

Commercial Applications of Machine Learning for Personalized Wireless Portals

Michael J. Pazzani

Abstract. Consumers and businesses have access to vast stores of information on the Internet, ranging from newspapers, shopping catalogs, restaurant guides, classified ads, jobs listings, dating services to discussion groups and e-mail. All this information is typically accessible only while users are in front of a computer at home or in an office. Wireless devices allow unprecedented access to information from any location at any time. The presentation of this information must be tailored to the constraints of mobile devices. Small screens, slower connections, high latency and limited input capabilities present new challenges. Agents that learn user's preferences and select information for the user are a convenience when displaying information on a 19-inch desktop monitor accessed over a broadband connection; they are essential on a handheld wireless device. This paper summarizes commercially deployed systems using machine learning methods for personalizing mobile information delivery.

1 The Need for Adaptive Personalization

Many wireless data services cope with small screens and limited capabilities of wireless devices by reducing the amount of information available to users. One news site makes four business articles available to mobile users and one wireless carrier limits the user to less than ten wireless games to download. This is equivalent to a newspaper having only the front page of each section available or an electronic game store having a single shelf of products.

However, it is rarely the case that a small offering can satisfy the needs of a diverse group of consumers. Here, we describe the experience that AdaptiveInfo has had with personalizing delivery of wireless news and downloadable content such as Java applications for mobile phones. The experience with wireless readers of the Los Angeles Times wireless web site is that users read much more than what appears on the front page of the print edition— only 31.6% of the content read wirelessly is among the top five stories in each section on the regular web site, showing that mobile users want more content than a one-size-fits-all approach can deliver. Similarly, in an analysis of downloadable content for mobile personal digital assistants, we found that the top 10 items account for less than 40% of the sales for games and less than 30% of the sales for utilities. While a limited one-size-fits-all approach is easy to use, if a mobile wireless experience is to meet the demands of consumers, it cannot limit its offerings to a small number of options.

Many wireless carriers offer a limited degree of customization, where users fill out forms with check boxes to select items of particular interest. Most carriers report that only 2%–5% of users customize their experience. Three factors contribute to the limited usage of such approaches. First, it is often complicated to create an account on the web and associate that account with a wireless device, limiting the audience to the technically sophisticated and to those who plan ahead of time. Second, the options tend to be very coarse-grained, allowing users to select from general categories such as “Business News” rather than specific topics such as “Inflation.” Third, topic profiles require maintenance to be useful. For example, many people don’t regularly follow politics but gain interest during elections. Others closely follow particular sports teams but change their reading habits during the Olympics. Few users are willing to continually maintain their customized profile on a regular web site for an optimal wireless web experience.

Explicit customization approaches, such as web-based questionnaires, place the burden of personalization on the user. We advocate adaptive personalization, an automated approach that uses machine learning techniques to construct a model of each user’s interests. It is important that adaptive systems learn from a few examples so that the benefits of personalization are available to new users. It is equally important that the personalization system adapt quickly to changing user interests so that they don’t frustrate long-term customers by preventing them from finding novel items. To be truly useful in a mobile context, user’s interests must be inferred implicitly from actions and not obtained exclusively from explicit content ratings provided by the user.

2 Machine Learning of User Models

We have pursued a hybrid approach to user modeling [1,2] that combines similarity-based machine learning algorithms [3], Bayesian algorithms [4] and community-based algorithms [5]. We uniquely (but anonymously) identify each user and keep a history of the items presented and how the user interacts with them. Simple rules use this implicit feedback on items to create positive and negative training examples passed to each machine learning algorithm. The similarity-based learning methods are particularly useful in identifying new trends while the Bayesian methods are useful for identifying general themes of interest. The hybrid approach combining recommendations from the above approaches achieves a balance of learning and adapting quickly to changing interests while avoiding brittleness. An additional use of the similarity-based component is to make sure that each screen does not contain too many similar items. Finally, in many domains, we also try to leverage existing efforts to define a baseline or default order to use in the absence of personalized information. For example, we take editorial input into account in ordering news stories by boosting the priority of lead stories. The effect of this boosting is that first-time users of the wireless news site see articles in the same order as on the unpersonalized wired web site and that all users always see the lead story in each section. This also allows the adaptive personalization engine to learn more about each

user. Users who elect not to receive personalized content for privacy reasons obtain this same default order.

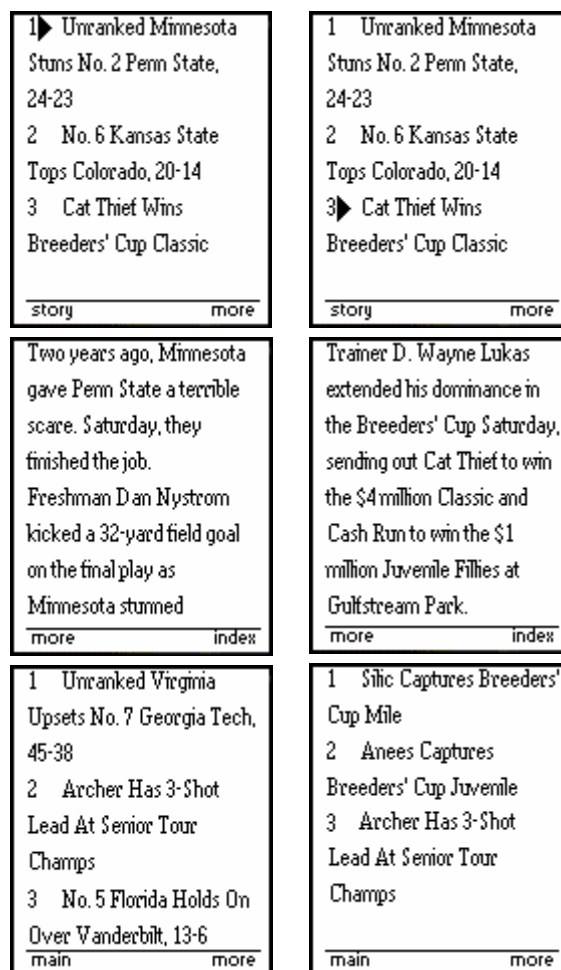


Fig. 1. (Left) Trace of a first-time user interested in football. After reading a football article, the user asks for the next set of headlines and is sent additional football headlines. (Right) Trace of a first-time user interested in horse racing.

Figure 1 shows an example of two different consumers using the same mobile news site for the first time. Both users are presented with the same initial set of articles. One user selects a college football story, and then when that user requests additional sports news, other college football stories are shown. In contrast, another user selects a horse racing story and receives additional horse racing headlines on the

third screen. Note that both users are also shown a headline of a golf story to make sure that a variety of content is shown.

Scalability of the machine learning algorithms is a key concern. Because our target mobile devices have limited processing power and memory and slow network access (cf. [4]), our personalization technology resides on a server that interacts with a browser on the devices (or a program implemented in Qualcomm's BREW or Sun Microsystem's J2ME). With inverted indices, caching of similarity calculations and distributing computation across multiple servers, we are able to meet the performance needs of personalized content delivery. An added benefit of the inverted index is that it does support full text search for content, and the similarity cache facilitates finding related content.

We have compared personalized and non-personalized news access under a variety of conditions to determine the benefits of personalization. We accomplish this either by alternating personalized and non-personalized news to the same user on different days, or sending one group of users personalized content and another group non-personalized content. The following two findings indicate that increased usage and increased loyalty among users are among the benefits of personalization:

1. In a visit to a wireless site, users read over 40% more content that is personalized to their interests [1].
2. Over a six-week period, viewers who received personalized content returned to the site over 60% more often than did users who had not received personalized content [2].

3 Applications of Adaptive Personalization

Adaptive personalization provides benefits any time there are more options than easily fit on the screen, and each individual explores only a subset of the available options. Many mobile applications are improved by the use of this technology.

Downloadable Content. A growing trend in the wireless industry is to allow users to download applications and content such as games, ring tones and images to cell phones. By providing a personalized interface to such a catalog, downloadable items that meet a user's preferences are shown first- simplifying shopping for applications and increasing the number of downloads.

News. An adaptive interface to news enhances the reader's experience. It makes it easy to follow news stories that unfold across several days. When the user is reading letters to the editor or opinion pages, items related to articles read in the past are displayed. By ensuring that a variety of articles of potential interest are displayed to the user, including breaking news and lead stories, users stay informed of important events while pursuing their particular interests.

Classified Ads. Wireless classified ads benefit considerably from adaptive personalization. It is difficult to page through long lists on most mobile devices. Therefore, it is critical to get the most personally relevant information to the user on the initial screens. Personalization also saves users from filling out many forms to specify their interests precisely. Instead, the user typically enters a more general specification and the adaptive interface learns about the user's precise requirements. Other applications related to classified ads include online dating services and online employment listings.

Restaurants and Entertainment Listings. Adaptive personalization automatically learns customers' tastes in restaurants, nightclubs, etc. Positive feedback is indicated by actions such as calling the restaurant or looking for directions. Many restaurant listing services simply list restaurants by distance from the customer. In contrast, an adaptive interface uses distance as one factor, but orders the nearby restaurants according to the customer's preferences. Furthermore, diversity is factored into the listings so that many restaurants of the same chain are not included in the same display.

4. Conclusion

False starts in the wireless industry have shown that a literal translation of wired web sites to mobile web sites results in service that is not usable. Merely squeezing data onto small screens detracts from the user experience on mobile devices. By adaptively learning users' preferences, all users can have easy access to information that is personally relevant while allowing a mobile operator to provide a large offering to meet a diverse set of needs.

References

1. Billsus, D. & Pazzani, M. (2000). User Modeling for Adaptive News Access, *User Modeling and User-Adapted Interaction*, 10(2/3): 147-180.
2. Billsus, D., Brunk, C., Evans, C., Gladish, B. & Pazzani, M.. (2002) . *Adaptive Interfaces for Ubiquitous Web Access*. Commun. ACM 45, 5.
3. Cover, T. & Hart, P. (1967). Nearest Neighbor pattern classification, *IEEE Transactions on Information Theory*, 13, pp. 21-27.
4. Duda, R. & Hart, P. (1973). *Pattern Classification and Scene Analysis*. New York, NY: Wiley and Sons.
5. Pazzani, M. (1999). A Framework for Collaborative, Content-Based and Demographic Filtering. *Artificial Intelligence Review* 13(5-6): 393-408