

ADAPTIVE INTERFACES FOR UBIQUITOUS WEB ACCESS

[DANIEL BILLSUS, CLIFFORD A. BRUNK,
CRAIG EVANS, BRIAN GLADISH, AND
MICHAEL PAZZANI]

Allowing mobile users to access any information at any time from any location.

The invention of the movable type printing press launched the information age by making the mass distribution of information both feasible and economical. Newspapers, magazines, shopping catalogs, restaurant guides, and classified advertisements can trace their origins to the printing process. Five and a half centuries of technological progress in communications networks, protocols, computers, and user interface design led to the Web, online publishing, and e-commerce. Consumers and businesses have access to vast stores of information. All this information, however, used to be accessible only while users were tethered to a computer at home or in an office. Wireless data and voice access to this vast store allows unprecedented access to information from any location at any time.

The presentation of this information must be tailored to the constraints of mobile devices [3]. Although browsing and searching are the acceptable methods of locating information on the wired Web, those operations soon become cumbersome and inefficient in the wireless setting and nearly impossible in voice interfaces. Small screens, slower connections, high latency, limited input capabilities, and the serial nature of voice interfaces pre-

sent new challenges. Agents that select information for the user [9] 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. To give one real-world example, the *Los Angeles Times* Web site contains approximately 1,000 current news articles at any point in time, 10,000 local restaurant listings, and

100,000 classified advertisements, whereas somewhere between three and nine headlines, restaurant names, or summaries fit on the screen of a wireless device. Adaptive personalization technology that makes all of this information available to the mobile user is discussed here. The goal of adaptive personalization is to increase the usage and acceptance of mobile access through content that is easily accessible and personally relevant.

The Need for Adaptive Personalization

Many wireless data and voice 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 sports, business, and other stories available to mobile users, and one restaurant site recommends 10 restaurants per city. This is equivalent to a newspaper's just having the front page of each section available. The experience

with wireless readers of the *Los Angeles Times* 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.

Many wireless carriers offer a limited degree of personalization, where users customize their menus by visiting a Web site and selecting the order in which items appear or indicating topics of particular interest. Most carriers report that only 2%–5% of users customize their interfaces. Several factors contribute to the limited usage of such approaches. First, it is fairly complicated to create an account on the Web and associate that account with a wireless device, which limits the user base to the technically sophisticated. Second, the categories tend to be very coarse-grained, allowing users to select from general categories such as “International News” rather than specific categories such as “telecom industry in Korea.” Third, topic profiles require maintenance to be useful. Many people don't regularly follow sports but gain interest in championship games. Others closely follow particular teams but change their reading habits during the Olympics. Few users are willing to continually maintain their cus-

tomized profile on a regular Web site for an optimal wireless Web experience. Finally, such approaches may prevent wireless users from gaining access to information outside their profile without going back to a wired Web site and adjusting the settings. This clearly is in conflict with the goal of being able to access any information at any time from any location.

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 artificial intelligence and statistical techniques to construct a model of each user's interests. It is important that adaptive systems learn from a few examples and adapt quickly to changing user interests. To be truly useful in a mobile context, user's interests must be




Figure 1. Personalized access to classified advertisements.

inferred implicitly from actions [5] and not obtained exclusively from explicit content ratings provided by the user. For example, the more a user reads of an email message or news story, the higher his or her level of interest. Obtaining explicit feedback about user interest requires a user interface consuming screen real estate and additional data to be transmitted. We also argue that the adaptive interface should not limit the user's choices by filtering information, but rather order it so that the most relevant information is at the top of each screen.

Figure 1 shows how adaptive personalization helps a user shop for a used car in wireless classified ads. The user enters a partial specification of a desired item, in this case, a Ford in the area of the 90210 zip code, and then selects “Find” (Figure 1, part a). A list of possibilities is shown (b), and the user selects the second car and looks at the description (c). Scrolling down in the same page (d), the user has a number of options, including placing a call, saving the ad, and looking at similar cars. In addition to helping the user, these options provide implicit feedback on the user's interest in the item: viewing an ad is a sign of moderate interest, whereas saving the ad or forwarding it via



Figure 2. A personalized restaurant finder locates nearby restaurants that match the user's interests. Other services simply list restaurants by distance. With a small screen, it's important to put the most personally useful options first: (a) location only; (b) location and personalization.

email are signs of strong interest. In this case, let's assume the user just goes back to the list of ads displayed in (b) and then hits 9 to see more ads. The next page of ads (e) has been personalized for the user and shows cars similar to the car viewed. Although, in many cases, quickly honing in on that model would be helpful to the user, when the user doesn't select one of the similar ads and looks for more ads, negative feedback is given to the adaptive interface and a more diverse group of cars is displayed (Figure 1, part f).

Requirements for Adaptive Personalization and Adaptive Interfaces

We believe that for adaptive personalization to meet the difficult test of user acceptance, it must be designed with certain constraints in mind:

- Provide a good initial experience and learn quickly for new users. The first use should provide an acceptable, nonpersonalized experience. The benefits of adaptive personalization must be apparent within the first few uses if users are to return. Furthermore, the transition from a nonpersonalized to a personalized experience must be smooth without going through a state of presenting random items to the user. Most of the results in information categorization, [11, 12] report on asymptotic performance, whereas in an adaptive interface, the performance on a handful of examples is very important. At the *Los Angeles Times* wireless site,

adaptive personalization is combined with explicit customization to provide a starting point for power users. Users may select sports teams or companies to follow. Fewer than 3% of the users have explicitly set their preferences, emphasizing the importance of adaptive methods.

- Adapt quickly to changing interests. Users often change their reading habits due to external events. A user who skips 160 stories on a baseball team during the regular season and reads about that same team's victory in four league championship games is probably more than 4/164ths interested in the World Series.
- Avoid tunnel vision. Personalization should not get in the way of finding important novel information or breaking news stories.
- Do not require hand-tagging of content with category labels. A process that requires the content provider to do extra work by hand such as adding metatags with category labels is not feasible with thousands of new items being added daily.
- Avoid brittleness. A single action, such as selecting something accidentally or skipping over an article on a topic (perhaps because one heard about the details on the radio or the wireless connection is dropped) should not have a drastic and unrecoverable effect on the presentation.
- Support multiple modes of information access. Although we have found that most of the time users access news through section headings such as sports, about 8% of the *Los Angeles Times* content served wirelessly is accessed by users searching for keywords. The order of search results is also personalized so that if there are many stories with the same keywords, the ones most similar to content read in the past are prioritized. Users can also access content related to the current story through "Related Stories" links that are automatically inserted. Approximately 12% of the *Los Angeles Times* content served wirelessly is accessed through related stories links.
- Respect individuals' privacy. Although personalization is designed to help users by minimizing the amount of work they must do to access relevant information, different users have different trade-offs for convenience and privacy. Users can opt to have the server save identifying information to avoid repeatedly typing the same address when using features such as forwarding items by email. However, users may remain anonymous while providing the benefits of personalization, turn off all personalization features, or delete their data from the server [8].

We found that a combination of similarity-based methods [6], Bayesian methods [7], and collaborative methods [1, 10] meets the preceding requirements by achieving a balance of learning and adapting quickly to changing interests while avoiding brittleness. However, we take editorial input into account in ordering content for delivery by boosting the priority of lead stories. The effect is that first-time users of the wireless site see articles in the same order as on the wired Web site and all users always see the lead story in each section, allowing the adaptive personalization engine to learn more about each user. Finally, the similarity-based methods also are used to make sure a variety of news articles, products, or restaurant listings is presented on each screen in much the same manner a newspaper does not fill the front page with articles on



Figure 3. Adaptive defaults save the user time when visiting frequently accessed sites: (a) wireless portal menu; (b) adaptive defaults on menu.

the same topic.

Of course, response time and scalability are also key concerns. Because our target mobile devices have limited processing power and memory and slow network access (see [4]), our personalization technology resides on a server that interacts with a browser on the devices (or a program implemented in Qualcomm's BREW). 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 searching for content, and the similarity cache facilitates finding related articles.

Applications of Adaptive Personalization

Adaptive personalization provides benefits anytime 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.

Classified Advertising. Wireless classified advertising benefits 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.

News. An adaptive interface to news enhances the reader's experience. It makes it easy to follow news stories that occur over several days. When the user is reading letters to the editor or editorial 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. The average *Los Angeles Times* reader views more than seven pages in a session on his or her mobile device, a figure that exceeds many wired Web sites. A significant number of readers of the *Los Angeles Times* are outside California and even outside the U.S. The personalization engine usually skips news specific to Southern California for these users and presents world and national news of interest. We have compared personalized and nonpersonalized news access under a variety of conditions to determine the benefits of personalization. We accomplish this either by alternating personalized and nonpersonalized news to the same user on different days, or sending one group of users personalized content and another group nonpersonalized content. The following two findings indicate increased usage and increased loyalty among users are among the benefits of personalization: In a visit to a wireless site, users read over 40% more content personalized to their interests [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.

Restaurants and Entertainment Listings. Adaptive personalization automatically learns customers' tastes in restaurants, nightclubs, and so forth. 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 (see Figure 2, part a). In contrast, an adaptive interface uses distance as one factor, but orders the nearby restaurants according to the customer's preferences (see Figure 2, part b). Furthermore, diversity is factored into the listings so that many restaurants of the same chain are not included

in the same display. Although we have illustrated this example with restaurants, adaptive personalization provides benefits for locating nightclubs, comedy clubs, and other points of interest.

Wireless Portals. The menu structure of a wireless portal divides sites into logical categories such as Mail, News, and Shopping (see Figure 3, part a). Selecting one of these menu items (such as Mail) takes a user to a menu of possible mail providers such as Yahoo and AOL. For example, selecting Yahoo allows users to access their Yahoo mail account. Because most users access a single mail account, seeing a list of alternative email providers each time is a waste of time and network bandwidth. However, we would still like users to have access to alternative accounts in case they change email providers or loan the phone or hand-held to a colleague to check email. An adaptive solution to this problem is to log which sites the user accesses from a hierarchical menu and make the most frequently accessed site the default for that user. As with the initial home screen, the menu of all options is transmitted and presented if the left button or number key is pressed, but the default for each category is available by simply clicking on the right button. This eliminates the need to see and transmit another menu, saving the user time in accessing frequently visited sites.

Voice Portals. Services such as Tellme and BeVocal allow users to access recorded or synthesized content over any telephone. Such services are also starting to appear in automobiles. Anyone who has interacted with an airline or bank automated phone system can understand the benefits of presenting information in an order that is relevant to the user. In a voice interface, adaptive personalization can change both the order in which categories are listed, as well as the order in which items within categories are spoken. Some users who listen to news prefer to hear business stories, whereas others prefer sports. Some sports users care about horse racing and others, tennis. By adaptively learning preferences, users can save time while accessing personally relevant information.

Wired Web. Concentrating on mobile access has forced us to think carefully about what capabilities and interface elements are essential. However, adaptive personalization technology is not tied to any delivery channel and provides benefits for Internet access from a desktop computer as well. For example, it's not uncommon to have a few hundred cars, apartments, or jobs match in an online classified ad system. Adaptive methods can be used to prioritize which ads are most likely to be of interest to each shopper and display them on the first screen. In an online news site, adaptive personalization puts the news stories most likely to

be of interest to the reader above the fold (that is, visible at the top of the screen without scrolling). AdaptiveInfo is currently working with a large U.S. media company to roll out an adaptive interface for a group of wired and wireless news Web sites.

The list of applications for adaptive interfaces on mobile devices is by no means comprehensive. Other applications of adaptive personalization include prioritizing email messages, finding applications to download onto 3G phones, keeping track of sales leads, shopping from product catalogs, searching through real estate listings, and browsing online dating services. Furthermore, the server solution can easily be extended to handle new devices as they become available in the rapidly evolving world of mobile communications.

Conclusion

A literal translation from the Web to wireless is inadequate. 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 a vast amount of information at any time. ■

REFERENCES

1. Billsus, D. and Pazzani, M. Learning collaborative information filters. In *Proceedings of the Fifteenth International Conference on Machine Learning* (Madison, WI, 1998), Morgan Kaufmann, 46–54.
2. Billsus, D. and Pazzani, M. User modeling for adaptive news access. *User Modeling and User-Adapted Interaction* 10, 2/3 (2000), 147–180.
3. Buchanan, G., Farrant, S., Jones, M., Marsden, G. Pazzani, M., and Thimbleby, H. Improving mobile Internet usability. In *Proceedings of the Tenth International World Wide Web Conference*, Hong Kong, 2001, 673–680.
4. Cheverst, K., Mitchell, K., and Davies, N. The role of adaptive hypermedia in a context-aware tourist GUIDE. *Commun. ACM* 45, 5 (May 2002).
5. Claypool, M., Le, P., Wased, M. and Brown, D. Implicit interest indicators. In *Proceedings of the International Conference on Intelligent User Interfaces*, (Santa Fe, 2001), 33–40.
6. Cover, T. and Hart, P. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory* 13, (1967), 21–27.
7. Duda, R. and Hart, P. *Pattern Classification and Scene Analysis*. Wiley, New York, 1973.
8. Kobsa, A. Personalized hypermedia and international privacy. *Commun. ACM* 45, 5 (May 2002).
9. Maes, P. Agents that reduce work and information overload. *Commun. ACM* 37, 7 (July 1994).
10. Pazzani, M. A Framework for collaborative, content-based and demographic filtering. *Artificial Intelligence Review* 13, 5/6 (1999), 393–408.
11. Pazzani M. and Billsus, D. Learning and revising user profiles: The identification of interesting Web sites. *Machine Learning* 27, (1997), 313–331.
12. Yang, Y. An evaluation of statistical approaches to text categorization. *Journal of Information Retrieval* 1, 1/2 (1999), 67–88.

DANIEL BILLSUS (billsus@adaptiveinfo.com) is Chief Technology Officer, **CLIFFORD BRUNK** (brunk@adaptiveinfo.com) is Vice President of Engineering, **CRAIG EVANS** (evans@adaptiveinfo.com) is a user-interface specialist, **BRIAN GLADISH** (gladish@adaptiveinfo.com) is a senior developer, and **MICHAEL PAZZANI** (pazzani@adaptiveinfo.com) is CEO of AdaptiveInfo in Irvine, CA. He is also a professor on leave from the University of California, Irvine.
