AI Magazine Volume 22 Number 2 (2001) (© AAAI)

Personalized Electronic Program Guides for Digital TV

Barry Smyth and Paul Cotter

■ Although today's world offers us unprecedented access to greater and greater amounts of electronic information, we are faced with significant problems when it comes to finding the right information at the right time—the essence of the information-overload problem. One of the proposed solutions to this problem is to develop technologies for automatically learning about the implicit and explicit preferences of individual users to customize and personalize the search for relevant information. In this article, we describe the development of the personalized television listings system (PTV),[1] which tackles the information-overload problem associated with modern TV listings data by providing an Internet-based personalized TV listings service so that each registered user receives a daily TV guide that has been specially compiled to suit his/her particular viewing preferences.

The term *information overload* has become almost synonymous with the internet and the World Wide Web, and today the internet's 300+ million users are finding it increasingly difficult to efficiently locate precisely relevant information content among its growing repository of 1+ billion pages. For example, modern search engines provide only a first cut through the information space, leaving the user with a significant search task to locate individual information items. This information overload is beginning to cause problems on the internet and is seen as a serious barrier to its future success.

This problem takes on even more significance when one considers the new generation of mobile phones, which offer users an alternative internet access route through the wireless application protocol (WAP). Web content (including text, graphics, forms, and hyperlinks) is displayed as wml (wireless markup language) encoded pages; WML is the equivalent of html (hypertext markup language) on WAP devices. These devices currently suffer from greatly reduced display sizes, limited bandwidth, and restricted on-board memory (figure 1). Under these conditions, it becomes even more important to be able to offer WAP users personalized information content because current WAP devices do not facilitate a "trawl" (search through large amounts of information) through even moderate quantities of information in conventional web terms.

Content personalization is one potential solution to the information-overload problem. It promises the precise delivery of user-targeted information by automatically learning about the preferences of individual users over time and by using this information to guide the search for, and presentation of, relevant information.

In this article, we focus on an emerging information-overload problem that is associated with the new generation of digital TV systems. We suggest that it will become almost impossible for people to cope with the promise of hundreds of TV channels and thousands of TV programs daily and that traditional TV guides will fail to provide any practical assistance. We present personal television listings (PTV) as a real solution to this problem. In short, PTV is an innovative internet service that uses content personalization techniques to automatically learn about the TV viewing preferences of individual users to provide them with highly customized and personalized daily TV guides. In particular, we focus on different versions of PTV, a fully deployed web-based system, and newly developed versions for WAP-based mobile phones and personal digital assistants (PDAs) (Cotter and Smyth 2000a, 2000b; Smyth and Cotter 2000a, 2000b, 1999).

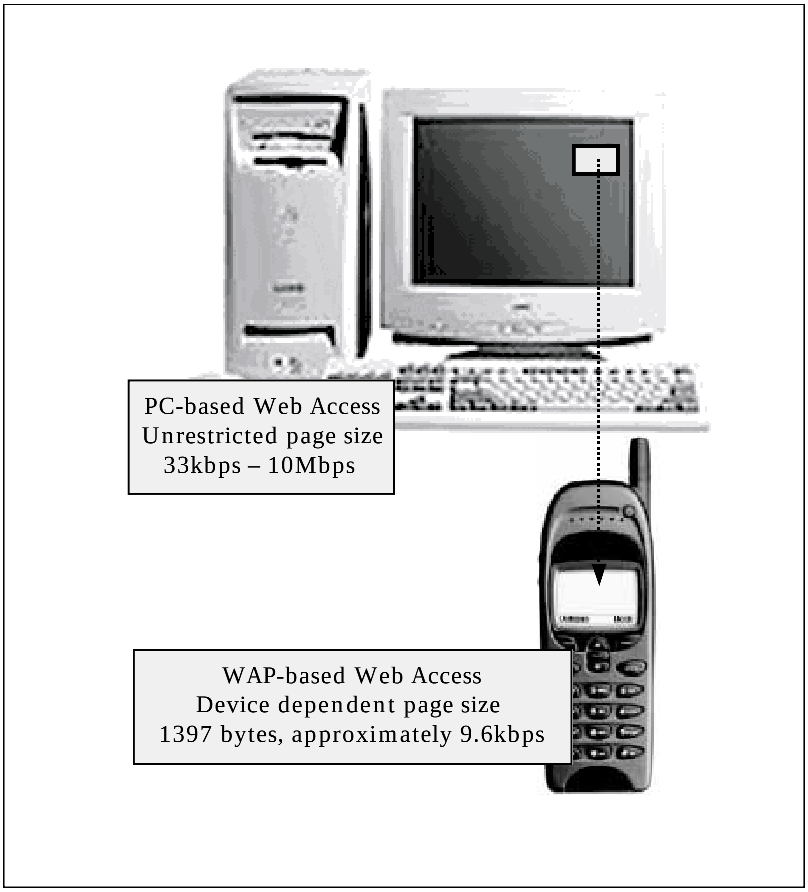

Figure 1. Compared to the Traditional Mode of Web Access, the PC, Wireless Application Protocol (WAP)–Based Access Devices, Such as Mobile Phones, Can Suffer from Greatly Reduced Screen Real Estate, Bandwidth, and Page Sizes.

Figure 2. A Sample Electronic Program Guide (EPG) Listing Programs for a Sample of Seven Channels over a One-Hour Time Slot (courtesy of ReplayTV, www.replaytv.com).

Problem Description

With the arrival of new cable and satellite TV services, and the next generation of digital TV systems, we will soon be faced with an unprecedented level of program choice (upwards of 200 channels and 4000 programs a day over the next 2 years). Navigating through this space represents a new variation on the information-overload theme, and it will become increasingly difficult to find out what relevant programs are showing on a given day.

Of course, the digital TV vendors are aware of these issues, and their current solution is the electronic program guide (EPG), providing users with on-screen access to online TV listings. However, simply providing an electronic equivalent of the paper-based TV guide is not a scalable solution to the problem. For example, a typical EPG might cover a 60-minute time slot for 5 to 10 channels in a single screen, which means that even a relatively modest lineup of 70 channels will occupy 10 to 15 screens of information for each 60-minute slot, or well over 160 screens for each viewing day (figure 2).

Of course, this information overload doesn't just introduce problems for digital TV users. The television channels themselves are faced with the significant problem of how to ensure that viewers will notice their programming content within a sea of alternatives. This problem is particularly difficult for the smaller channels and could ultimately have a negative impact on their ability to attract advertising revenue. In all likelihood, if a solution to this information-overload problem is not forthcoming, then users will probably focus their attention on a small number of larger channels, essentially avoiding the smaller channels.

This state of affairs is depicted in figure 3, which charts the level of personalization required to support different levels of content in a digital TV setting. We are currently positioned near the origin, with current levels of TV content pushing the limits of what traditional nonpersonalized TV guides (hard copy and online) can hope to usefully handle. The so-called zone of usefulness is wide in this portion of the chart, and many existing EPG solutions remain useful and fall within this zone. However, as we move to the future, and the number of TV channels increases (along with the available content), the zone of usefulness narrows rapidly. Traditional, nonpersonalized solutions rapidly move out of this zone, indicating that they are no longer capable of coping with the increased content levels.

Indeed, we maintain that the only effective solution is to provide a fully personalized EPG

(PEPG) that is capable of automatically learning about the viewing needs and preferences of individual users and alerting these users to the right programs at the right times. If successful, this type of EPG will remove the traditional channel boundaries to offer viewers their own personalized television channel, drawing together relevant programming content from across the full range of available channels no matter how small or big. In this way, viewers are guaranteed to receive the right information at the right time, and even the smallest channels will benefit from viewership as long as their program content is relevant to viewers.

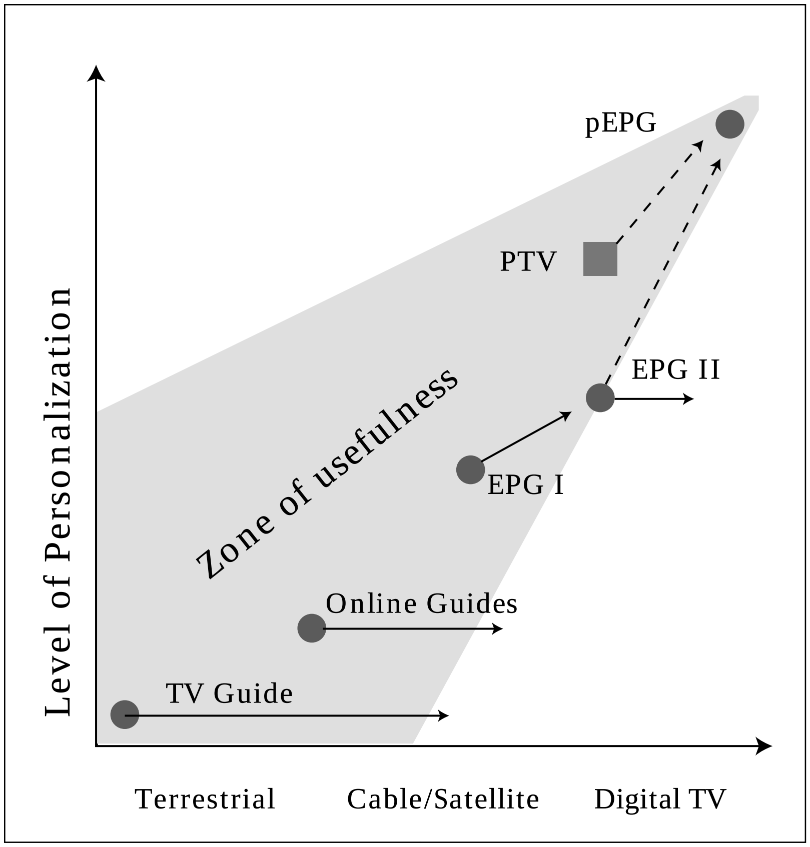

Figure 3. Personalization versus Content in the Digital TV Domain.

Application Description

The PTV project is motivated by the belief that the TV listings domain can benefit greatly from an EPG that incorporates content personalization techniques as a means of filtering and customizing TV listings information for individual users. In this section, we describe the PTV system, focusing in particular on how it produces personalized TV guides by integrating user profiling, case-based reasoning, and collaborative profiling techniques.

Hardware and Software

PTV is a JAVA-based client-server system and includes a specially designed, optimized, multithreaded server and dynamic HTML/WML/XML page generator plus all the AI and user-profiling components necessary for personalization. It currently runs on LINUX on an Intel 450-megahertz processor with 64 megabytes of random-access memory (RAM) and has been stress-tested beyond 7 million hits a month without any substantial performance degradation.

System Architecture

PTV users can register, log in, and view their personalized TV guides as specially customized HTML pages (for conventional PC-based access) or as WML pages (for mobile phone access). The architecture of PTV (figure 4) does not depend on the mode of access (PC versus WAP-based device), and all user interaction is handled by HTTP. The heart of the system lies with its server-side components, which handle all the main information-processing functions such as user registration and authentication, user profiling, guide compilation, and the all-important program recommendation and grading.

Profile Database and Profiler

The key to PTV's personalization facility is an accurate database of user profiles. Each profile encodes the TV preferences of a given user, listing channel information, preferred viewing times, program and genre preferences, guide preferences, and so on (figure 4). Preliminary profile information is collected from the user at registration time to bootstrap the personalization process. However, the majority of information is learned from grading feedback provided by the user; each recommended program is accompanied with grading icons or links that allow the user to explicitly evaluate the proposed recommendation.

Program Case Base

This database contains the program content descriptions (program cases). Each entry describes a particular program using features such as the program title, genre information, the creator and director, cast or presenters, the country of origin, and the language; an example program case for the comedy *Friends* is shown in figure 4. This information repository is crucial for the content-based (case-based) recommendation component of PTV.

Schedule Database

This database contains TV listings for all supported channels. Each listing entry includes details such as the program name; the viewing channel; the start and end times; and, typically, some text describing the program in ques-

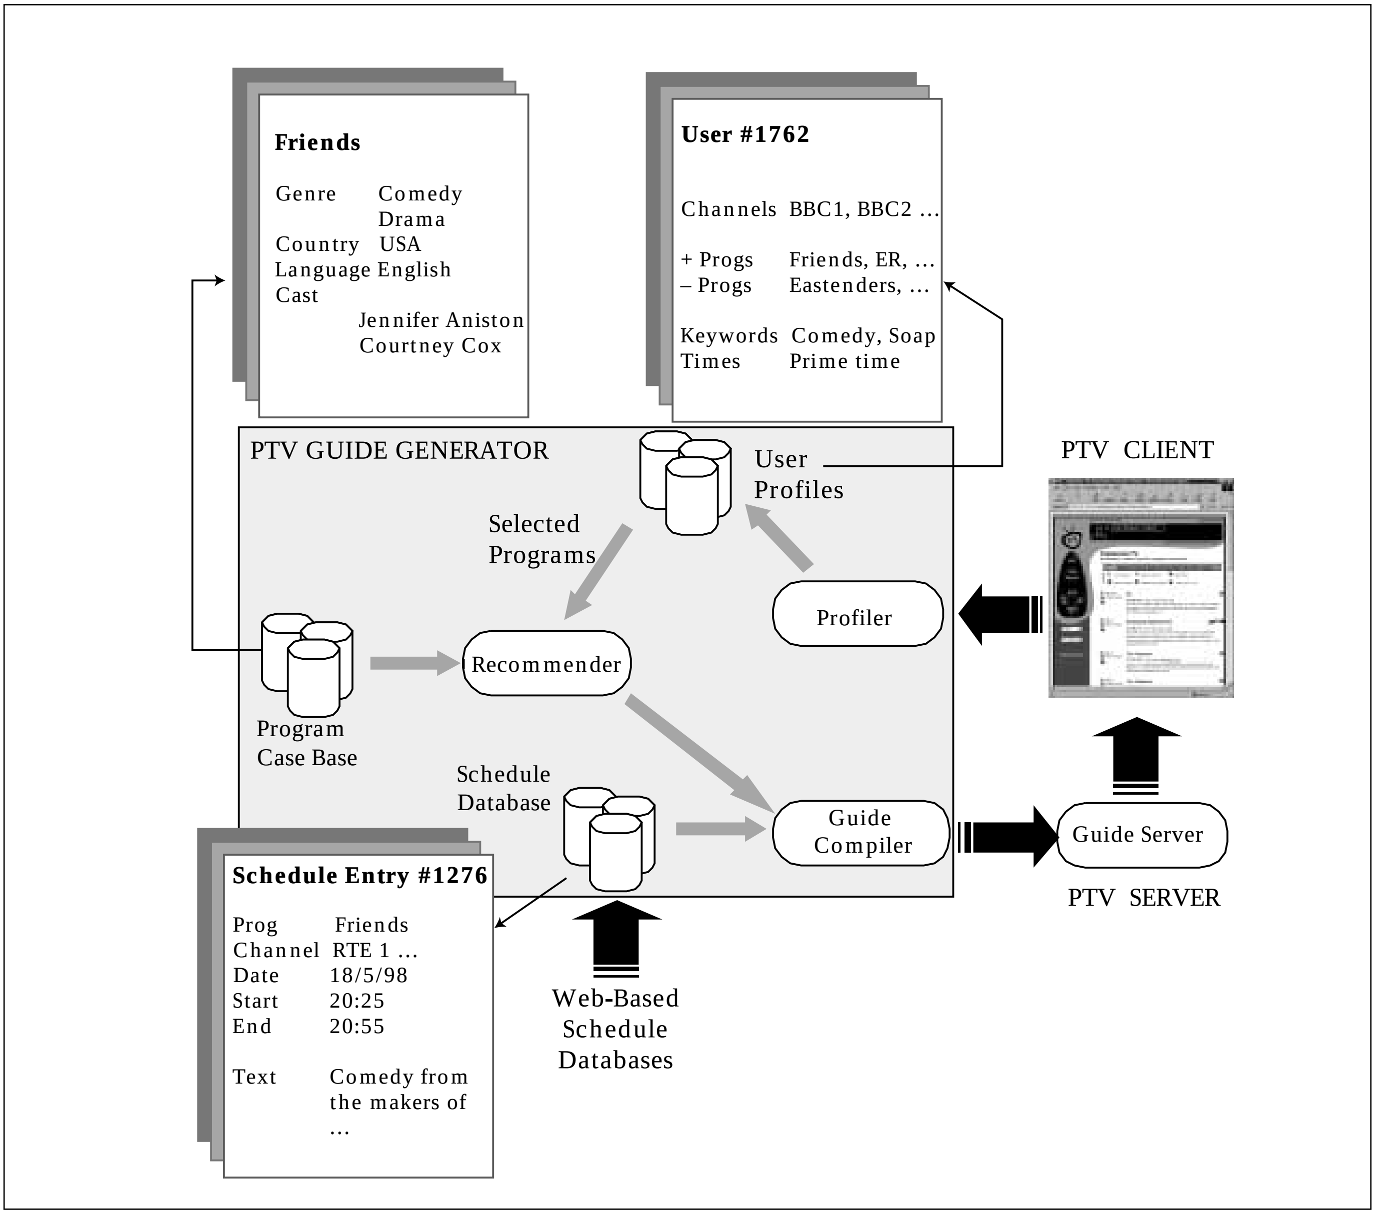

Figure 4. An Overview of the PTV System Architecture.

tion (see the schedule entry example in figure 4). The schedule database is constructed automatically from electronic schedule resources.

Recommender

The *recommender component* is the intelligent core of PTV. Its job is to take user-profile information and select new programs for recommendation to a user. In the next section, we explain how PTV uses a hybrid recommendation approach that combines content-based and collaborative recommendation strategies.

Guide Compiler

To compile a personalized guide for a user, PTV uses two program lists: (1) programs listed as positive in the user's profile, along with those programs selected for recommendation (that do not occur in the profile), and (2) a list of programs to be aired on the specified date by channels listed in the user's profile. The intersection of these lists is the set of programs that will finally appear in the personalized guide.

Guide Translator

The guide compiler produces a generic guide format, which is automatically converted into an HTML, XML, or WML page by the guide translator, as appropriate. Although individual guides are converted into single HTML pages for the web, they are converted into multiple WML pages (or cards) for mobile phone use; this conversion is necessary to solve the problems of limited presentation space (and memory space) that exist on current WAP phones.

Problem Description

AI techniques are central to the success of the PTV system. Specifically, the ability to accurately personalize the TV guide of an individual user relies on the availability of an accurate model of this user (user profiling) and an ability to relate this profile to relevant program content (program recommendation). In this section, we outline PTV's user-profiling component and its content recommendation strategies.

Acquiring User Profiles

The success of PTV depends ultimately on the quality of its personalized guides, which depends largely on the quality of the user profiles and their ability to represent the viewing preferences of users (Jennings and Higuchi 1993; Kay 1995; Perkowitz and Etzioni 1997). In PTV, each user profile contains two types of information: (1) domain preferences and (2) program preferences. *Domain preferences* describe general user preferences such as a list of available TV channels, preferred viewing times, subject keywords and genre preferences, and guide format preferences. *Program preferences* are represented as two lists of program titles, a positive list containing programs that the user has liked in the past and a negative list containing programs that the user has disliked.

Profile information is gathered in two ways. Users are encouraged to manually update their profiles directly by specifying viewing preferences. However, although manual profile editing has its advantages (usually in terms of profile accuracy), it is a burden for the users. In particular, we have found that users are happy to provide fairly complete domain preferences but tend to provide only limited program preferences. For this reason, PTV includes a profile update facility that is driven by direct user feedback through a set of grading icons listed beside guide programs. PTV's profiler uses this information to automatically alter a user's profile in a number of ways. The simplest modification is to update the program preference lists by adding positively or negatively graded programs to the appropriate list. However, the domain preferences can also be altered. For example, viewing time preferences can be adjusted if a user frequently prefers prime-time programs to morning shows. In general, this long-term feedback connection between user and system is vital if PTV is to maintain an accurate picture of each user over time.

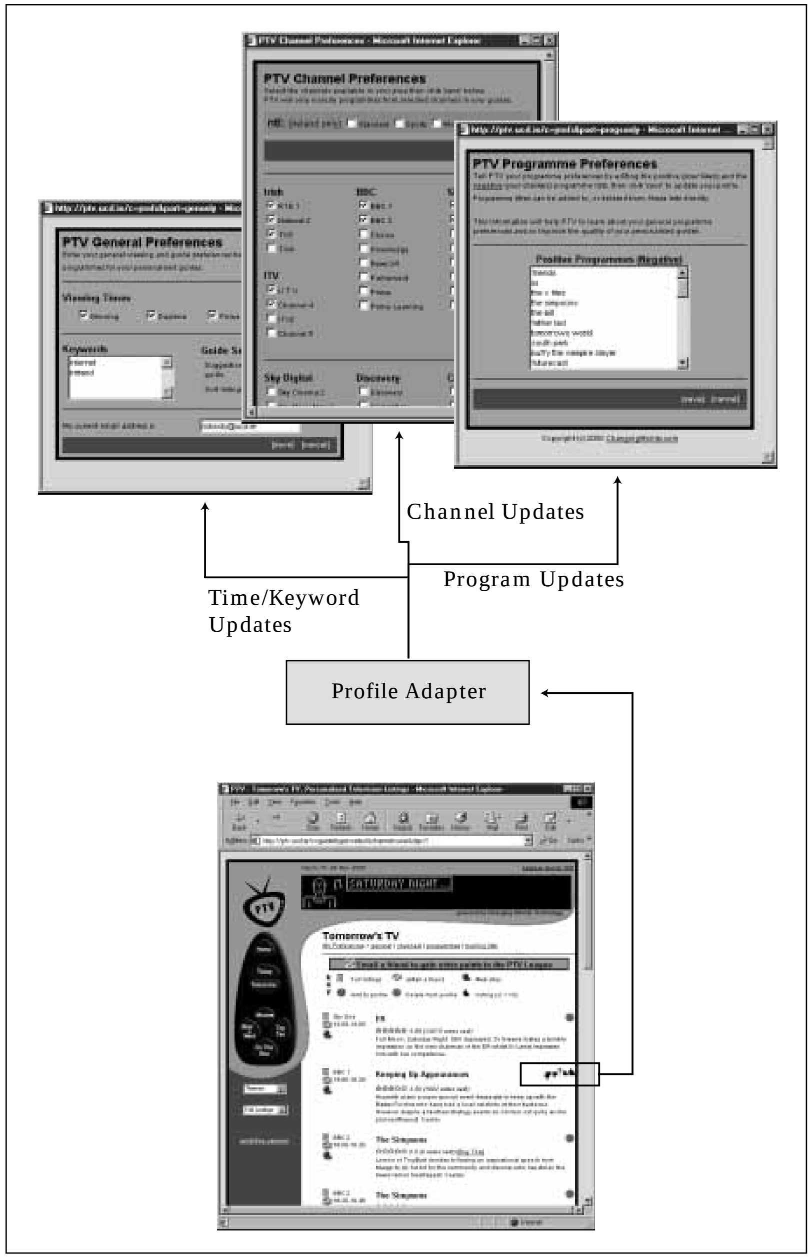

Figure 5. User Profiles and Feedback in Personalized Television Listing (PTV) (Web Based).

Figure 5 outlines how user profiles are updated in the web-based PTV system. A similar scenario operates in the WAP-based version of PTV except that preference and grading options require a number of individual pages rather than have a single preferences page or integrate the grading icons with the main guide pages, as in the web-based version.

A Content-Based Recommendation Approach

Ultimately in PTV, personalizing a given user's TV guide boils down to recommending the right programs for the user given his/her various viewing constraints. PTV harnesses two complementary recommendation strategies to base its recommendations on the programs that a given user has liked in the past (case based or content based) and on the programs that similar users like (collaborative). In this section, we look at the more traditional content-based (or case-based) approach (Watson 1997), and in the following subsection, we look at the complementary collaborative recommendation strategy.

The basic philosophy in content-based recommendation is to recommend items that are similar to those items that the user has liked in the past (Smyth and Cotter 2000a, 2000b, 1999; Balabanovic and Shoham 1997; Hammond, Burke, and Schmitt 1996). For PTV, this means recommending programs that are similar to the programs in the positive program list and dissimilar to those in the negative program list. Three components are needed for content-based recommendation: (1) content descriptions for all TV programs (see the program case base in the section entitled Program Description and figure 3), (2) a compatible content description of each user's profile, and (3) a procedure for measuring the similarity between a program and a user.

PTV's program case base has already been outlined, and an example case is shown in figure 4. Each case is described as a set of features, and the similarity between two cases can be defined as the weight sum of the similarity between corresponding case features. However, there is no direct means of computing the similarity between a case and a user profile because user profiles are not described as a set of case features. Instead, each raw user profile is converted into a feature-based representation called a *profile schema*. Basically, the profile schema corresponds to a content summary of the program preferences contained in a user profile, encoded in the same features as the program cases. The similarity between a profile and a given program case can then be computed using the standard weighted-sum similarity metric, as shown in equation 1, where

$f_i^{Schema(u)}$ and f_i^{p}

are the ith features of the schema and the program case, respectively.

$$PRGSIM\left(Schema(u), p\right) = \sum w_i \bullet SIM\left(f_i^{Schema(u)}, f_i^{p}\right) \qquad (1)$$

A problem with content-based methods is the knowledge engineering effort required to develop case representations and similarity models. Furthermore, because content-based methods make recommendations based on item similarity, the newly recommended items tend to be similar to the past items, leading to reduced diversity. In the TV domain, this reduced diversity can result in narrow recommendation lists, for example, a lot of comedies if the majority of profile programs are comedies.

A Collaborative Recommendation Approach

Collaborative recommendation methods, such as automated collaborated filtering, are an alternative to content-based techniques. Instead of recommending new items that are similar to the ones that the user has liked in the past, they recommend items that other similar users have liked (Billsus and Pazzani 1998; Konstan et al. 1997; Balabanovic and Shoham 1997; Maltz and Ehrlich 1995; Shardanand and Maes 1995; Goldberg, Nichols, and Oki 1992). In addition, instead of computing the similarity between items, we compute the similarity between users or, more precisely, the similarity between user profiles. In PTV, the recommendations for a target user are based on the viewing preferences of the k-most similar users.

PTV computes user similarity by using a simple graded difference metric, shown in equation 2; where $p(u)$ and $p(u')$ are the ranked programs in each user's profile, and

$r\left(p_i^{u}\right)$

is the rank of program p_i in profile u. The possible grades range from –2 to +2, and missing programs are given a default grade of 0. Of course, this technique is just one of several similarity techniques that have proved useful in PTV, and any number of techniques could have been used, for example, statistical correlation techniques such as Pearson's correlation coefficient (for example, Billsus and Pazzani [1998]).

$$PRFSIM\left(u, u'\right) = \frac{\sum_{p(u) \cup p(u')} \left| r\left(p_i^{u},\right) - r\left(p_i^{u'}\right) \right|}{4 \bullet \left| p(u) \cup p(u') \right|} \qquad (2)$$

$$PRGRANK\left(p, u\right) = \sum_{u' \in U} PRFSIM\left(u, u'\right) \qquad (3)$$

Once PTV has selected k similar profiles for a given target user, a recommendation list is formed from the programs in these similar profiles that are absent from the target profile. This list is then ranked, and the top r programs are selected for recommendation. The ranking metric is shown in equation 3; U is the subset of k-nearest profiles to the target that contain a program p. This metric biases programs according to their frequency in the similar profiles and the similarity of their recommending user.

In this way, popular programs suggested by similar users tend to be recommended.

Collaborative filtering is a powerful technique that solves many of the problems associated with content-based methods. For example, there is no need for content descriptions or sophisticated case-similarity metrics. In fact, high-quality recommendations that would ordinarily demand a rich content representation are possible. Moreover, recommendation diversity is maintained because relevant items that are dissimilar to the items in a user profile can be suggested.

Collaborative filtering does suffer from some shortcomings. There is a startup cost associated with gathering enough profile information to make accurate user-similarity measurements. There is also a latency problem, which means that new items will not be recommended until these items have found their way into sufficiently many user profiles. This latency problem is particularly difficult in the TV domain because new and "one-off programs" (special single-episode programs such as concerts) occur regularly and do need to be considered for recommendation even though these programs will not have made it into any user profiles. The key to PTV's success is the use of a combined recommendation approach. For a given guide, a selection of programs is suggested; some are content-based recommendations (including new or one-off programs), and others are collaborative recommendations. In particular, recommendation diversity is ensured through the use of collaborative filtering, and the latency problem can be solved by using content-based methods to recommend new or one-off programs.

System Demonstration

In this section, we look at the use of the PTV system by a new user, stepping through each of the basic stages from initial registration through guide viewing. To take advantage of PTV's personalization facilities, each new user must register an account with PTV.

An example of a personalized guide produced by PTV is shown in figure 6, which shows four separate listings for three programs, *Friends*, *Married with Children*, and *Ally McBeal*. This guide has been produced for a user with a strong interest in American sitcoms. The user has previously expressed an interest in programs such as *Friends*, and PTV has further recommended *Married with Children* and *Ally McBeal*, which the user has not encountered before but that PTV feels are relevant.

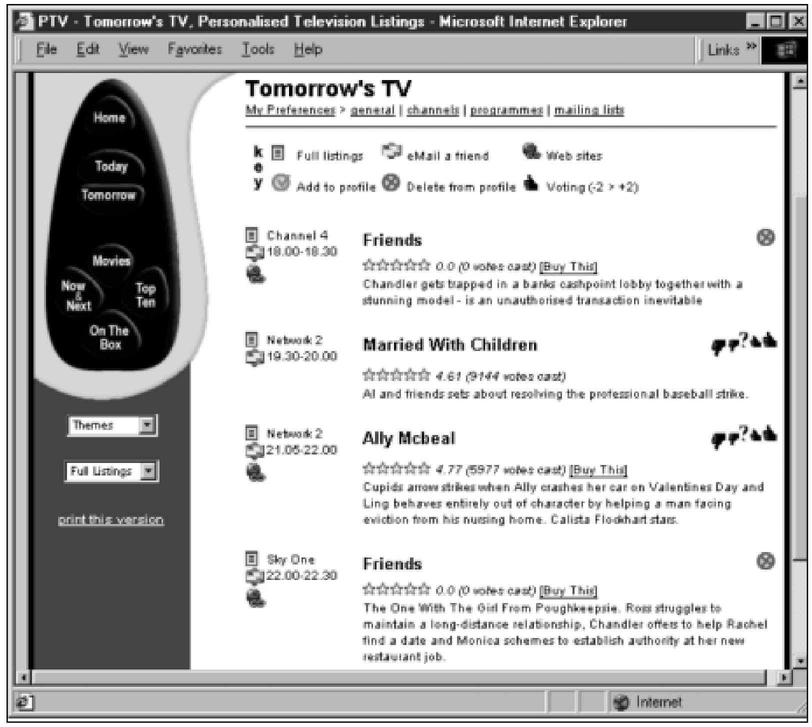

Figure 6. A Sample PTV Personalized TV Guide.

Because *Married with Children* and *Ally McBeal* are recommendations from PTV, the user is afforded the opportunity to rate these suggestions by clicking on the grading icons (thumbs-up and thumbs-down icons) beside these programs. Importantly, this type of information allows PTV to learn about a user's specific and general viewing preferences. For example, by rating *Ally McBeal* positively (small or large thumbs-up), PTV will learn not only that the user is interested in this particular program but also that the user is interested in a range of similar programs, including other American sitcoms and courtroom dramas. Moreover, PTV will also learn about more general viewing preferences, such as the fact that the user likes to watch shows on network 2, which airs during primetime.

Of course, the ultimate judgment of a user's interest in a program is whether the user actually watches it, but in the current incarnation of the PTV system, there is no way of capturing this information directly by monitoring a user's online behavior. However, ultimately, users will be able to access systems such as PTV through their television sets, and in this case, it will be possible to recognize whether a user watches a recommended show, thereby doing away with the need to elicit direct feedback from the user.

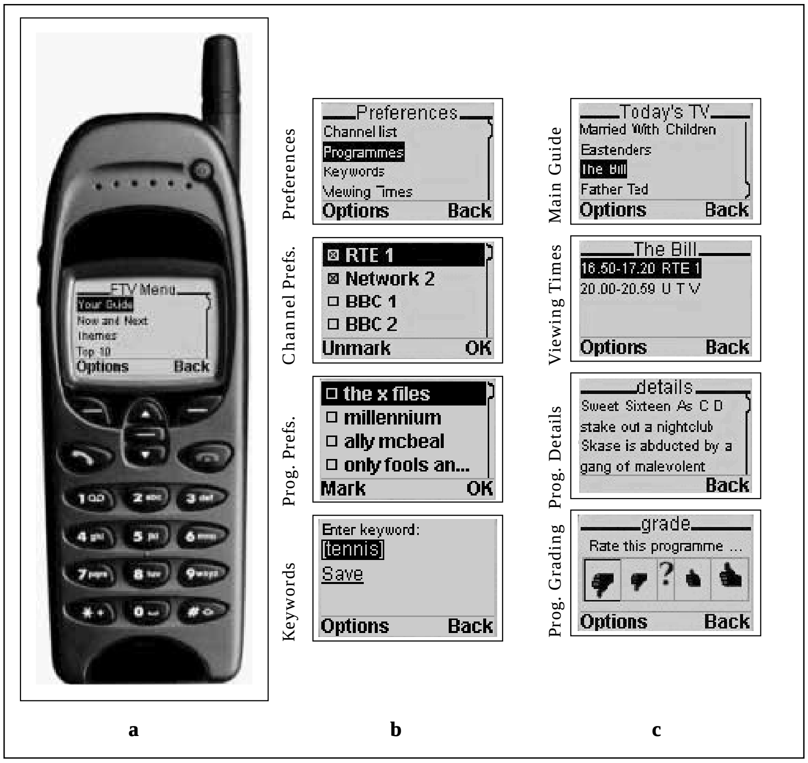

Figure 7. Screen Shots of PTV for Mobile Phones.

A. The PTV wireless application protocol main menu.
B. The preferences screens. C. A part of a personalized guide, including viewing times, program description, and grading screens.

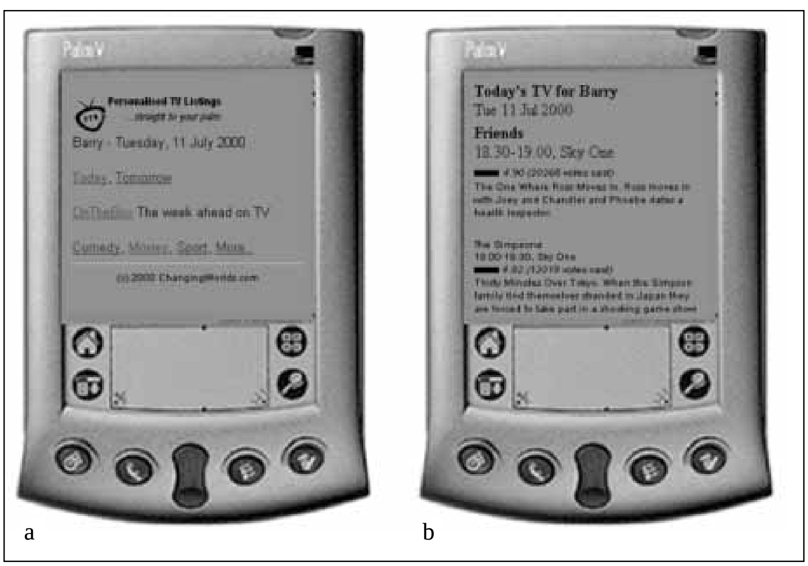

Figure 8. Screen Shots of PTV for Handheld Computing Devices.

A. The PTV personal digital assistant (PDA) main menu.
B. A PTV PDA personalized guide, including program ratings.

Multiple-Delivery Touchpoints

PTV has recently been adapted for a variety of internet touch points, including WAP devices, PDAs, and TV "set-top boxes" (cable company receivers), and in this section, we provide a brief overview of these new services. Obviously, when one considers the limitations associated with non-PC modes of internet access, such as restricted screen space, the importance of personalization takes on a whole new meaning. For example, current WAP devices such as WAP-enabled mobile phones offer a screen area that is as little as 1/200th that of a typical PC monitor, so it is vitally important that valuable screen "real estate" is not wasted on irrelevant content.

Figures 7 and 8 present screen shots from the WAP and PDA versions of the PTV system (see also Cotter and Smyth [2000a, 2000b]). Both versions offer the same functions as the web-based version of PTV but are specially customized for the WAP and PDA environments. In addition, PTV has also been customized to work with a variety of internet-enabled set-top-boxes, including the NetGem set-top-box range.[2,3]

Deployment, Evaluation, and Maintenance

PTV was originally developed as part of a three-year basic research project in the Department of Computer Science at University College Dublin, Ireland. The resulting personalization technology was reimplemented as a commercially viable personalization engine during 1999 (approximately nine person-months). A well-developed set of tools and systems now exist for rapidly developing new commercial versions of PTV to suit a wide range of client needs. For example, the latest development of the WAP-based version of PTV required only eight person-weeks of development time.

Application Use and Payoff

PTV went live in January 1999, and the number of registered users has grown to nearly 20,000, with as many as 200,000 personalized guides generated each month. Furthermore, PTV has not been publicized; so, its current popularity is based largely on word of mouth and unsolicited press coverage. In fact, since the launch, the site has received a Yahoo! Site of the Month Award as well as many favorable press reviews from Irish and U.K. magazines and newspapers, including *ComputerActive, Dot.ie, PC Live, Business and Finance, the Sunday Business Post,* and the *Irish Times*. PTV was a finalist in the Irish Golden Spider Internet Awards under the best use of technology category. Moreover, ChangingWorlds has recently won the Isiah Software Associations Technology Innovation Award for 2000 for its CLIXSMART personalization technology.

PTV has been so successful that a campus company called ChangingWorlds has been established in University College Dublin to market the PTV system and, in particular, its underlying personalization technology.[4] This technology has led to the development of the CLIXSMART personalization engine and a range of personalization tool kits that have been customized for the mobile-wireless and digital TV domains (Smyth and Cotter 2000). Personalized applications, based on the CLIXSMART engine, have already been licensed to the following clients: (1) Ireland.com portal site, Ireland's largest portal, run by the *Irish Times* newspaper group;[5] (2) Unison, a leading Irish internet service provider that is unique in delivering internet services through the TV;[6] and (3) Eircell, Ireland's leading mobile operator.[7]

End-User Evaluation

Of course, from a scientific viewpoint, the big question concerns the accuracy of PTV's personalization facility. From January to June 1999, real users carried out an extensive and detailed study on PTV. In total, 310 users completed a comprehensive questionnaire regarding all aspects of PTV, including its personalization quality, speed, and ease of use, the results of which are summarized in the pie charts shown in figure 9; see Smyth and Cotter (1999) for further experimental results.

Clearly, the results are extremely positive. Critically, only 3 percent of users found the personalized guides to be of poor quality, and 99 percent of users found the site to be easy to use as a source of TV listings. Moreover, 88 percent of users found the response time to be acceptable, which we view very positively, especially because PTV's pages are created dynamically, and today's internet has a limited speed.

Maintenance

The PTV systems are designed to need minimal ongoing maintenance. For example, the maintenance of the user profiles and the program schedules is fully automatic. In fact, in our experience, the only manual maintenance that is required involves the addition of new program cases and the addition of new channels and cable regions. Even a relatively inexperienced user can manage both of these maintenance options by using PTV's built-in tools.

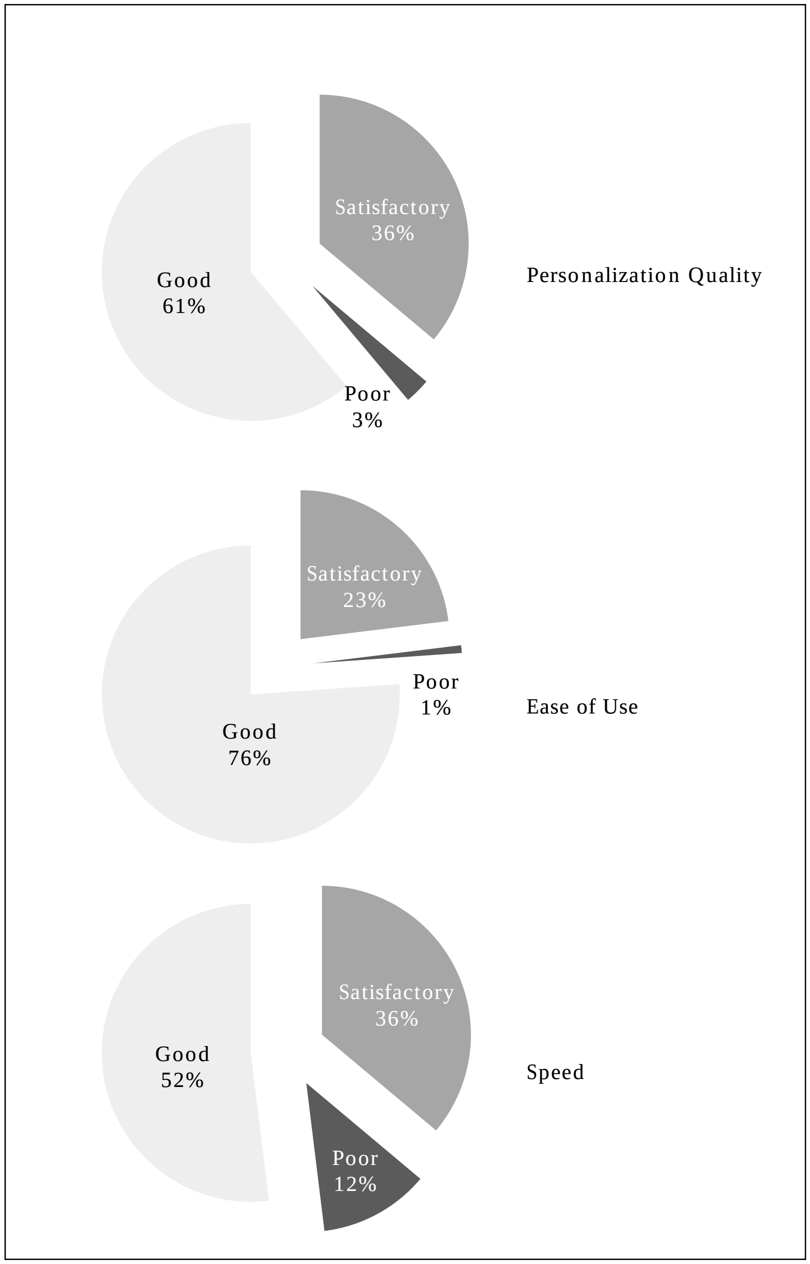

Figure 9. User Survey Summary Results.

Conclusions

We believe that PTV represents a convergence of technologies that provides an effective solution to the very real problem of providing people with relevant TV listings information as digital TV becomes a reality. PTV personalizes TV information to meet the viewing preferences of individual users by integrating a range of different information-filtering strategies, case-based reasoning and collaborative filtering, with user-profiling techniques. The resulting hybrid personalization technique allows program recommendations to be made according to the type of programs a target user has enjoyed in the past as well as the programs that other similar users have enjoyed.

To date, this technology has been deployed on a number of leading Irish web sites, and these applications have proven to be successful, with widespread adoption across the Irish internet market. Recently, similar success has been forthcoming in the mobile domain as WAP users recognize the real benefits of high-quality content personalization on their restricted mobile handsets. In fact, we argue that traditional TV listings services are not appropriate given the screen and bandwidth limitations of the current generation of WAP devices—a personalized service such as PTV is the best available solution.

Of course, PTV's personalization technology is not restricted to personalizing TV listings content. The PTV systems are built around the CLIXSMART personalization engine, which can readily be adapted to practically any source of information content. ChangingWorlds is currently using this technology to develop the next generation of intelligent, personalized information services.

Notes

1. www.ptv.ie.
2. www.netgem.com.
3. The interested reader is referred to www.unison.ie/tv for an online demonstration.
4. www.changingworlds.com.
5. www.ireland.com.
6. www.unison.ie.
7. www.e-merge.ie.

References

Balabanovic, M., and Shoham Y. 1997. FAB: Content-Based Collaborative Recommender. *Communications of the ACM* 40(3): 66–72.

Billsus, D., and Pazzani, M. J. 1998. Learning Collaborative Information Filters. In *Proceedings of the International Conference on Machine Learning*, 46–54. San Francisco, Calif.: Morgan Kaufmann.

Cotter, P., and Smyth, B. 2000a. Personalization Technologies for the Digital TV World. In *Proceedings of Prestigious Applications of Intelligent Systems, (PAIS'00)*, 701–705. Berlin, Germany: IOS.

Cotter, P, and Smyth, B. 2000b. Wapping the Web—A Case Study in Content Personalization for WAP-Enabled Devices. In *Proceedings of the International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, (AH2000)*, 98–108. Berlin, Germany: Springer-Verlag.

Goldberg, D.; Nichols, D.; Oki, B. M.; Terry, D. 1992. Using Collaborative Filtering to Weave an Information Tapestry. *Communications of the ACM* 35(12): 61–70.

Hammond, K. J.; Burke, R.; and Schmitt, K. 1996. A Case-Based Approach to Knowledge Navigation. In *Case-Based Reasoning Experiences: Lessons and Future Directions*, 125–136. Cambridge, Mass.: MIT Press.

Jennings, A., and Higuchi, H. 1993. A User Model Neural Network for a Personal News Service. *User Modeling and User-Adapted Information* 3(1): 1–25.

Kay, J. 1995. Vive la Difference! Individualized Interaction with Users. In Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, 978–984. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Konstan, J. A.; Miller, B. N.; Maltz, D.; Herlocker, J. L.; Gordan, L. R.; and Riedl, J. 1997. GROUPLENS: Applying Collaborative Filtering to Usenet News. *Communications of the ACM* 40(3): 77–87.

Maltz, D., and Ehrlich, K. 1995. Pointing the Way: Active Collaborative Filtering. In Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '95), 202–209. New York: Association of Computing Machinery.

Perkowitz, M., and Etzioni, O. 1997. Adaptive Web Sites: An AI Challenge. In Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, 16–23. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Shardanand, U., and Maes, P. 1995. Social Information Filtering: Algorithms for Automating "Word of Mouth." In Proceedings of the Conference on Human Factors in Computing Systems (CHI95), 210–217. New York: Association of Computing Machinery.

Smyth, B., and Cotter, P. 2000a. A Personalized Television Listings Service. *Communications of the ACM* 43(8): 107–111.

Smyth, B., and Cotter, P. 2000b. A Personalized TV Listings Service for the Digital TV Age. *Journal of Knowledge-Based Systems* 13(2–3): 53–59.

Smyth, B., and Cotter, P. 1999. Surfing the Digital Wave: Generating Personalized TV Listings Using Collaborative, Case-Based Recommendation. In *Proceedings of the International Conference on Case-Based Reasoning*, 561–571. Berlin, Germany: Springer-Verlag.

Watson, I. 1997. *Applying Case-Based Reasoning: Techniques for Enterprise Systems*. San Francisco, Calif.: Morgan Kaufmann.

Barry Smyth is a lecturer in the Department of Computer Science at University College Dublin in Ireland and the cofounder of ChangingWorlds, a company specializing in the development of advanced personalization technologies. Smyth received his Ph.D. in 1996 from Trinity College Dublin and has worked as a research scientist for Hitachi Ltd. His research interests include intelligent internet technologies, content personalization, case-based reasoning, and machine learning. He has made fundamental contributions to a number of research areas, including case-based reasoning and content personalization, and this work has won a number of prestigious international research awards, including best paper awards at the International Joint Conference on Artificial Intelligence (IJCAI-95), the International Conference on Case-Based Reasoning (ICCBR-95), and the British Computer Society Conference on Knowledge-Based Systems and Applied Artificial Intelligence. Smyth has published in over 80 refereed publications in scientific journals, books, and conferences and has presented his work worldwide. His e-mail address is Barry.Smyth@ucd.ie

Paul Cotter is the cofounder and chief technical architect at ChangingWorlds. His research interests include automatic content personalization technologies. Cotter has been responsible for the development of the CLIXSMART personalization engine and the successful PTV personalized TV listings portal. Cotter has previously worked as a researcher in the Department of Computer Science at University College Dublin, and his published research has won a number of international awards. His e-mail address is Paul.Cotter@changingworlds.com.