2009

Strategic Learning In Recommendation Systems

Monica Johar
The University of North Carolina at Charlotte, msjohar@uncc.edu

Pelin Atahan
pxa041000@utdallas.eduThe University of Texas at Dallas, pxa041000@utdallas.edu

Sumit Sarkar
The University of Texas at Dallas, sumit@utdallas.edu

STRATEGIC LEARNING IN RECOMMENDATION SYSTEMS

Johar, Monica, The University of North Carolina at Charlotte, 9201 University City Blvd, Charlotte, NC 28223-0001, USA, msjohar@uncc.edu
Atahan, Pelin, The University of Texas at Dallas, 800 West Campbell Road, Richardson, TX 75080-3021, USA, pxa041000@utdallas.edu
Sarkar, Sumit, The University of Texas at Dallas, 800 West Campbell Road, Richardson, TX 75080-3021, USA, sumit@utdallas.edu

Abstract

Effective personalization can help firms reduce their customers' search costs and enhance customer loyalty. The personalization process consists of two important activities: learning and matching. Learning involves collecting data from a customer's interactions with the firm and then making inferences from the data about the customer's profile. Matching requires identifying which products to recommend or links to provide for making a sale. Prior research has typically looked at each activity in isolation. For instance, recent research has studied how a user's profile can be inferred quickly by offering items (links) that help discriminate user classes. Research on matching has typically assumed that all the recommendations in an interaction are made to generate immediate sales. We examine the problem of identifying items to offer such that both learning and matching are taken into consideration, thereby enabling the firm to achieve higher payoffs in the long run.

Keywords: *Bayesian Learning, Recommendation, Personalization, Electronic Commerce*

1 INTRODUCTION

Effective personalization can help firms reduce their customers' search costs and enhance customer loyalty. This, in turn, translates into increased cash inflows and enhanced profitability (Ansari and Mela 2003). Extant research has also shown that in electronic shopping environments, personalized product recommendation enable customers to identify superior products with less effort (Häubl and Trifts 2000). These works have demonstrated that personalization can be an effective tool for firms.

The personalization process consists of two important activities — learning and matching. Learning involves collecting data from a customer's interactions with the firm and then making inferences from the data about the customer's profile. For instance, the relevant profile for a customer may be her membership in one of several possible demographic or psychographic segments, which could be based on age, gender, zip code, income, political beliefs, etc. (Montgomery et al. 2001, Wall Street Journal October 17 2007). Matching is the process of identifying products to recommend based on what is known about the customer's profile.

Naturally, the quality of a customer's profile impacts the ability of the firm to provide high quality recommendations targeted towards sales (viz., the matching ability). The profiling activity is usually conducted first followed by matching that leads to specific recommendations. For instance, a profile may be explicitly obtained by directly asking the customer through registration forms or surveys, before the customer can gain access to the site (Miller et al. 2001). However, in online environments, firms are also able to dynamically profile a customer based on the customers' click-stream data (Padmanabhan et al. 2001). The profile is thus learnt implicitly by tracking the customer's interactions with the website where each link clicked on by a customer, or product purchased, represents a choice amongst a set of alternatives provided to the customer. Further, the firm could customize the offerings (offer set) on a page to learn the profile as quickly as possible.

Prior research has typically looked at each activity in isolation. For instance, Atahan and Sarkar (2007) have shown how a user's profile can be inferred quickly by offering those items (links) that help discriminate user segments (classes). Another approach to learn interests of new users is by presenting them with a number of items to rate. Studies have explored ways to address this problem by determining the most informative items that users can rate (Rashid et al 2002, Yu et al 2004). Research on matching has typically assumed that all the recommendations in an interaction are made to generate immediate sales (Breese et al 1998, Bodapati 2008). Interestingly, the interactive nature of web-based commerce enables a firm to actively consider both goals at the same time. However, extant research has not considered how a firm can strategically manage both concerns simultaneously – learning a customer's profile while also trying to sell products, so that the improved profile in earlier interactions can lead to an improved matching ability in later interactions. We examine the problem of identifying items to offer such that both learning and matching are taken into consideration, thereby enabling the firm to achieve higher payoffs in the long run.

The rest of the paper is organized as follows. In the next section, we provide a brief overview of related literature. We model the interaction process in Section 3, and discuss how the necessary probability parameters can be estimated in practice. In Section 4, we discuss how a firm should choose an offer set such that both learning and matching are taken into consideration. In Section 5, we conclude the paper by discussing the implications of this work for firms, as well as identifying areas of ongoing and future research.

2 LITERATURE REVIEW

As already mentioned, prior research on web-based personalization has looked at learning and matching activities in isolation. Several researchers have focused on which products to recommend or links to provide for making a sale (Breese et al 1998, Bodapati 2008). These recommendations are based on the matching piece of the personalization process, using techniques such as collaborative filtering and rule-based systems (Adomavicius and Tuzhilin 2005, Sandvig et. al. 2007). An assumption implicit in these works is that all the recommendations in an interaction are made to generate immediate sales.

Learning profiles has attracted considerable attention from researchers in the context of information retrieval and intelligent agents (Billsus and Pazzani 2000; Liu et al. 2004; Middleton et al 2004; Pazzani and Billsus1997; Widyantoro et al 2001; Wong and Butz 2000). These works show how to learn user preferences and interests in order to help users cope with the information overload on the internet. In these studies, the profiles are represented as feature (term) vectors with weights, and are updated using algorithms that update the weight of these vectors based on user feedback on documents viewed. Baglioni et al. (2003) attempt to learn the gender of visitors to a website based on their navigational history by applying various classification models. Montgomery (2001) studies learning demographic profiles based on user's web traversals using probability calculus.

Atahan and Sarkar (2007) present a learning approach that is most relevant to this work. They present a probabilistic technique to infer the relevant profile of a customer by tracking the items selected for examination by the customer at each interaction. Since we examine the problem of identifying items to offer such that both learning and matching are taken into consideration, we incorporate their learning approach in our analysis.

3 THE FRAMEWORK

Our objective is to develop a model that will enable us to analyze how a site can combine learning and matching strategically. In other words, the firm should take into account the impact of learning in one interaction on the probability of making a sale in future interactions. We model the interactions between a visitor and a site as shown in Figure 1. The figure shows the choices faced by the visitor at each interaction.

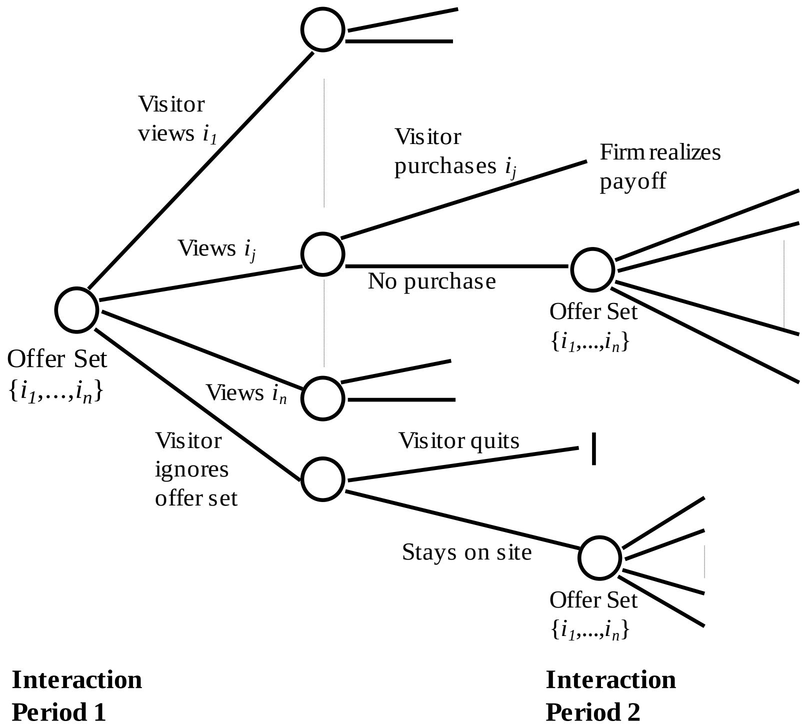

Figure 1. Interactions between a visitor and the site

The interactions are iterative in nature, with the firm offering a set of items (e.g., links on a page) to the visitor at each interaction. Given the offered items (that we call the offer set), the visitor can either choose to view detailed information on one of the items offered, or ignore the offer set. When the visitor clicks on a link associated with an item or ignores the offer set, the site can make inferences about the visitor's profile provided it tracks some pertinent information.

When the visitor views information on one of the items (say i_j) by clicking on the appropriate link, the site provides detailed product information for i_j, along with a new offer set (i.e., a new set of recommendations) in case the visitor does not like the product. If, on viewing the information on item i_j, the visitor decides to purchase that item, it results in a payoff to the firm. If the visitor does not purchase that item, then the visitor has the option of selecting an item from the new offer set for further evaluation, and the process repeats.

Two outcomes are possible if the visitor ignores the offer set. One is where the visitor decides to quit the site, thereby ending the session. The other is where the visitor stays on the site, moving to another page without selecting any of the recommended items/links (e.g., the visitor could use the back button, or use a search engine on the site to locate other products, etc.). The site continues to make recommendations to the visitor on the page the visitor lands on.

The above framework can be used to evaluate the expected payoff for the firm under different scenarios. If the firm is interested in maximizing its immediate payoff (i.e., payoff at each interaction), then it would focus only on matching. In that case, the offer set would be determined by optimizing the expected payoff at the end of each period, based on what the firm knows about the visitor's profile when composing the offer set. If, on the other hand, the firm is strategic, it would compose the offer set differently. It would consider not just the expected payoff from that offer set, but also how the offer set can enable better learning of the visitor's profile. Thus, the firm could look ahead and examine how by improving the profile learnt, the firm may make better recommendations in future interactions. The potential for increased payoffs in future interactions from this improved matching ability would then be factored into the determination of the offer set.

To operationalize this framework, the firm would need a method to learn and update the visitor's profile based on the selections made by the visitor at each interaction. We incorporate the probabilistic technique presented by Atahan and Sarkar (2007) to learn the profile information of a customer. The firm would also need to estimate the probability parameters associated with the choices made by the visitor. These include the following:

- The probability distribution (profile) associated with a visitor based on the observed interactions at any stage of the interaction process.
- The probability that a visitor associated with a given profile will view item i_j when presented with an offer set $O=\{i_1,\ldots,i_n\}$.
- The probability that such a visitor will purchase item i_j after viewing information on that item.
- The probability that a visitor does not find any of the offered items to be of interest.
- The probability that a visitor quits the site when s/he does not find any of the offered items to be interesting.

One approach to obtain the necessary parameters is by directly estimating them based on the historical data on customer interactions at that site. While that could be feasible for some of above parameters, it would be very difficult for others. A potential problem is that the number of feasible offer sets would be typically very large. As a result, it will often not be possible for the firm to directly estimate all of these probability parameters. However, by judiciously tracking some item-level interactions, it is possible to estimate the above parameters. The profile learning approach presented by Atahan and Sarkar (2007) also involves determining some of the probability parameters used in our model based on item-level interactions. We describe how the probability parameters can be estimated based on item level interactions in the remainder of the section.

3.1 Probability Parameters Associated with Selecting an Offered Link

We illustrate using an example the determination of the probability parameters associated with a visitor's decision to view an item. Consider a firm (e.g., Amazon) that wants to learn the gender of a visitor who is traversing its website in order to make gender appropriate recommendations. The attribute gender (G) takes on values male (m) and female (f). Imagine that the site has offered to the visitor links to two items i_1 and i_2. Given this offer set, the site would need to determine the probability the visitor will click on either of the items or ignore the offer set. In this case, the information needed are (i) the current profile of the visitor, represented as $P(m/IH)$ and $P(f/IH)$, where IH (short for *item history*) denotes the visitor's choices prior to the current interaction, (ii) the likelihood that an item is of interest to an average visitor, denoted by $P(i_j)$, (iii) the prior distribution of a visitor's gender, $P(m)$ and $P(f)$, and (iv) the typical profile of visitors who view each item, denoted by $P(m|i_j)$ and $P(f|i_j)$, respectively.

The current profile is represented as the probability the visitor belongs to the possible classes, m and f in this case, given the visitor's prior interactions (we discuss how the site learns the profile based on the visitor's interactions with the site in the subsequent section). $P(i_j)$ is the likelihood that information on item i_j is of interest to an average visitor, and may be estimated as the proportion of visitors who have viewed the information on item i_j when it has been offered as a link to them at some point in their interactions with the site. This information can be obtained from the server's log files. The prior distribution for the class (i.e., gender in this example) refers to the proportions of visitors to the site who are male and female, respectively. $P(m|\ i_j)$ is the proportion of visitors viewing item i_j that are male. The statistics $P(m)$ and $P(m|\ i_j)$ can be obtained by sampling a subset of visitors, or obtained from professional market research agencies such as comScore networks. Such agencies collect personal information from a large number of users and track their online activities in order to estimate such probabilities.

Let $O = \{i_1, i_2\}$ denote the offer set. The probability of interest is $P(m|i_1,O)$ (note that $P(f|i_1,O)$ is equal to 1- $P(m|i_1,O)$). To obtain this, it is first necessary to derive the probability that the visitor will view item i_1 given the visitor is male, and the probability given the visitor is female. Assuming that the visitor views one of the two items, Atahan and Sarkar estimate these probabilities as follows:

$$P(i_1|m,O,V) = \frac{P(m|i_1)P(i_1)}{P(m|i_1)P(i_1) + P(m|i_2)P(i_2)}.$$

The conditioning event V in the above expression recognizes the assumption that one of the items is viewed. Consequently, $P(i_2|m,O,V) = 1 - P(i_1|m,O,V)$. The corresponding probabilities for female visitors are similarly obtained as:

$$P(i_1|f,O,V) = \frac{P(f|i_1)P(i_1)}{P(f|i_1)P(i_1) + P(f|i_2)P(i_2)} \text{ and } P(i_2|f,O,V) = 1 - P(i_1|f,O,V).$$

Atahan and Sarkar are interested in how profile learning works, and therefore assume that the visitor views one of the offered items; if the visitor views neither of them, the belief regarding the gender of the visitor remains unchanged. In our analysis, we also need to consider the situation where the visitor does not view either item, as that will impact the probability that they eventually purchase one of the offered items. The probability that the visitor does not view either item (in other words, does not click on either link) is interpreted as the probability that the visitor is not interested in either item. Let $P(i_1|m)$ be the probability that item i_1 is of interest to a visitor who is male. Using the parameters discussed earlier, we can estimate this as:

$$P(i_1|m) = \frac{P(m|i_1)P(i_1)}{P(m)}.$$

$P(i_2|m)$ can be obtained analogously. Then, the probability that a male visitor is not interested in either of the items offered, denoted $P(\phi|m,O)$, can be estimated as follows:

$$P(\phi|m,O) = (1 - P(i_1|m)) * (1 - P(i_2|m)).$$

This assumes that the probability a visitor does not have any interest in item i_1 is conditionally independent of the probability that the visitor would be interested in item i_2, given the visitor's gender. The probability that a male visitor will view one of the offered items is then $P(V|\ m,O) = 1 - P(\phi|m,O)$. The probability that a female visitor will view one of the items offered, $P(V|\ f,O)$, can be estimated similarly.

The probability that a male visitor will view item i_1 when the assumption that the visitor views an offered item is relaxed, $P(i_1|m,O)$, is then obtained as:

$$P(i_1|m,O) = P(i_1|m,O,V) * P(V|m,O).$$

Similarly, $P(i_1|f,O) = P(i_1|f,O,V) * P(V|f,O)$.

Then, the probability that the visitor will view item i_1, $P(i_1|O)$, is

$$P(i_1|O) = P(i_1|m,O) * P(m/IH) + P(i_1|f,O) * P(f/IH).$$

The probability that the visitor will not view either item, $P(\phi|O)$, is

$$P(\phi|O) = P(\phi|m,O) * P(m/IH)) + P(\phi|f,O) * P(f/IH).$$

3.2 Learning a Visitor's Profile

We have discussed that the profile is represented by the possible values (classes) a profile attribute can take and the visitor's likelihood of belonging to each of these classes given the visitor's prior interactions with the site. For a new visitor, initially the item history would be empty and the profile would be equivalent to the population priors, $P(m)$ and $P(f)$. When a visitor makes a selection, either views an offered item or ignores the offer set, the site can update the beliefs regarding the visitor's profile. Atahan and Sarkar show how the beliefs are revised when the visitor views an offered item. For instance, the probability that a visitor who has item history *IH* is male if the visitor views an additional item i_1 is

$$P(m|IH,i_1,O) = \frac{P(i_1|m,O)P(m|IH)}{P(i_1|m,O)P(m|IH)+P(i_1|f,O)P(f|IH)} \ .$$

Therefore, the probability that the visitor is female if item i_1 is viewed is $P(f|i_1,O) = 1- P(m|i_1,O)$.

Likewise, the site can revise the likelihood that the visitor is male when the visitor ignores the offer set as

$$P(m|IH,\phi,O) = \frac{P(\phi|m,O)P(m|IH)}{P(\phi|m,O)P(m|IH)+P(\phi|f,O)P(f|IH)} \ .$$

The belief revision process is particularly efficient if the probability of viewing information on an offered item is assumed to be independent of the items already viewed given the gender.

3.3 Other Probability Parameters

Two parameters remain to be discussed. The first remaining parameter needed is the probability that a visitor will purchase item i_j after viewing information on that item. A firm can estimate this probability by observing what proportion of visitors who view information on each item go on to purchase that item. This information can be obtained from the site's log files. If the site has sampled user behavior for different classes separately, then this estimate can be further refined by incorporating the class specific estimates, as shown below.

$$P(Purchase|IH,O,i_j) = P(Purchase|i_j,m)P(m|IH,i_j) + P(Purchase|i_j,f)P(f|IH,i_j)$$

This assumes that the effect of the item history viewed by the visitor directly affects the belief about the visitor's class, and the probability of purchasing the item i_j is conditionally independent of the item history given the class and the fact that the item has been viewed.

The remaining parameter of interest is the probability that a visitor will quit the site (instead of staying on the site) when none of the items in the offer set are of interest to the visitor. Once again, this can be easily estimated from the site's log files. A straightforward way to estimate this is to determine the proportion of offer sets to all visitors (over a reasonably large period of time) that led to the visitor quitting the site, as compared to the visitor ignoring the offer set but staying on the site. Again, if the site has sampled user behavior for different classes separately, then this probability can be refined using the class specific estimates.

4 DETERMINING THE OFFER SET

As outlined in Section 3, a firm can strategically compose an offer set that enables it to learn a visitor's profile quickly so that it can better match the recommendations to the visitor's profile in subsequent interactions (periods). Ideally, the firm would like to take into account the benefit from all future periods when calculating the expected payoff from an offer set. However, considering multiple future periods would be computationally intensive, and therefore difficult to implement in real time. Therefore, we limit the scope of the current analysis to a *one-interaction look-ahead* model, i.e., the firm takes into account the impact of the composition of the offer set on learning in interaction *t* and the expected payoff in interactions *t* and *t+1*. This helps to capture the effect of learning from an offer set in the current interaction to better recommendations (i.e., greater likelihood of purchase) in the following interaction.

5 DISCUSSION AND CONCLUSIONS

We have presented a methodology that online firms could use to incorporate learning and selling in a strategic manner in order to improve their long term payoffs. While the methodology is rigorous, obtaining optimal solutions may be difficult for sites that have very large number of items to offer given the computational complexity involved. In future work, we will develop heuristics that can provide good solutions in an efficient manner, so that firms can employ this methodology effectively in real time. We plan to conduct experiments to determine the benefits to firms when using our proposed approach as compared to benchmark approaches that do not consider the strategic aspects of learning the user profile quickly.

The proposed methodology has several important implications for online firms. In addition to improved payoffs, it also provides a firm the opportunity to provide better recommendations in the long run. This should eventually lead to higher customer satisfaction and increased loyalty. Future research should examine the benefits that could accrue to firms from improved long term interactions.

References

Adomavicius, G., and Tuzhilin, A (2005). Personalization Technologies: A Process Oriented Perspective. Communications of the ACM, 48(10), 83-90.

Ansari, A., and Mela, C. F. (2003). E-Customization. Journal of Marketing Research. Vol. XL, 131-145.

Atahan, P., and Sarkar, S. (2007) Designing Websites to Learn User Profiles Quickly. In Proceedings of the WITS, Montreal, Canada.

Baglioni, M., Ferrara, U., Romei, A., Ruggieri, S., and Turini, F. (2003). Preprocessing and Mining Web Log Data for Web Personalization. In Proceedings of the AI*IA 2003: Advances in Artificial Intelligence, 8th Congress of the Italian Association for Artificial Intelligence, p. 237, Pisa, Italy.

Billsus, D., and Pazzani, M. J. (2000). User Modeling for Adaptive News Access. User Modeling and User-Adapted Interaction, 10(2-3), 147-180.

Bodapati, A. (2008). Recommendation systems with purchase data. Journal of Marketing Research. XLV, 77-93.

Breese, J. S., Heckerman, D., and Kadie, C. (1998). Empirical Analysis of Predictive Algorithms for Collaborative Filtering. In Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence p. 43, Madison, WI.

Häubl, G., and Trifts, V.(2000). Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids. Marketing Science, 19(1), 4-21.

Linden, G. (2008).People who read this article also read … . ,” IEEE Spectrum, 45(3).

Liu, F., Yu, C., and Meng, W. (2004). Personalized Web Search for Improving Retrieval Effectiveness. IEEE Transactions on Knowledge and Data Engineering. 16(1), 28-40.

Middleton, S. E., Shadbolt, N. R., and De Roure, D. C. (2004). Ontological User Profiling in Recommender Systems. ACM Transactions on Information Systems, 22(1), 54–88.

Miller, T.W., and Dickson, P. R (2001). On-line market research. International Journal of Electronic Commerce, 5(3), 139-167.

Montgomery, A.L. (2001). Applying Quantitative Marketing Techniques to the Internet. Interface, 30(2), 90-108.

Padmanabhan, B., Zheng, Z., and Kimbrough, S.O. (2001). Personalization from incomplete data: What you don’t know can hurt. In Proceedings of the Seventh ACM SIGKDD International Conference on KDD, San Francisco, CA.

Pazzani, M.J., and Billsus, D. (1997). Learning and Revising User Profiles: The Identification of Interesting WebSites. Machine Learning, 27(3), 313–331.

Sandvig, J. J., Mobasher, B., and Burke, R. (2007). Robustness of collaborative recommendation based on association rule mining. In Proceedings of the 2007 Conference on Recommender Systems, p. 105, Minneapolis, MN.

Widyantoro, D. H., Ioerger, T. R., and Yen, J. (2001). Learning User Interest Dynamics with a Three-Descriptor Representation. Journal of the American Society of Information Science and Technology (JASIST), 52(3), 212-225.

Wong, S. K. M., and Butz, C. J. (2000) A Bayesian Approach to User Profiling in Information Retrieval. Technology Letters ,4(1), 50-56.

The Wall Street Journal. Firm Mines Offline Data to Target Online Ads. October 17, 2007.