

Intelligent Content Discovery on the Mobile Internet

Experiences and Lessons Learned

Barry Smyth, Paul Cotter, and Stephen Oman

■ *The mobile Internet represents a massive opportunity for mobile operators and content providers. Today there are more than 2 billion mobile subscribers, with 3 billion predicted by the end of 2007. However, despite significant improvements in handsets, infrastructure, content, and charging models, mobile users are still struggling to access and locate relevant content and services. An important part of this so-called content-discovery problem relates to the navigation effort that users must invest in browsing and searching for mobile content. In this article we describe one successfully deployed solution, which uses personalization technology to profile subscriber interests in order to automatically adapt mobile portals to their learned preferences. We present summary results, from our deployment experiences with more than 40 mobile operators and millions of subscribers around the world, which demonstrate how this solution can have a significant impact on portal usability, subscriber usage, and mobile operator revenues.*

The mobile Internet represents a massive opportunity for mobile operators and content providers. With 3 billion subscribers predicted by the end of 2007,¹ the mobile market dwarfs the PC market. A recent study (Wright 2006) shows that 28 percent of mobile subscribers worldwide have browsed the Internet using their mobile phone. Moreover, recent growth is fueled by older users (age 35 and over), suggesting that the traditional early adopter segment (young males) no longer drives mobile Internet access. In terms of market size, the largest mobile carriers in the United States, Cingular Wireless, Sprint Nextel, T-Mobile, and Verizon Wireless, generated combined data revenues in excess of \$6.3 billion for the first half of 2006, putting them on track to realize a 75 percent increase over 2005 revenues (Reardon 2006). However, for mobile operators to capitalize fully on the mobile Internet opportunity they must help their subscribers to easily access relevant mobile content and services.

Despite major improvements in handsets, infrastructure, and billing, helping users (that is, mobile subscribers) to access mobile content remains a significant challenge; we refer to this as the content-discovery challenge. Most mobile users access content through their operator's mobile portal, navigating through portal menus to locate their favorite content and services. However, navigating through a complex portal structure, on a small-screen handheld device, can be very time consuming (Smyth 2002), lead-

Photo © Luis Pedrosa

ing to limited user satisfaction and low usage levels overall.

In this article we will describe how ChangingWorlds Ltd.² has used artificial intelligence technologies to help some of the world's leading mobile operators meet this content-discovery challenge. ChangingWorlds was founded in 1999 to bring advanced personalization technology to new content markets, including the mobile Internet market. We will focus on how the ClixSmart intelligent portal platform uses personalization technology to deliver mobile portals that are automatically adapted to the learned needs and preferences of individual mobile users. Today the platform has been deployed in more than 40 mobile operators around the world, personalizing tens of millions of portals for mobile users every day. Relevant technical details have been presented previously (Smyth and Cotter 2002a, 2002b), and in this article we will focus more on our deployment experiences and specifically on how the technology has had a dramatic effect on portal usability, driving usage and growing revenues for operators.

The Content-Discovery Challenge

Today the vast majority of mobile Internet content is accessed through the portals of mobile operators. For example, recent research (Church et al. 2007) has highlighted how more than 90 percent of mobile subscribers use their operator's portal as their primary source of content. Fewer than 10 percent of users avail themselves of search engines to locate off-portal content, despite the recent interest in mobile search. At the same time, mobile Internet usage has remained at relatively low levels with portal usability one of the most critical barriers affecting user satisfaction and usage. In particular, the protracted navigation times associated with mobile portals are cited as one of the main sources of user frustration.

Mobile portals are examples of hierarchical menu systems (HMSs) (Marsden and Jones 2001), and long before the arrival of the mobile Internet different forms of hierarchical menu systems were studied extensively with respect to their general usability and navigation characteristics (Larson and Czerwinski 1998, Miller 1981, Zaphirs 2000). Certainly the scale of the usability and navigation problems associated with mobile portals today, along with the mismatch between user expectations and realities, is highlighted by a number of recent studies (Chittaro and Cin 2002, Ramsey and Nielsen 2000). For instance, Chittaro and Cin (2002) examine two important Wireless Application Protocol (WAP) user-interface design choices (single-choice menu selections and navigation among cards) with respect to novice users. They

provide evidence that exploiting such navigation links and single-choice selections can significantly improve usability. However, another study highlights the pitfalls of too many navigation links and claims that while the average user expects to be able to access content within 30 seconds, the reality is closer to 150 seconds (Ramsey and Nielsen 2000).

Given that the majority of mobile content is accessed through the mobile portal, and the importance of the navigation problem in overall portal usability, we hypothesize that solving this navigation problem will have a major impact on end-user satisfaction and, ultimately, operator revenue. In ChangingWorlds we have tackled this problem head on by using artificial intelligence techniques to automatically learn subscriber content needs in order to automatically restructure portals on a user-by-user basis and so to offer users greatly reduced navigation times. To this end we have developed the so-called *click-distance* model of navigation effort (Smyth and Cotter 2002a, 2002b). This model measures the number of navigation steps required to locate a given content item from within the portal (typically from the portal home page). With the current generation of mobile phones, there are two basic types of navigation step. The first is the *menu select*: the user clicks to select a specific menu option. The second is a *menu scroll*: the user clicks to scroll up or down through a series of options. Accordingly, an item of content, i , within a mobile portal, can be uniquely positioned by the sequence of selects and scrolls needed to access it, and the navigation effort associated with this item can be modeled as click-distance, the number of these selects and scrolls (see equation 1).

$$ClickDistance(i) = Selects(i) + Scrolls(i) \quad (1)$$

Although this simple model of navigation effort equally weights the scrolls and selects, when we evaluate click-distance in comparison to navigation time by analyzing the behavior of live users on commercial mobile portals, we find a near-perfect correlation. Indeed we have found that the above model is easily adapted for the new generation of handsets, which accommodate other forms of navigation (such as sideways and diagonal navigation or touch-screen-based navigation).

Thus, large click-distances are indicative of protracted navigation times, and recent studies illustrate the extent of the click-distance problem. For example, an analysis of 20 European mobile portals reported an average click-distance in excess of 16 (Smyth 2002). In other words, a typical European mobile portal user can expect to have to make 16 or more clicks (scrolls and selects) to navigate from his or her portal home page to a typical content target. Moreover, on average, European portals are organized such that fewer than 30 per-

cent of content sites are within 10 to 12 clicks of the portal home page; 10 to 12 clicks corresponds to a navigation time of about 30 seconds (Smyth and Cotter 2002a, 2002b), which is expected by mobile portal users (Ramsey and Nielsen 2000). To put this another way, more than 70 percent of mobile portal content is essentially invisible to users because of its positioning within its parent portal.

Toward an Intelligent Mobile Portal Platform

The ClixSmart Intelligent Portal Platform is a carrier-grade mobile portal platform that provides operators with a complete multiaccess portal solution, combining portal management, content integration, device management, and business intelligence features to enable a superior online experience. In addition, the ClixSmart platform provides unique subscriber profiling and portal personalization features, and in the remainder of this article we will describe how these features have helped mobile operators around the world to significantly enhance their mobile portal offerings and grow subscriber usage.

“One-Size-Fits-All” Versus Personalization

Large click-distances are a fundamental feature of a one-size-fits-all approach to portal design, and the only sustainable solution to the usability problem this entails is to break with this tradition. Ultimately, portal click-distance can be greatly minimized by tailoring the portal for the needs of an individual user so that the content and services that are of interest to this user are near to the portal home page and thus accessible with a minimum number of clicks. Less relevant content and services can be relegated to the outskirts of the portal. Achieving this is no trivial task. It means developing a separate portal for each individual user, an unacceptably expensive task for the portal operator. Of course instead of forcing the operator to be responsible for individualizing the portal, the user could be provided with a facility that allows for manual customization. However, to date such initiatives, whereby users can manually reconfigure the portal according to their needs, have failed to attract users in sufficient numbers to be successful; even those users who do initially spend time customizing their portal rarely maintain it in line with their changing interests.

Fortunately, a practical solution is at hand that avoids the need for manual user-based customization or major operator expense by using artificial intelligence techniques to automatically optimize a portal for individual users. Recent research has made it possible to use user profiling and personalization techniques to learn about the preferences of individual users, and this information can be

used to adapt the structure of the portal on a user by user basis. For example, if a given user regularly accesses his or her local cinema’s listings then this content service can be made available from the portal home page (or at least nearby to the home page) rather than languishing deep within the portal structure. Thus, our strategy for decreasing navigation effort is to reduce the click-distance of the content items that a given user is likely to be interested in by *promoting* these items (or the links that lead to them) to higher positions within the portal menu structure. In general, personalization research seeks to develop techniques for learning and exploiting user preferences, to deliver the right content to the right user at the right time (see Bill-sus, Pazzani, and Chen [2000]; Fu, Budzik, and Hammond [2000]; Perkowski [2001]; Perkowski and Etzioni [2000]; Reiken [2000]; and Smyth and Cotter [2000]), and these ideas can be applied to the personalization of a portal structure to aid navigation effort (see also Anderson, Domingos, and Weld [2001] and Smyth and Cotter [2002b, 2002a]).

The core idea behind our personalized navigation technique is to use a probabilistic model of user navigation preferences to predict the likelihood that some portal/menu option o will be selected by a user u , given that the user is currently in menu m , and based on his or her past navigation history. We wish to compute $P_u(o|m)$, the access probability of o given m for user u , for all options o accessible from m (either directly or indirectly, through descendant menus). Put simply, when a user arrives at menu page m , we do not necessarily return its default options, o_1, \dots, o_n , which have essentially been hard-coded by the portal editor. Instead we compute the options, o'_1, \dots, o'_k , that are most likely to be accessed by the user from m ; that is, the k menu options, accessible from m , which have the highest access probabilities. This can mean promoting certain menu options into m , options that by default belong to descendants of m . The size of the final personalized menu is constrained by some maximum number of options, k , and the constituent options of m are ordered according to their access probabilities.

Subscriber Intelligence and Navigation Profiles

As users access a portal over time, they build up a navigation history, and this history can be very revealing with respect to their content preferences and information needs. For example, frequent accesses to the same content and services lead to well-traveled paths through the portal. By recording these access patterns (that is, by recording each sequence of menu options that the user accesses) it is possible to construct an accurate picture of an individual user’s navigation history (see also Herder

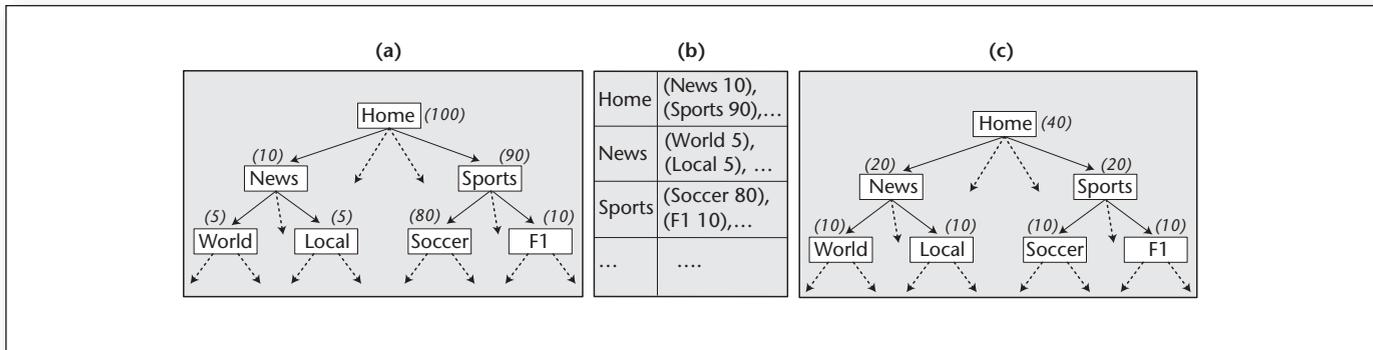


Figure 1. Menu Hierarchy, Hit Table Entries, and Menu Tree.

(a) The menu hierarchy. (b) Hit table entries corresponding to a sequence of visits by a given user. (c) Menu tree corresponding to the static, default portal structure.

[2003]) as the basis for a comprehensive user profile. The so-called hit table data structure is an efficient way of storing this information for a given user; see figure 1a and 1b for an example of a partial menu tree and corresponding hit table entries. A hit table can be thought of as a simple hash table, keyed according to the menu identifier and storing the number of accesses made by that user to options within that particular menu. For example, figures 1a and 1b reflect how one particular user has accessed the “news” section of his or her portal’s home page 10 times and the “sports” section 90 times over a series of sessions. The hit table entries can be used directly to compute the basic probabilities that a given menu option will be accessed within the portal.

In fact, there are two important types of hit table. The user hit tables reflect the access patterns for each individual user. In addition there is also a static hit table maintained to reflect the portal’s default structure. This static table makes it possible to deliver the standard (default) menu structure (as developed by the portal designer) early on, but this will eventually be overridden by the personalized portals as the access probabilities build. Moreover, the default hit values that are set in the static hit table make it possible to control the personalization latency—low static values mean that personalization takes effect very quickly, while large values make the system less sensitive to user activity. For example, figure 1c shows a sample default portal corresponding to a static hit table that gives equal weight to each of the menu options shown.

Building a Personalized Menu

The key then to personalizing the navigation structure of a mobile portal relies on an ability to reconstruct individual portal menus to reflect the navigation history of a given user. For example, if a user regularly navigates from the entertainment

menu of a portal, through a series of submenus, in order to access his or her local cinema listings, then perhaps this local cinema option should be promoted to the entertainment menu.

All personalization is performed in real time on receipt of each user request for a particular portal page. The basic process model is presented as figure 2 and includes the following sequence of 9 steps:

- (1) The user requests a menu page from his or her mobile handset.
- (2) The request is forwarded by the WAP Gateway with the user’s unique ID (MSISDN number) to the Device Manager, which ultimately optimizes the content according to the features of the target handset.
- (3) The Device Manager recognizes the device type and then forwards the request to the Navigator Server.
- (4) The Navigator Server examines the portal and requests the default menu content.
- (5) The Navigator Server examines the user profile database and requests the user’s current profile if it has not already been downloaded.
- (6) The Navigator Server is responsible for the portal personalization and combines the static portal with the user’s profile in order to construct the personalized portal menu by re-ordering or promoting content links.
- (7) The Device Manager reads the device style sheet for the user’s device.
- (8) The Device Manager formats the personalized menu for the appropriate device and sends the response to the WAP Gateway.
- (9) The WAP Gateway forwards the personalized page to the user.

Obviously step 6 is the critical part of the process from a portal personalization standpoint: it is here that the personalized version of the particular menu, *m*, is generated. To perform this step, the Navigator Server component must determine how the default options of *m* should be ordered and whether any of the menu options that appear below *m* merit promotion. Since menu size is usually limited by portal style guides, a means of ordering

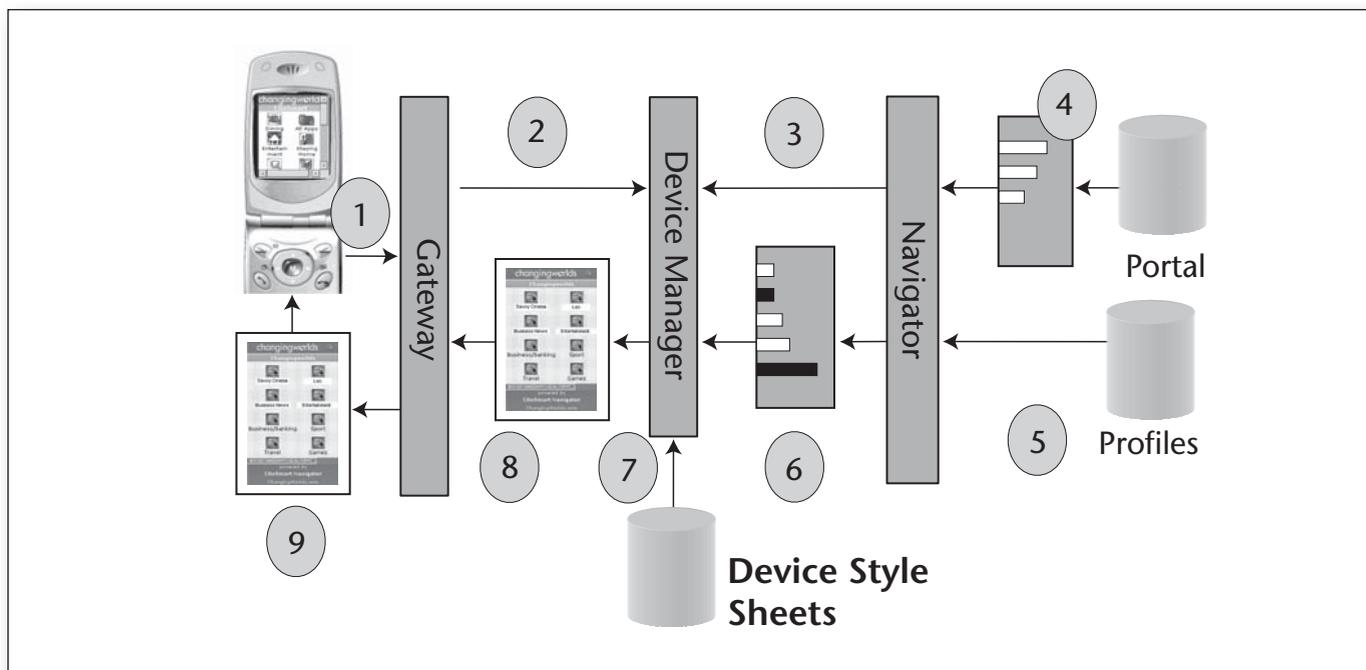


Figure 2. A Typical Integration of ClixSmart Navigator (and Device Manager) with a Mobile Operator's Gateway.

eligible options is required. One solution is to compute the k most probable options from m ; that is the k options with the highest $P_u(o|m)$. Thus, the k options that are most likely to be accessed, given that the user is currently accessing menu m , are added to m . To do this we take account of the hit values listed for each option in both the static and user hit tables by using the recorded access frequencies as a way to estimate the necessary access probabilities. For example, given the data shown in figure 1, $P_u(\text{News}|\text{Home})$ is calculated as the combined relative frequency accesses, taking genuine user accesses and default static hit values into account. Thus $P_u(\text{News}|\text{Home}) = (20 + 10)/(40 + 100) = 0.214$. Similarly, $P_u(\text{World}|\text{Home})$ is calculated by combining chaining access probabilities so that $P_u(\text{News}|\text{Home}) P_u(\text{World}|\text{News}) = (20 + 10)/(40 + 100) (5 + 10)/(10 + 20) = 0.107$. In this way we can calculate the access probabilities for all of the menu options that are accessible from m (in this case m is the portal home page). For the current example, in descending order of access probability (or desirability) we have *Sports*, *Soccer*, *News*, *FL*, *World*, and *Local*. And for $k = 3$, *Sports*, *Soccer*, and *News* are selected, in order, for addition to the requested *Home* menu.

Content Promotions in Practice

This personalized navigation method supports two basic types of menu adaptations. Firstly an option may be reordered within its default menu. That is, the relative position of an option within a parent

menu may be changed so that options are ordered in descending order of their access probabilities; this reduces the amount of scrolling needed during navigation. Alternatively, if there is sufficient evidence, a menu option may be promoted from its default menu to some higher-level menu. For instance, in the worked example above, the *Soccer* menu is promoted from its default location within the *Sports* menu to the portal home page. Thus promotion is the second type of menu adaptation and influences click-distance by reducing the number of menu selects needed to access a content item.

By way of an example, figure 3 presents a series of portal pages leading the user to his or her local cinema listings (Ster Century) through a number of intermediate menu options. Assuming that this becomes a well traveled path for the user in question, we can expect the portal to promote the Ster Century service to a more prominent position in the portal for that user. An example promotion scenario is presented in figure 3b to illustrate this. The Ster Century service has been promoted to the top position within the *Entertainment* menu, reducing its click-distance significantly by eliminating a number of intermediate portal levels. In addition, the *Entertainment* menu within the portal home page has been promoted from position 5 to position 1.

In this way, menu reorderings and promotions (and conversely demotions) are side effects of the access probability calculations and provide a fluid personalization scheme that gracefully adapts the

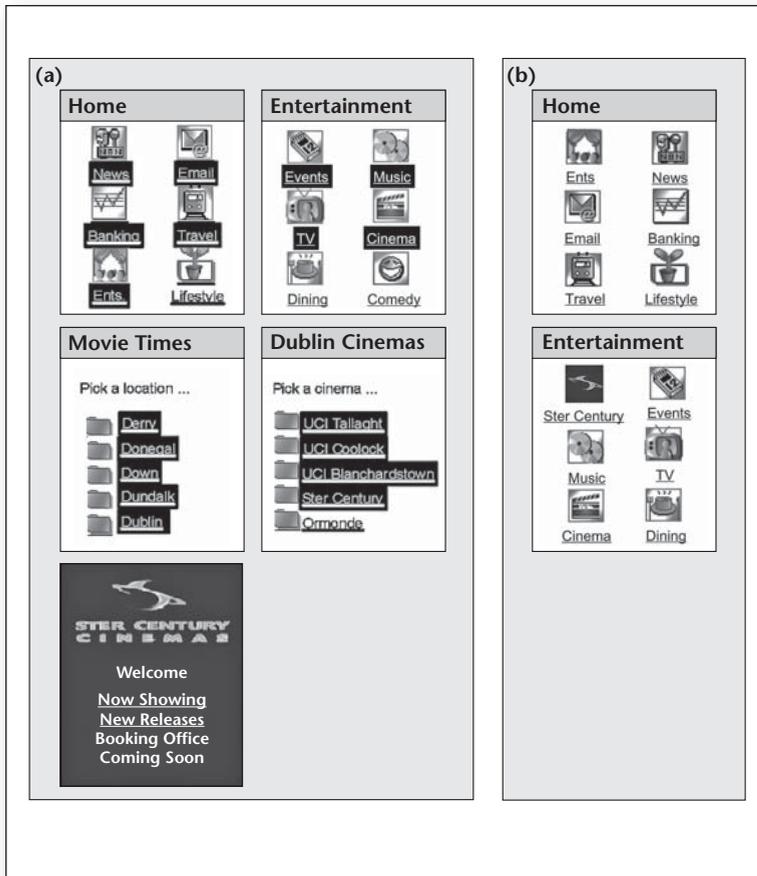


Figure 3. A Series of Portal Pages.

(a) An example sequence of navigation steps through a series of portal pages leading users to their local cinema listings (Ster Century) through a number of intermediate menu options. (b) An example personalization scenario is presented in which the Ster Century service has been promoted to the top position within the Entertainment menu. In addition, the Entertainment menu has been promoted within the portal home page from position 5 to position 1.

navigation structure of a portal in response to a user's access patterns. The examples here have been kept simple for reasons of clarity, focusing on the promotion of single items, for example. Of course in reality there may be a number of content services competing for a limited number of promotion slots. In theory options can be promoted from anywhere deep within the portal structure once their probabilities build sufficiently, although in practice certain limits may be necessary to control the speed and scope of personalization as we will discuss briefly.

Efficiently Computing Promotions

The efficiency of the proposed personalization method depends on the complexity of the process that identifies the k most probable options for the

menu, m . This can mean examining not just the options of m but also all the options contained in menus that are descendants of m . This leads to an exponential growth in the number of menu options that must be considered as more and more portal levels are considered during personalization, which is clearly not practical in order to deliver real-time personalization. One option is to constrain this growth by limiting the number of levels to look ahead during personalization; for example, promotion candidates for m may only be drawn from menus that are up to say two levels deep from m . Of course this limits promotion opportunities as there may well be compelling promotion candidates that exist three or four levels deep within the portal.

Fortunately, a more efficient algorithm is possible once we recognize that, by definition, $P_u(o|m)$ is always greater than or equal to $P_u(o'|m)$ where o' is an option of a menu, m , with m' itself being a descendant of m through o . This means that we can find the k most probable nodes for menu m by performing a depth-limited, breadth-first search over the menu tree rooted at m . And we only need to expand the search through an option o' if $P_u(o'|m)$ is greater than the k^{th} best probability so far found; (see Smyth and Cotter [2002b] for further technical details). In practice this modification can result in significant reductions in search effort allowing probabilities to be computed on the fly without introducing any significant delays in content delivery.

Application Use and Payoff

During the course of the past five years ChangingWorlds has had the opportunity to work closely with the world's leading mobile operators in order to deploy and evaluate the ClixSmart Intelligent Portal Platform in a variety of comprehensive evaluation scenarios. In particular, this has included an opportunity to evaluate the impact of personalization during a number of long-term field trials involving live users and separate control groups. In this section we will describe a typical trial scenario and a subset of results as they relate to portal personalization, focusing on the ability of ClixSmart personalization to improve portal usability by reducing navigation times and so increase portal usage. In addition we will highlight a number of more recent results from a tier 1 operator that relate to the ultimate value attributed to portal personalization.

Live User Trials

An important challenge for ChangingWorlds over the past five years has been the need to educate operators about the value of the personalization opportunity. To begin with, the concept of automatic portal personalization (that a mobile portal can automatically adapt its structure better to better re-

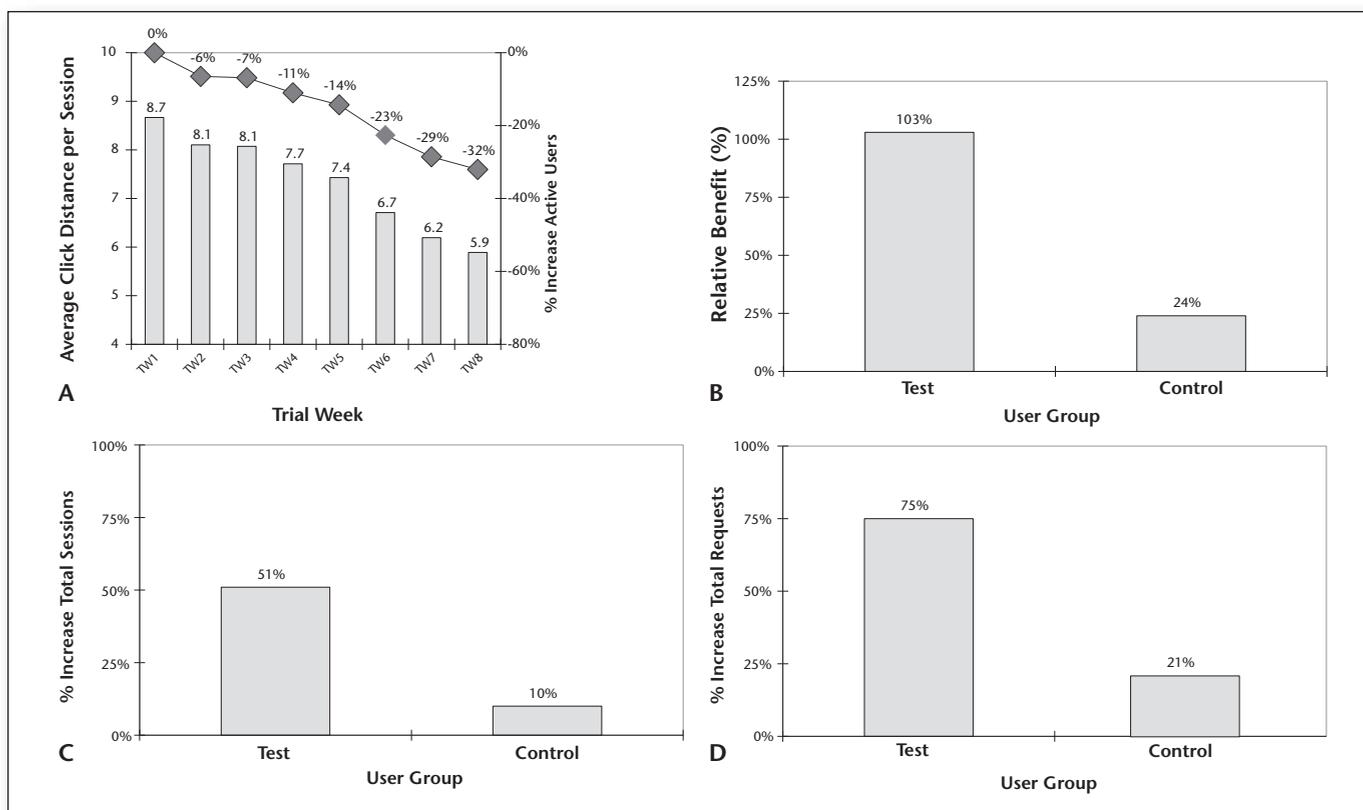


Figure 4. Key Evaluation Metrics.

(a) Average session click-distance for test users during each of the eight trial weeks. (b) Average comparative increase in the number of users accessing the portal on a weekly basis. (c) Average comparative increase in the number of user sessions. (d) Average comparative increase in the number of page requests.

flect the usage patterns and likely interests of an individual subscriber) was somewhat alien to the “one-size-fits-all” mindset of mobile operators. Indeed many operators initially expressed concerns that personalizing portals may serve to confuse subscribers, especially if their favorite sites and services begin to change position from one session to the next. As already mentioned, the challenge of effective personalization is to ensure that portal changes are made in a way that does not confuse or disorientate subscribers, and in this regard a number of features are included to safeguard against such problems. For example, operators can control the sensitivity of individual menu options with respect to user activity, thereby facilitating different rates of promotion for different types of content. In turn, operators can control the scope of personalization, such as whether an option can be promoted beyond its parent menu and, if so, how far such a promoted option can travel through the portal hierarchy. These controls provide a very effective mechanism for ensuring that users do not become disoriented in practice. It also allows oper-

ators to influence the promotion of content within the portal in line with their business needs and forms the basis for a range of targeted marketing features that are beyond the scope of this article.

Satisfying mobile operators about the ability of portal personalization to deliver a more efficient and effective mobile portal experience was achieved by using a comprehensive trial methodology in order to evaluate the behavior of a group of test users (with access to a personalized version of the portal) relative to a control group (with access to the regular portal) for a sustained period of time, typically two or more months. The results presented in this section are taken from one such trial with a major European operator. For the purpose of the trial, a mirror of the standard operator portal was managed by the ClixSmart platform, offering portal personalization to a group of almost 900 test users who were selected at random. The usage patterns of these test users were tracked during an eight-week trial period and compared to the usage of the remaining subscriber base, which served as a control group.

The usage results are presented in figures 4a–d and show a very significant benefit associated with the activity levels of the test group relative to the control. For example, in figure 4a we see how the test group enjoys a gradual decline in its average session click-distance (the number of interactions needed to access an item of content within a given session) over the trial period. To begin with, the typical user required an average of 8.7 clicks to access content, but by the end of the trial this had dropped by more than 30 percent to 5.9; indeed our studies indicate that on average click-distance will fall by about 50 percent over a three-month period.

Figures 4b–d highlight certain key activity indicators of particular interest to mobile operators. In this instance we have graphed the average increase in activity for each group of users during the eight-week trial compared to the previous eight-week pretrial period. For example, in figure 4b we compare the change in the average number of users accessing the portal during a typical week. The results show that the test group using the personalized portal increased their activity levels by more than 100 percent over the eight-week pretrial period during the eight-week trial period. During the eight-week pretrial period, an average of just over 140 of the test users accessed the portal on a weekly basis, and this rose to just under 290 users during the trial period. In contrast, during the same period of time the average increase in the activity of the control group increased by only 24 percent.

Similar increases are seen across other key metrics such as number of sessions (figure 4c) and total requests (figure 4d). For instance, we see a 54 percent (75 percent – 21 percent) relative increase in the average total weekly requests generated by the test users compared to the control group. In other words, despite the fact that click-distance was falling for the test users during the trial (so they were generating fewer navigation requests), these users were generating an increased proportion of content requests. It is worth noting that at the time of this trial the operator in question employed a request-based charging model, whereby users were charged on the basis of requests. Therefore this benefit can be linked directly to an expected uplift in revenue for the operator; directly after the trial, the operator in question rapidly deployed the ClixSmart solution across its entire subscriber base.

The results presented relate to one particular trial for one particular operator. Similar advantages have been observed across a range of comparable trials, and averaging across these different trials we find the following high-level benefits as a direct result of portal personalization: (1) Click-distance falls by an average of 50 percent within three months. (2) The average number of user sessions

increases by 30 percent. (3) The average number of requests per session increases by 35 percent. (4) The average number of failed sessions (sessions where users fail to locate content) falls by approximately 50 percent.

The Value of Personalization

Ultimately operators are focused on the ability of this technology to drive revenue, and a key question relates to the long-term ability of portal personalization to deliver a sustained increase in usage across a subscriber population so as to deliver a significant increase in mobile Internet revenue. To answer this question a third-party independent analysis of the benefits of ClixSmart portal personalization was recently commissioned by a leading mobile operator group. The aim was to assess the business impact of portal personalization some two years postdeployment. The results of this analysis clarify how personalization has had a significant and measurable impact on portal usability, usage, and operator revenue. Specifically, the study drew the following conclusions:

First, there was a significant increase in portal usage across all major content areas; for example, browsing was found to have increased by 28 percent, ring-tone downloads were up by 14 percent, video content downloads were increased by almost 30 percent, and gaming downloads were increased by more than 35 percent.

Second, personalization had a positive impact on customer satisfaction and reduced the likelihood of subscriber churn (the likelihood that a user will switch to an alternative operator); for example, 31 percent of users reported an increase in their satisfaction levels, and more than 30 percent of users indicated that they would be less likely to churn.

Third, the direct revenue impact of portal personalization over a 12-month period for the commissioning operator was found to be in excess of \$15 million per annum. That is, the analysis concluded that personalization was directly responsible for delivering an increase in usage across the subscriber base worth \$15 million in direct revenue.

Application Development and Deployment

ClixSmart is a Java- and XML-based carrier-grade mobile portal platform designed to deal with the demands of the world's largest operators. The platform provides for full horizontal and vertical scalability and is commercially deployed with service-level agreements (SLAs) requiring 99.7 percent scheduled uptime (or better). Moreover, the platform is designed to offer real-time personalization without negatively affecting overall portal per-

formance. For example, a single ClixSmart Navigator Server, running on a 4xCPU, SunFire 480R with 16 GB of RAM is capable of handling 400 requests for personalized pages per second.

In all deployments the ClixSmart software has been successfully integrated (in loosely or tightly coupled configurations) within the complex service delivery infrastructure that makes up a modern mobile operator's data service layer. The deployment option that is chosen depends upon a combination of the specific project requirements and the nature of the overall portal solution, including the availability of existing portal components and device management functionality.

Lessons Learned

A number of valuable lessons have been learned during the course of the development and deployment of ClixSmart. These may be of interest to others when it comes to future deployments of AI and related technologies, and so in this section we review a number of these lessons.

One traditional approach to the development of AI technologies attempts to produce what might be termed a "black-box style" AI solution, one in which the various AI techniques are largely hidden from the end user with a view to providing a fully automated solution. Indeed this was the approach taken early on in our own work, largely on the assumption that mobile operators would prefer a fully automated personalized portal, which would adapt to the changing preferences of users without any need for explicit intervention. This assumption turned out to be at least partially flawed. In truth, while a high degree of automation was desirable, it was equally important for mobile operators to retain a level of control and to be afforded the opportunity to adapt and adjust portal personalization in line with their business needs and their view of their subscribers' preferences. As a result a significant degree of effort has been invested in developing a personalization engine that can be easily administered by portal managers and that provides portal managers with a way to control and influence portal personalization. The resulting ClixSmart portal management interface provides operators with the ability to control high-level personalization features (such as whether certain sections of the portal are available to personalization) as well as lower-level details, such as the sensitivity of certain content items to user interactions. In this way mobile operators can develop personalized portals that are designed to take best advantage of their subscriber demographics and content portfolios, while at the same time addressing their business needs.

Patrick Winston (Massachusetts Institute of Technology) is often reported to have referred to practical AI systems as "raisin-bread systems,"

highlighting how, in practice, AI technology tends to become embedded in larger systems and mixed with conventional technologies. The AI technology does not have to carry much volume, but it may carry a significant portion of the value; without the raisins, raisin bread is just, well it's just bread. What is important is how we come to blend AI technologies with more traditional technologies to provide useful new services and systems that would otherwise not be possible. We have observed a similar phenomenon at work during the development of the ClixSmart platform as an increasing amount of development work is devoted to more traditional technologies. Nevertheless personalization has remained at the core of everything we do, and AI technologies continue to play a vital role in the ClixSmart personalization engine. In our experience, AI technology has proven to be a key differentiator for the company and its products within the mobile marketplace.

Finally, it is worth highlighting perhaps the most important lesson learned during the early years of commercializing this technology. In the beginning there was always a high degree of optimism that the AI technology would, in a sense, largely sell itself: that mobile operators would be eager to be among the first to offer their subscribers a fully personalized portal, for example. Of course it soon became clear that the reality was considerably different, as mobile operators had learned hard lessons about investing in new technology for technology's sake during the late 1990s. And so while there was a significant interest in personalization technology, it was vital to be able to present proven business benefits to operators. Ultimately a qualified evaluation of the technology's return on investment (ROI) proved to be a much more important sales tool than any precision or recall studies that would normally have been presented. As such, significant effort was invested in the evaluation of the ClixSmart platform with early reference customers to provide robust ROI arguments and detailed user experience studies to demonstrate the true value of the technology to mobile operators and their subscribers. These large-scale studies involved real mobile subscribers and were conducted over extended periods of time in genuine usage scenarios. Studies such as the one presented earlier in this article provided the raw data necessary to develop the personalization business model that operators required.

These lessons, and the market experience gained as a result, have helped ChangingWorlds to significantly enhance the value of its ClixSmart platform, which today is recognized as the class leader in the mobile sector. As mentioned previously, the platform has been deployed in more than 40 mobile operators around the world, ranging from tier 1 operators (with in excess of 10 million sub-

scribers) to tier 2 (5–10 million subscribers) and tier 3 operators (with fewer than 5 million subscribers); for example, it currently handles over 50 percent of the United Kingdom’s mobile Internet traffic.

Beyond Portal Navigation: From Content Recommendation to Mobile Advertising

So far we have outlined the use of AI techniques to automatically adapt the navigation structure of a mobile Internet portal to reduce the effort that a subscriber has to invest in order to reach content. Of course portal navigation is just one mode of information access, and the availability of user preference data in the form of subscriber profiles provides many additional opportunities for the provision of highly personalized services. In this section we highlight how complementary personalization techniques can be used to inform other modes of content discovery including content recommendation and mobile advertising, both of which involve proactive approaches for selecting and suggesting content items to mobile users.

Personalized Content Recommendation

Personalizing the navigation structure of a mobile portal is a powerful way to help users to navigate to familiar parts of an information space more efficiently. But oftentimes the goal is to help users to discover new types of information or content items that they have not yet had an opportunity to experience. Recommender systems provide this capability by selecting new items to recommend to users based on their learned preferences. For example, recommender systems such as Amazon’s use purchase history information to suggest new items to some target user based on the purchases of other similar users to the target. Amazon’s ubiquitous “people who bought *X* also bought *Y*” recommendations have proved to be a remarkably effective cross-selling tool and provide a strong case for one particular brand of recommendation technology, *automated collaborative filtering* (ACF), which is based on the ratings patterns of users rather than any metadata descriptions of the items being recommended (Schafer et al. 2006). While ACF techniques can be effective among large populations of active users they are less so when dealing with new users or immature profiles where ratings scarcity can limit the accuracy and coverage of recommendations. In such circumstances, other types of recommendation strategies must be used (for example, *content-based recommendation* can be used to suggest items to target users that are *similar* to items they have preferred in the past (Smyth 2006), but only if item metadata is available as the basis for similarity), but these too are limited in many practical circumstances.

In general, no single recommendation strategy is appropriate in all circumstances, and as a result the ClixSmart Recommendation Framework has been developed to accommodate a range of different strategies to provide access to a range of complementary core recommendation strategies including collaborative filtering and content-based and rule-based techniques, for example. Briefly, the framework incorporates two main functional components:

Component One—The Candidate Generator. This component is responsible for generating a broad set of recommendation candidates for a given target user. These candidates are generated based on an ensemble of different core recommendation techniques as appropriate to the recommendation domain. For example, in a music downloads domain, where there is a history of past user downloads and a reasonably rich set of metadata associated with music genres, artists, and so on, then a combination of collaborative filtering and content-based techniques can be used.

Component Two—The Candidate Selector. This component is responsible for selecting a short-list of candidate recommendations for target users based on their current context. Making such contextual recommendations is essential if recommendations are to “make sense” to users, and a number of techniques can be used to filter candidates according to the user’s context. For example, if the user is currently viewing a Bruce Springsteen ring tone then candidate ring tone recommendations that have similar artist or genre features may be suggested.

Figure 5 shows a simple example of how recommendations can be contextualized based on the user’s current portal activities. In figure 5a, for example, the target user is viewing a high-level portal page that provides links to ring tones, games, and wallpaper (images) content, and as a result, recommendations have been selected to cover these three different types of content. In this case, the target user’s profile suggests recommendations for music by Britney Spears and the Killers, images of Britney Spears as device wallpaper, and a couple of mobile games (Football Manager and Tetris). In figure 5b, the same user is presented with a different set of recommendations. In this case the user is currently considering a mobile version of Pacman and so is presented with a set of game-focused recommendations (Space Invaders, Football Manager, and Tetris). This second set of recommendations has been generated by selecting candidates that are related to the current Pacman item.

Personalized Mobile Advertising

Perhaps one of the most important recent developments in the mobile data space is the recognition of some of the unique features of the mobile

space from an advertising point of view. Mobile devices provide advertisers with a range of different advertising channels, from Short Message Service (SMS) and Multimedia Messaging Service (MMS) messaging to mobile Internet and customer-care interactions. In addition, the subscriber intelligence that is available to many mobile operators enables a far more precise targeting of advertisements than has previously been possible. For example, in addition to coarse-grained targeting information such as consumer demographics, time of day, and so on, mobile operators now have access to rich preference data based on the behavior of mobile users online. In combination, this information can be harnessed to ensure that advertising messages are more precisely targeted to the right individuals at the right time thereby ensuring a greater recognition and click-through rate and an improved experience for mobile subscribers.

Recently ChangingWorlds has used its recommendation framework as the basis for the development of the ClixSmart AdPersonalizer solution, which integrates with the existing advertising value chain and provides community-based targeting of mobile advertising content across a variety of channels. Initial live-user trial results point to high click-through rates of approximately 15 percent for targeted mobile advertisements, as compared to fractions of a percent for more traditional online advertisements. Certainly, if mobile advertising is to succeed then operators need to ensure that the limited bandwidth available for advertisements is used wisely as a way to enhance the mobile experience, and there is little doubt that the spamming of mobile subscribers with “quick-buck” advertising content will drive many away from the mobile channel. Once again, understanding and catering for the needs of users will play a key part in any successful mobile advertising proposition.

Conclusions

In this article we have focused on how personalization technology can be used to automatically adapt the structure of a mobile portal in line with learned subscriber preferences (intelligent navigation). We have demonstrated the significant benefits of this solution based on our experiences with more than 40 of the world’s leading mobile network operators. This application of artificial intelligence technology in the mobile sector has had a significant positive impact on operator revenues and subscriber satisfaction. Every day millions of mobile subscribers enjoy an enhanced mobile experience as a result of our personalization technology, and every day these subscribers consume more content than users of traditional one-size-fits-all mobile portals.

As an added-value benefit of this technology,

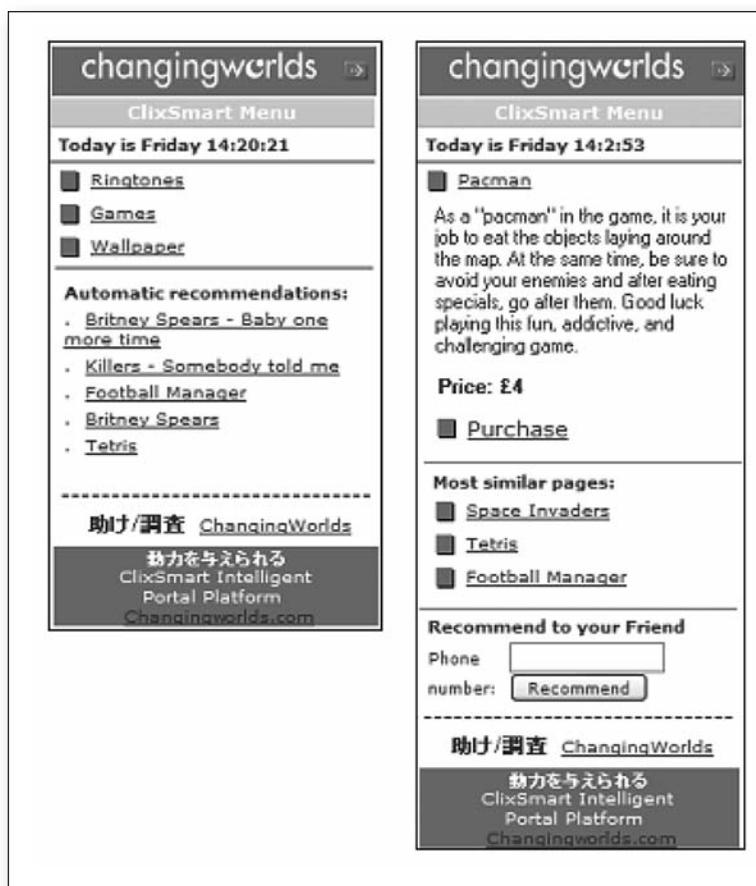


Figure 5. Sample Recommendations for a Target User.

(a) A set of general recommendations that cover a range of content types including ring tones, games, and images. (b) A more focused set of recommendations based on the user’s current gaming context.

mobile operators have also captured unique and valuable business intelligence about the content preferences of their subscribers. Personalizing portal navigation is a vital first step on the road to mobile content discovery, and the availability of this rich repository of subscriber intelligence will pave the way for a new generation of personalized information services, from smarter mobile search to targeted mobile advertising, helping to ensure that mobile subscribers continue to enjoy content and services that are relevant to their true needs.

Notes

1. See *Cellular News*: 2.5 Billion Mobile Phones in Use (www.cellular-news.com/story/19223.php); last checked 18 September 2006.
2. www.changingworlds.com.

References

- Anderson, C.; Domingos, P.; and Weld, D. 2001. Adaptive Web Navigation for Wireless Devices. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, 879–884. San Francisco: Morgan Kaufmann Publishers.
- Billsus, D.; Pazzani, M.; and Chen, J. 2000. A Learning Agent for Wireless News Access. In *Proceedings of Conference on Intelligent User Interfaces*, 33–36. New York: Association for Computing Machinery.
- Chittaro, L., and Cin, P. D. 2002. Evaluating Interface Design Choices on WAP Phones: Navigation and Selection. *Personal and Ubiquitous Computing* 6(4): 237–244.
- Church, K.; Smyth, B.; Cotter, P.; and Bradley, K. 2007. Mobile Information Access: A Study of Emerging Search Behavior on the Mobile Internet. *ACM Transactions on the Web* 1(1).
- Fu, X.; Budzik, J.; and Hammond, K. 2000. Mining Navigation History for Recommendation. In *Proceedings of Conference on Intelligent User Interfaces*, 106–112. New York: Association for Computing Machinery.
- Herder, E. 2003. Modeling User Navigation. In *User Modeling 2003: 9th International Conference, UM 2003*, ed. P. Brusilovsky, A. Corbett, and F. de Rosis, 417–419. Berlin: Springer.
- Larson, K., and Czerwinski, M. 1998. Web Page Design: Implications of Memory, Structure, and Scent for Information Retrieval. In *Proceedings of the CHI'98 Human Factors in Computer Systems*, 25–32. New York: ACM Press.
- Marsden, G., and Jones, M. 2001. Ubiquitous Computing and Cellular Handset Interfaces: Are Menus the Best Way Forward? Paper presented at the South African Institute of Computer Scientists and Information Technologists Annual Conference (SAICSIT), 25–28 September, Pretoria, South Africa.
- Miller, D. 1981. The Depth/Breadth Trade-off in Hierarchical Computer Menus. Paper presented at the 25th Annual Meeting of the Human Factors and Ergonomics Society, October, Rochester, NY.
- Perkowitz, M. 2001. Adaptive Web Sites: Cluster Mining and Conceptual Clustering for Index Page Synthesis. Ph.D. Dissertation, Department of Computer Science and Engineering, University of Washington, Seattle, WA.
- Perkowitz, M., and Etzioni, O. 2000. Towards Adaptive Web Sites: Conceptual Framework and Case Study. *Artificial Intelligence Journal* 18(1–2): 245–275.
- Ramsey, M., and Nielsen, J. 2000. *The WAP Usability Report*. Fremont, CA: Nielsen Norman Group.
- Reardon, Marguerite. 2006. *The Mobile Internet: Are We There Yet?* San Francisco: CNET Networks, Inc. (www.news.com/2100-1039_3-6110100.html); last checked 19 December 2007.
- Reiken, D., ed. 2000. Special Issue on Personalization. *Communications of the ACM* 43(8).
- Schafer, J. B.; Frankowski, D.; Herlocker, J.; and Sen, S. 2006. Collaborative Filtering Recommender Systems. In *The Adaptive Web: Methods and Strategies of Web Personalization*, Lecture Notes in Computer Science, Vol. 4321, ed. P. Brusilovsky, A. Kobsa, and W. Nejdl, 291–324. Berlin: Springer Verlag.
- Smyth, B. 2002. *The Plight of the Mobile Navigator*. Malmö, Sweden: MobileMetrix.
- Smyth, B. 2006. Case-Based Recommendation. In *The Adaptive Web: Methods and Strategies of Web Personalization*, Lecture Notes in Computer Science, Vol. 4321, ed. P. Brusilovsky, A. Kobsa, and W. Nejdl, 342–372. Berlin: SpringerVerlag.
- Smyth, B., and Cotter, C. 2000. Wapping the Web: A Case-Study in Content Personalization for WAP-Enabled Devices. In *Adaptive Hypermedia and Adaptive Web-Based Systems: International Conference, AH 2000*, Lecture Notes in Computer Science 1892, ed. P. Brusilovsky, O. Stock, and C. Strapparava, 98–108. Berlin: Springer.
- Smyth, B., and Cotter, P. 2002a. Personalized Adaptive Navigation for Mobile Portals. In *Proceedings of the 15th European Conference on Artificial Intelligence—Prestigious Applications of Artificial Intelligence*. Amsterdam, The Netherlands: IOS Press.
- Smyth, B., and Cotter, P. 2002b. The Plight of the Navigator: Solving the Navigation Problem for Wireless Portals. In *Adaptive Hypermedia and Adaptive Web-Based Systems: International Conference, AH 2002*, Lecture Notes in Computer Science 2347, 328–337. Springer-Verlag.
- Wright, Adam. 2006. Mobile Phones Could Soon Rival the PC as World's Dominant Internet Platform. News Release. Paris: Ipsos Corporation (www.ipsos-na.com/news/pressrelease.cfm?id=3049); last checked 19 December 2007.
- Zaphirs, P. 2000. Depth Versus Breadth in the Arrangement of Web Links. In *Proceedings of 44th Annual Meeting of the Human Factors and Ergonomics Society*, 139–144. Santa Monica, CA: Human Factors and Ergonomics Society.



Barry Smyth holds the digital chair of computer science in the School of Computer Science and Informatics at University College Dublin. He has a Ph.D. from Trinity College Dublin and his research interests include personalization, recommender systems, case-based reasoning, machine learning, and information retrieval. Smyth is also a cofounder of ChangingWorlds Ltd., a leading provider of mobile content-discovery solutions.



Paul Cotter holds a B.Sc. in computer science from University College Dublin. He is a cofounder of ChangingWorlds Ltd. where he currently serves as the chief technology officer. Cotter has more than 10 years' experience in the design and development of personalization technology for the mobile sector.



Stephen Oman holds a B.Sc. in computer science and an M.Sc. in artificial intelligence from Trinity College Dublin. He is a program director with ChangingWorlds with overall responsibility to bring ChangingWorlds's personalization technologies to mobile operators. He has successfully delivered AI-based personalization solutions to some of the world's largest mobile operators worldwide.