Enabling Intelligent Content Discovery on the Mobile Internet

Barry Smyth and **Paul Cotter** and **Stephen Oman**
ChangingWorlds Ltd.
South County Business Park, Leopardstown, Dublin, Ireland.
{firstname.lastname}@changingworlds.com
Tel. +353-1-4359800, Fax. +353-1-4359810

Abstract

The mobile Internet is a massive opportunity for mobile operators and content providers, but despite significant improvements in handsets, infrastructure, content, and charging models, mobile users are still struggling to access and locate relevant content and services. The core of this so-called *content discovery* problem is the navigation effort that users must invest in browsing and searching for mobile content. In this paper 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.

Introduction

The mobile Internet represents a massive opportunity for mobile operators and content provides. With 3 billion subscribers predicted by the end of 2007 (Cellular-News 2006), the mobile market dwarfs the PC market. A recent study (IposInsight 2006) shows that 28% of mobile subscribers worldwide have browsed the Internet using their mobile phone. Moreover, recent growth is fueled by older users (age 35+), 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 US, 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 realise a 75% increase over 2005 revenues (ZDNet 2006). However, for mobile operators to capitalise 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 via their operator's mobile portal, navigating through portal menus to locate their favourite content and services. However, navigating through a complex portal structure, on a small-screen handheld device, can be very time consuming (Smyth 2002), leading to limited user satisfaction and low usage levels overall.

In this paper we will describe how ChangingWorlds Ltd. (*www.changingworlds.com*) 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 $ClixSmart^{TM}$ 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 & Cotter 2002a; 2002b) and in this paper 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 via the portals of mobile operators. For example, recent research (Church *et al.* 2007) has highlighted how more than 90% of mobile subscribers use their operator's portal as their primary source of content. Less than 10% of users avail 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 most critical barriers impacting 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 (HMS) (Marsden & 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 &

Czerwinski 1998; Miller 1981; Zaphirs 2000). Certainly the scale of the usability and navigation problems associated with mobile portals today, and the mismatch between user expectations and realities, is highlighted by a number of recent studies (Chittaro & Cin 2002; Ramsey & Nielsen 2000). For instance, (Chittaro & Cin 2002) examine two important 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 & Nielsen 2000).

Given that the majority of mobile content is accessed via 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 offer users greatly reduced navigation times. To this end we have developed the so-called *click-distance* model of navigation effort (Smyth & Cotter 2002a; 2002b). This model measures the number of 'navigation steps' required to locate a given content item from a 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 modelled as click-distance, the number of these selects and scrolls (see Equation 1).

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

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 behaviour 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 (e.g., 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 their portal home page to a typical content target. Moreover, on average, European portals are organised such that less than 30% of content sites are within 10-12 clicks of the portal home page; 10-12 clicks corresponds to a navigation time of about 30 seconds (Smyth & Cotter 2002a; 2002b), which is expected by mobile portal users (Ramsey & Nielsen 2000). To put this another way, more than 70% of mobile portal content is essentially invisible to users because of its positioning within its parent portal.

An Intelligent Mobile Portal Platform

The ClixSmart Intelligent Portal Platform is a carrier-grade mobile portal platform that provides operators with a complete multi-access 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 paper 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" vs. 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 minimised by tailoring the portal for the needs of an individual user, so that 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 the user 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. It 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 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 with 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 po-

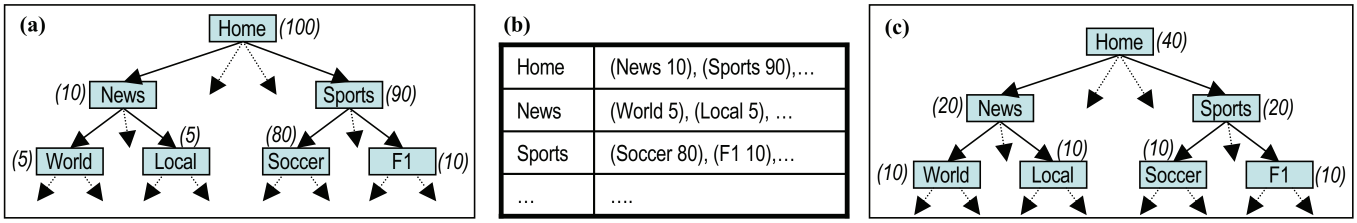

Figure 1: The menu hierarchy (a) and hit table entries (b) corresponding to a sequence of visits by a given user. (c) menu tree corresponding to the static, default portal structure.

sitions 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 (Billsus, Pazzani, & Chen 2000; Fu, Budzik, & Hammond 2000; Perkowitz 2001; Perkowitz & Etzioni 2000; Reiken 2000; Smyth & Cotter 2000) — and these ideas can be applied to the personalization of a portal structure to aid navigation effort; see also (Anderson, Domingos, & Weld 2001; Smyth & 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 they are currently in menu m, and based on their 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,..., o_n, which have essentially been hard-coded by the portal editor. Instead we compute the options, o'_1,...,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 descendans 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 & 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-travelled paths through the portal. By recording these access patterns—that is, by recording each sequence of menu options that are accessed—it is possible to construct an accurate picture of an individual user's navigation history (see also (Herder 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 1(a & b) 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, Figure 1(a & b) reflects how one particular user has accessed the *News* section of their 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 over-ridden 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 1(c) 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 sub-menus, in order to access their local cinema listings, then perhaps this local cinema option should be promoted 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 steps:

1. The user requests a menu page from their mobile handset.
2. The request is forwarded by the WAP Gateway with the users 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 recognises 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 users profile in order to construct the personalized portal menu by reordering and/or promoting content links.
7. The Device Manager reads the device style sheet for the users 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 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 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, $P_u(News|Home)$ is calculated as the combined relative frequency accesses, taking genuine user accessees and default static hit values into acoun. Thus $P_u(News|Home) = (20 + 10)/(40 + 100) = 0.214$. Similarly, $P_u(World|Home)$ is calculated by combining chaining access probabilities so that $P_u(News|Home)P(World|News) = (20 + 10)/(40 + 100) * (5 + 10)/(10 + 20) = 0.107$.

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 $Sports$, $Soccer$, $News$, $F1$, $World$, 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.

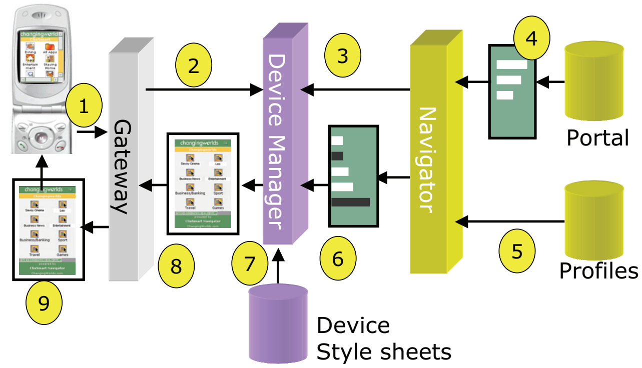

Figure 2: Constructing a personalized portal page.

By way of an example, Figure 3 presents a series of portal pages leading the user to their local cinema listings (*Ster Century* via a number of intermediate menu options. Assuming that this becomes a well travelled path for the user in question then 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 3(b) to illustrate this. The *Ster Century* service has been promoted to the top position within the *Entertainment* menu, reducing its click-distance significantly, by eliminanting 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 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 below.

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 descendents of m. This leads to an exponential growth in the number of menu options that must be considered as more an 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

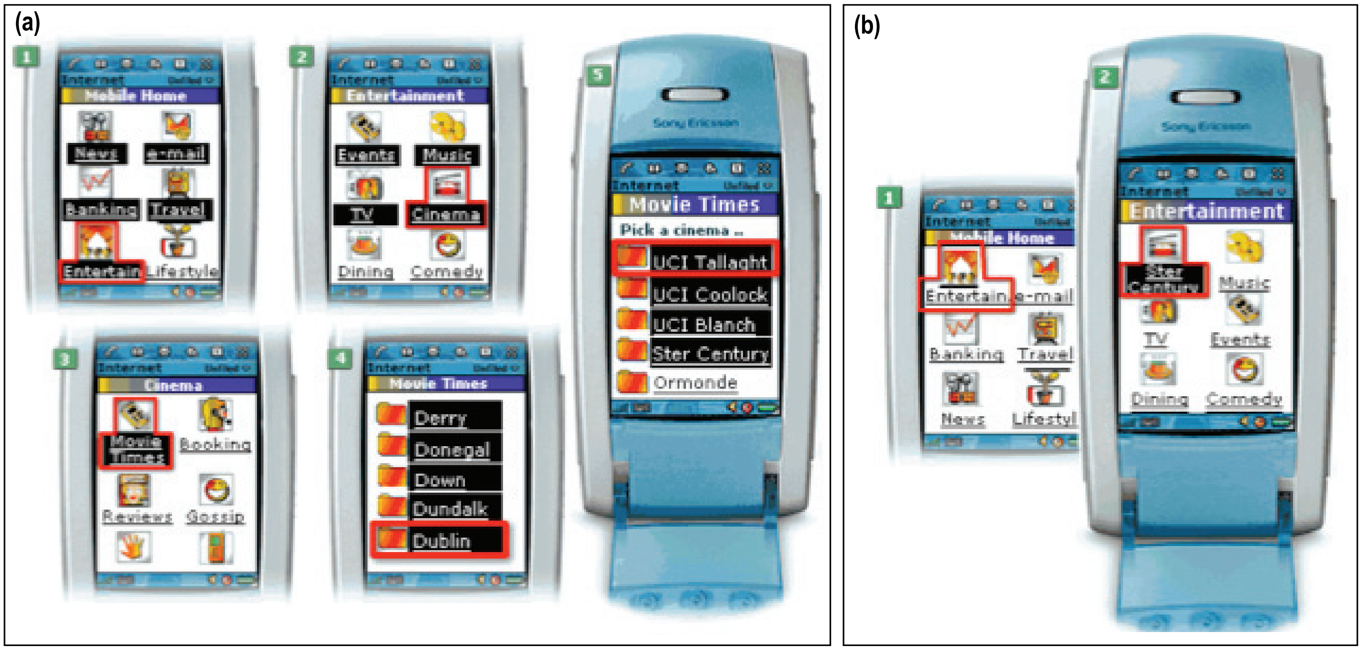

Figure 3: (a) An example sequence of navigation steps through a series of portal pages leading the user to their local cinema listings (*Ster Century*) via a number of intermediate menu options. (b) An example personalization scenario is presented in which the *Ster Century* service has been promoted 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

look-ahead during personalization; for example, promotion candidates for m may only be drawn from menus that are up to say 2 levels deep from m. Of course this limits promotion opportunities as their may well be compelling promotion candidates that exist 3 or 4 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 descendent 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 & 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 5 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 5 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 to better reflect 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 favourite 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 an promoted option can travel through the portal hierarchy.

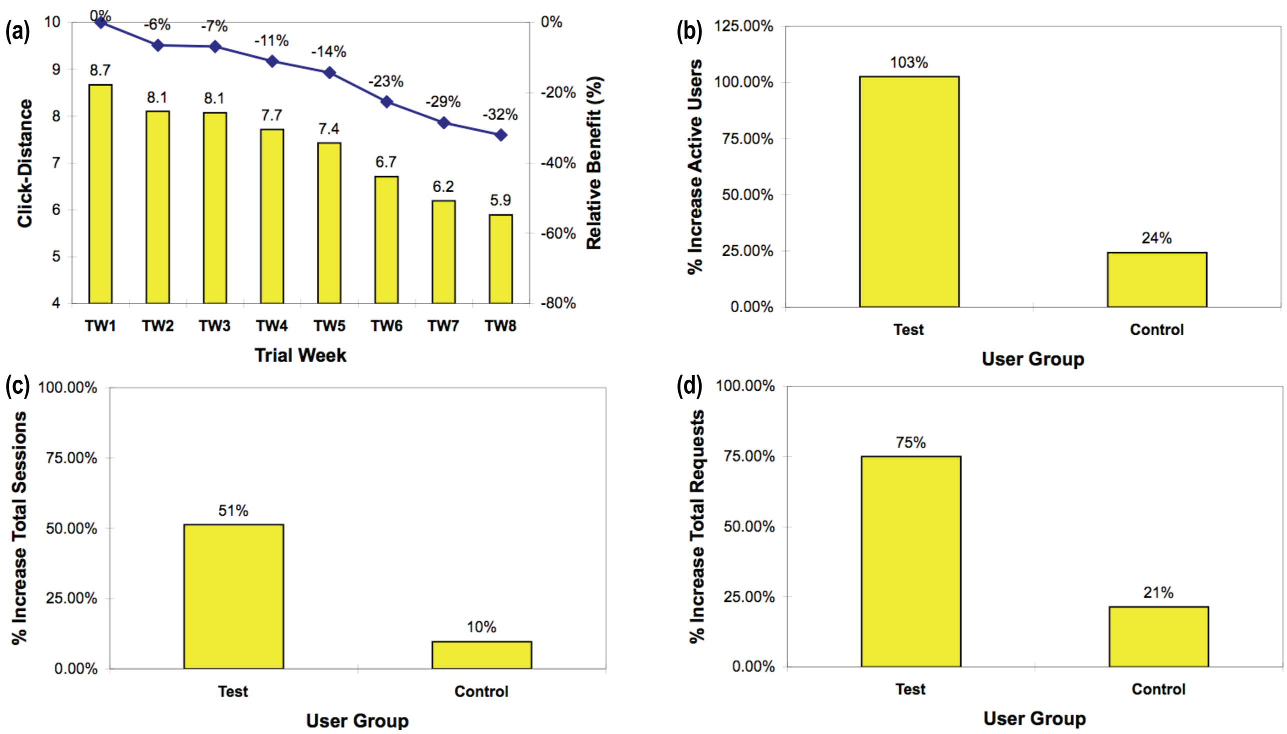

Figure 4: Key evaluation metrics: (a) average session click-distance for test users during each of the 8 trial weeks; (b) average comparative increase in the number of users accessing the portal on a weekly basis; (c) average comparative increase in the number of user sessions ; (d) average comparative increase in the number of page requests.

These controls provide a very effective mechanism for ensuring that users do not become disoriented in practice. It also allows operators 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 paper.

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 behaviour 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 2+ 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 8-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 4(a-d) and show a very significant benefit associated with the activity levels of the test group relative to the control. For example, in Figure 4(a) we see how the test group enjoys a gradual decline in their 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% to 5.9; indeed our studies indicate that on average click-distance will fall by about 50% over a 3 month period.

Figures 4(b-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 8-week trial period compared to the previous 8-weeks pre-trial. For example, in Figure 4(b) 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% between the 8-week pre-trial period and the 8-weeks during the trial; during the 8-week pre-trial period an average of just over 140 of the test users access 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%.

Similar increases are seen across other key metrics such as number of sessions (Figure 4(c)) and total requests (Figure 4(d)). For instance, we see a 54% (75% - 21%) 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 their entire subscriber base.

The results presented above relate to one particular trial for one particular operator. Similar advantages have been observed across 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% within 3 months.
2. The average number of user sessions increases by 30%.
3. The average number of requests per session increases by 35%.
4. The average number of *failed sessions* (sessions where users fail to locate content) falls by approximately 50%.

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 post-deployment. 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:

1. There was a significant increases in portal usage across all major content areas; for example, browsing was found to have increased by 28%, ringtone downloads were up by 14%, video content downloads were increased by almost 30%, and gaming downloads were increased by more than 35%.
2. 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% of users reported an increase in their satisfaction levels and more than 30% of users indicated that they would be less likely to churn.
3. 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 $15m in direct revue.

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. As mentioned previously the platform has been deployed in more that 40 mobile operators around the world, ranging from tier 1 operators (with in excess of 10 million subscribers), to tier 2 (5-10 million subscribers), and tier 3 operators (with less than 5 million subscribers). For example, it currently handles over 50% of of the UK's mobile Internet traffic.

The platform provides for full horizontal and vertical scalability and is commercially deployed with Service Level Agreements (SLAs) requiring 99.7% scheduled uptime (or better). Moreover, the platform is designed to offer real-time personalization without negatively impacting overall portal performance. For example, a single ClixSmart Navigator Server, running on a 4xCPU, SunFire 480R with 16GB 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.

Conclusions

Since 2001 ChangingWorlds has been helping mobile operators to understand and solve the usability and content discovery problems that frustrate many mobile subscribers. Our solution is unique in the mobile market, offering operators class-leading portal management and content delivery functionality alongside innovative personalization technology.

In this paper we have focused on how this 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 operators. The bottom-line is that this application of artificial intelligence technology in the mobile sector has had a significant positive impact on operator revenues and subscriber satisfaction. Everyday millions of mobile subscribers enjoy an enhanced mobile experience as a result of our personalization technology and everyday these subscribers consume more content than users of traditional one-size-fits-all portals.

As an added-value benefit of this technology, mobile operators have also captured unique and valuable business intelligence about the content preferences of their subscribers; for example, one operator now has in excess of 30 million subscriber profiles. 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 enjoy content and services that are relevant to their true needs.

References

Anderson, C.; Domingos, P.; and Weld, D. 2001. Adaptive Web Navigation for Wireless Devices. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, 879–884.

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.

Cellular-News. 2006. 2.5 Billion Mobile Phones In Use. http://www.cellular-news.com/story/19223.php. Last Checked 18th September 2006.

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 behaviour 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.

Herder, E. 2003. Modelling User Navigation. In *Proceedings of the 9th International Conference on User Modelling*, 417–419. Springer-Verlag.

IposInsight. 2006. Mobile Phones Could Soon Rival the PC As Worlds Dominant Internet Platform. http://www.ipsos-na.com/news/pressrelease.cfm?id=3049. Last Checked 18th September 2006.

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. ACM Press.

Marsden, G., and Jones, M. 2001. Ubiquitous computing and cellular handset interfaces are menus the best way forward? In *Proceedings of THE South African Institute of Computer Scientists and Information Technologists Annual Conference (SAICSIT)*, 111–119.

Miller, D. 1981. The Depth/Breadth Tradeoff in Hierarchical Computer Menus. In *Proceedings of the 25th Annual Meeting of the Human Factors and Ergonomics Society*, 296–300.

Perkowitz, M., and Etzioni, O. 2000. Towards adaptive web sites: Conceptual framework and case study. *Journal of Artificial Intelligence* 18(1-2):245–275.

Perkowitz, M. 2001. *Adaptive Web Sites: Cluster Mining and Conceptual Clustering for Index Page Synthesis*. PhD Thesis, Department of Computer Science and Engineering. University of Washington.

Ramsey, M., and Nielsen, J. 2000. *The WAP Usability Report*. Neilsen Norman Group.

Reiken, D. 2000. Special issue on personalization. *Communications of the ACM* 43(8).

Smyth, B., and Cotter, C. 2000. Wapping the Web: A Case-Study in Content Personalization for WAP-enabled Devices. In *Proceedings of the 1st International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH'00)*, 98–108.

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*. IOS Press.

Smyth, B., and Cotter, P. 2002b. The Plight of the Navigator: Solving the Navigation Problem for Wireless Portals. In *Proceedings of the 2nd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH'02)*, 328–337. Springer-Verlag.

Smyth, B. 2002. *The Plight of the Mobile Navigator*. MobileMetrix.

Zaphirs, P. 2000. Depth vs. Breadth in the Arrangement of Web Links. In *Proceedings of 44th Annual Meeting of the Human Factors and Ergonomics Society*, 139–144.

ZDNet. 2006. The mobile Internet: Are we there yet? http://news.zdnet.com/2100-1035_22-6110100.html. Last Checked 18th September 2006.