommendations, scalability, and incorporation of content. If recommender systems are able to surmount these challenges, they have the potential to become an essential component of doing business in E-commerce.

Introduction

Recommender systems were developed to help users sort through large information spaces to find items of most value to them. Their initial applications were to information spaces such as document spaces in corporate workgroups (Goldberg et al. 1992), Usenet News (Resnick et al. 1994), video recommendations (Shardanand et al. 1995, Hill et al. 1995), and Web pages (Terveen et al. 1997). These initial applications demonstrated the potential to develop systems that use the opinions of a large group of users to help other users identify the items in which they will be most interested.

Taxonomy

Recommender systems have also been applied to E-commerce. Three applications have proven most successful: direct product recommendations, gift centers, and cross-sell recommendations. The most direct applications of a recommender system to E-commerce is to make direct product recommendations to help individual customers find products they would like to purchase. For instance, Amazon.com (www.amazon.com) has a BookMatcher recommender system in which they ask customers to rate a number of books. Amazon.com then matches those ratings with other customers to find other books each customer is likely to enjoy. Amazon.com reports that recommender systems help sell more books,

Copyright © 1999, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

and help ensure customer loyalty, though they have not published more specific results.

A second application of recommender systems to E-commerce is to gift centers. This application is intended for a customer who visits a site with the intent of purchasing a product as a gift for someone else. The challenge in this application is to learn enough about the gift recipient to effectively make a recommendation, without requiring too much effort from the customer. One such system is used in the Cdnw gift center (www.cdnw.com). Here, the customer is asked to enter three musical artists the recipient likes. CDnow then searches for albums by other artists that are frequently liked by other customers who like those first three artists. The power of this application is that it is anonymous with respect to the recipient, yet still manages to deliver helpful recommendations to the customer.

The final application of recommender systems to E-commerce that we will discuss is cross-sell. Cross-sell in E-commerce involves recommending items to a customer based on other items that that user has already selected for purchase in the current visit. For instance, many E-commerce sites have suggestions for other products at checkout. Most sites are making these suggestions based on simple cross-sell lists of products that are judged by analysts to be "compatible" with products in your shopping basket. Recommender systems could be used to personalize these cross-sell lists, making the cross-sell more effective.

Recommender systems work by analyzing databases of customer interaction with Web sites. This data is often in the form of purchase information (i.e., what items has this customer purchased). Other forms of data that are often available include the contents of shopping baskets and explicit rating data (i.e., data gathered via some form of questionnaire, generally indicating the degree of like or dislike for certain products).

One of the most successful recommender technologies to date is automatic collaborative filtering. Collaborative filtering systems work by measuring distances between people in "taste space", and predicting interest in untried items based on a weighted sum of nearby users' impressions of the untried items. Collaborative filtering systems are powerful tools for recommending individual items in a large space of items in which individual taste is important. For instance, several music stores are using

collaborative filtering to recommend CDs to their customers.

Recommender systems can also be developed using neural network technology (Lang 1995), Bayesian belief networks (Breese et al. 1998), rule induction systems like Ripper (Cohen et al. 1995) and many other technologies. The defining properties for recommender systems in E-Commerce are:

- (1) They operate on databases of interactions between customers and products.
- (2) They produce recommendations of which products a given customer will like.
- (3) They learn over time to make better recommendations based on continued interaction with the customer.
- (4) They operate interactively, providing recommendations while the customer is visiting the site, and modifying the recommendations in real-time based on the interactions of the customer with the site.

Some recommender systems use substantial offline learning algorithms, which are permitted under this definition as long as they can also perform some learning interactively (Breese et al. 1998).

Challenges

The e-commerce environment provides a number of interesting challenges to the recommendation system developer.

One major challenge is in the area of so-called hybrid systems. The currently dominant systems generate recommendations utilizing only one type of input data about customer preferences for products (for example, explicit ratings data or purchase data). The goal of a hybrid system is to take all available preference data simultaneously, and use it in an intelligent way to provide recommendations.

This simple sounding problem actually contains a number of interesting possible solutions. The different types of available data can have different interpretations, different scales, and different representations. How does one go about merging the data into some reasonable single format for the system to utilize? Or should the system treat each type of data separately, and fuse the recommendation results of several different operations? Should all of the data be used throughout the recommendation process, or should each stage decide what pieces of data to use?

One specialized case of a hybrid system would be one that uses RFM (recency, frequency, monetary) data for a customer, and merges that with recommendations. RFM information is data that marketers have available, understand, and would like to use with recommendation systems. RFM data is how recently the customer has purchased from the Web site, how often the customer has purchased from the Web site, and how much the Web site has earned from the customer's purchases in the past. RFM analysis attempts to predict how much value the customer brings to the E-commerce site, and, hence, how much effort the site should be willing to put into

convincing the customer to purchase a product on this visit. Combining RFM analysis with a recommender system would enable a site to choose how much to discount a product for a customer based on the expected value of that customer to the site, and the expected value of the product to the customer. In a sense, current recommender systems are too focussed on the customer to make most marketing people happy. Marketing people would prefer recommender systems that are focussed on making the most money possible for the E-commerce site. As pressure comes from marketing organizations to create systems like these, important ethical issues are raised. To what extent is a recommendation from an E-commerce site a guarantee that the site really thinks the product is good for the customer? The situation is even more complicated if the site is paid by the producer of the product to make the recommendation!

Another challenge is providing confidence in recommendation quality. Generally speaking, confidence in the accuracy of a recommendation is most important in determining when to start actually presenting recommendations to a customer. Bluntly, no e-commerce site wants to look stupid to a customer. Since the cost of customer acquisition is high, many sites are more concerned about the risks of providing poor recommendations than about the benefits of providing good recommendations. A confidence metric would enable a site to determine when the system is ready to begin making reasonable recommendations. However, there are many difficulties to overcome.

Many popular algorithms for making recommendations do not easily lend themselves to producing confidence measures (Resnick et al. 1994, Shardanand et al. 1995, Lang 1995, Konstan et al. 1997). Even those that do often have the problem of producing low confidence values, even when the recommendations would be considered 'good enough'. The problem becomes one of not only providing a mechanism to determine the confidence of a particular recommendation, but also of providing a reasonable way for that confidence to be interpreted.

Scalability is yet another challenge facing current e-commerce recommendation systems. Large sites may have millions of users, hundreds of thousands of pages, and tens of millions of page views in a day. These same sites want to maintain the reactivity and responsiveness of smaller sites. Many commercial recommender systems have never handled a database that large, or a load that high. Many AI algorithms that work well for small-scale problems are too inefficient to be used for very large problems. New algorithms are needed that can handle very large-scale problems while maintaining the accuracy and near real-time response of smaller scale systems.

An important opportunity is to extend recommender systems to combine content and interest information. Content algorithms are able to determine which of a set of products are similar according to syntactic features of the products. For instance, books might be considered similar if they have similar keywords in their text. E-commerce sites have successfully used content algorithms, such as Amazon's "Eyes" program that lets users register to receive email when books with certain keywords are

published. These systems are different from recommender systems that rely on individual preferences for their recommendations. Several recent systems that combine recommender systems and content algorithms exist in the domain of content (Balabanovic et al. 1997, Sarwar et al. 1998, Basu et al. 1998), but we know of no such system for E-commerce. For instance, such a system might notice that a user tends to like books that have certain keywords in the publisher's description, and increase the prediction for other books with those keywords beyond the pure interest-based prediction. One such approach in the content domain combines both personalized agents (Maes 1995) and humans in a recommender system framework (Good et al. 1999). Schemes like these have the potential to combine the best features of both information filtering agents (Belkin et al. 1992), which are always ready with a content-based recommendation, and recommender systems, which are able to incorporate subjective factors like user interest and taste in their recommendations.

From E-Commerce to Commerce

One exciting new avenue of exploration that has grown up from the use of recommendation systems in e-commerce is the transition of the technology from the Internet to offline, more traditional, commerce. Call center, retail point-of-sale, direct mail marketing, and even in store kiosks are some of the new areas that recommendation systems are reaching:

Call Centers. Inbound call centers are using recommender systems to provide cross-sell recommendations to telephone operators. The call center software communicates the shopping basket to the recommender system, which suggests other products that the shopper will probably like. The call center operator then offers the products to the shopper over the phone. Outbound call centers can use recommender systems to choose customers to call for a particular marketing campaign. The recommender system chooses customers who are likely to buy the products being sold in the campaign. If the software is good, it should never recommend calling those of us who don't want to be called!

Retail Point-of-Sale: Retail point of sale applications involve cash registers connected to a computer. The cash registers can use the recommender system to suggest products to the clerk to recommend to the customer, or can modify the coupons printed by the coupon printer to print coupons for the products that will be most interesting to the customer.

Direct Mail Marketing: Sending a direct mail marketing letter is expensive, especially because the response rate is typically very low (1-2%). The recommender system can be used to select customers to receive the mail, or products to suggest to previously selected customers, increasing the response rate. We have heard anecdotally of response rates to direct email doubling with the use of a recommender system to recommend products to customers.

In-store Kiosks: Many customers walk out of stores without finding a product they want to purchase. An in-

store kiosk could recommend interesting products to the customers, making the customers feel better because they find products they like, and making the store owners feel better because they sell more products. For instance, in a video rental store a kiosk could suggest movies that a customer is likely to like, has not seen yet, and that are in stock.

Moving recommender systems off the Web is one of the first examples of a technology developed on the Internet transitioning to a traditional environment, and brings a unique set of challenges. Data collection in particular becomes a difficult problem. Purchase information will likely be the dominant data source initially, since it is easy to capture for credit card or check customers, or for members of a frequent buyer service. Many companies will likely want to make use of the richer range of options available with more types of interest data. Explicit ratings data will be more difficult to collect, because users will be interacting with a salesperson rather than a computer-driven Web site. An important challenge for the recommender system will be to infer interest data from the interaction with the customer through the customer interaction software. Other types of data will be easier to derive, such as whether a particular cross-sell was driven by a recommendation from the system. Having this information could lead to new types of self-tuning systems, provide better feedback to users, and help with confidence measures. Most traditional stores also have a vast amount of internal data representing relationships with their products and inventories. New ways of using this data will need to be investigated and discovered to open further inroads.

Conclusion

Recommender systems are demonstrating the practical importance of AI techniques, ranging from machine learning to personalized agents, in E-commerce. They are one of the most visible successes of AI technology on the Internet, and are in daily use on many of the most popular E-commerce sites. Recommender systems have the potential to become ubiquitous in E-commerce. In the not too distant future, every E-commerce transaction may be involve recommendations to the customer – if not before the purchase, then certainly after the purchase to increase cross-sell opportunities.

However, recommender systems face important challenges before they can become ubiquitous. The AI technologies underlying them must be extended to provide the following capabilities:

- 1) Hybrid data. How can recommender systems make use of all of the data available about a customer in making the best possible product recommendation?
- 2) Predictable quality. How can recommender systems tell in advance how accurate their recommendations are likely to be?
- 3) Scalability. How can recommender systems provide recommendations in hundredths of a second on

databases that involve millions of customers and millions of products?

- 4) Content. How can recommender systems be developed that use not only information about customer interests, but also syntactic information about the products being sold?

The next few years will be an exciting time in the evolution of recommender systems both in E-commerce and in traditional commerce applications. We look forward to discussing these issues at the workshop.

References

Balabanovic, M., and Shoham, Y. 1997. Fab: Content-Based, Collaborative Recommendation. *Communications of the ACM* 40(3):66-72.

Basu C., Hirsh H., and Cohen, W.W. 1998. Using Social and Content-Based Information in Recommendation. In *Proceedings of the AAAI-98*,:AAAI Press.

Belkin, N., and Croft, B.W. 1992. Information Filtering and Information Retrieval: Two Sides of the Same Coin?. *Communications of the ACM* 35(12):29-38.

Breese, J. Heckerman, D., and Kadie, C. 1998. Empirical Analysis of Predictive Algorithms for Collaborative Filtering. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, Madison, WI.

Cohen, W.W. 1995. Fast Effective Rule Induction. In *Proceedings of the Twelfth International Conference on Machine Learning*, Lake Tahoe, CA.:AAAI Press.

Goldberg, D., Nichols, D., Oki, B.M. and Terry, D. 1992. Using Collaborative Filtering to Weave an Information Tapestry. *Communications of the ACM* 35(12):61-70.

Good, N., Schafer, B., Konstan, J., Borchers, A., Sarwar, B., Herlocker, J., and Riedl, J. 1999. Combining Collaborative Filtering with Personal Agents for Better Recommendations. In *Proceedings of AAAI '99*, to appear.

Hill, W., Stead, L., Rosenstein, M., and Furnas, G., 1995. Recommending and Evaluating Choices in a Virtual Community of Use. In *Proceedings of ACM CHI'95*, 194-201. Denver, CO.: ACM.

Konstan, J.A., Miller, B.N., Maltz, D., Herlocker, J.L., Gordon, L.R., and Riedl, J. 1997. Applying Collaborative Filtering to Usenet News. *Communications of the ACM* 40(3):77-87.

Lang, K. 1995. NewsWeeder: Learning to Filter Netnews. *Proceedings of the Twelfth International Conference on Machine Learning*, 331-339, Tahoe City, CA, July.

Maes, P. 1995. Agents that Reduce Work and Information Overload. In *Readings in Human-Computer Interaction, Toward the Year 2000*, :Morgan Kauffman.

Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P. and Riedl, J. 1994. GroupLens: An Open Architecture for Collaborative Filtering of Netnews. In *Proceedings of 1994 Conference on Computer Supported Collaborative Work*, 175-186.: ACM.

Sarwar, B.M., Konstan, J.A., Borchers, A., Herlocker, J.L., Miller, B.N., and Riedl, J. 1998. Using Filtering Agents to Improve Prediction Quality in the GroupLens Research Collaborative Filtering System. In *Proceedings of CSCW '98*, Seattle, WA.: ACM.

Shardanand, U., and Maes, P. 1995. Social Information Filtering: Algorithms for Automating "Word of Mouth". In *Proceedings of ACM CHI '95*, Denver, CO.: ACM.

Terveen, L., Hill, W., Amento, B., McDonald, D., Creter J. 1997. PHOAKS: A System for Sharing Recommendations. *Communications of the ACM* 40(3):59-62.