Explicit vs Implicit Profiling - A Case-Study in Electronic Programme Guides

Derry O'Sullivan, Barry Smyth, David Wilson
Smart Media Institute, University College Dublin, Ireland
{dermot.osullivan, barry.smyth, david.wilson}@ucd.ie

Abstract

In this paper, we evaluate the use of implicit interest indicators as the basis for user profiling in the Digital TV domain. Research in more traditional domains, such as Web browsing or Usenet News, indicates that some implicit interest indicators (e.g., read-time and mouse movements) are capable of serving as alternative to explicit profile information such as user ratings. Consequently, the key question we wish to answer relates to the type of implicit indicators that can be identified within the DTV domain and the extent to which they can accurately reflect a user's true preferences.

1 Introduction

Rccommcndcr systems combine user profiling, machine learning and information retrieval methods to help users locate information by learning their preferences and making personalized recommendations to individuals ([Hammond *et al,* 1996; Balabanovic and Shoham, 1997; Smyth and Cotter, 2000; Good *et al,* 1999; Goldberg *et al.,*1992]). There are two basic profiling strategies. *Explicit profiling* means asking the user direct questions about their preferences, usually in the form of an item rating; e.g. in PTV-Plus (www.ptvplus.com), a TV show recommender [Smyth and Cotter, 2001; 2000], users rate individual shows on a five point scale from 'strong dislike' to 'strong like'. Explicit profiling has been used in a range of scenarios, from movies [Resnick *et al.,* 1994] to jobs [Rafter *et ai,* 2000], and the resulting profiles are assumed to be reliable indicators of user preferences. However, users are lazy and do not like responding to explicit requests for feedback. As a result explicit profiles often contain few ratings, which can limit the degree to which they will accurately reflect a user's changing preferences over time.

Implicit profiling methods construct user profiles by *inferring* user ratings from so-called *interest indicators* based on user's interactions with the system [Claypool *et al.,* 2001; Goecks and Shavlik, 1999]. For example, [Morita and Shinoda, 1994; Konstan *et ai,* 1997; Claypool *et al,* 2001] report a positive correlation between reading times (for Usenet News and Web pages) and content ratings. [Claypool *et ai,* 2001] shows a correlation between content ratings and both mouse movements and scrolling although mouse movement was only found to be a useful rating predictor at lower rating levels. Interestingly, [Claypool *et ai,* 2001] finds no correlation between user ratings and the number of clicks to links on a Web page.

In this paper we seek to evaluate the use of implicit profiling in a very different application domain, Digital TV (DTV). DTV users are finding it increasingly difficult to identify relevant TV shows among the hundreds of available channels and recommender systems such as PTVPlus have been forwarded as a potential solution [Smyth and Cotter, 2000], PTVPlus uses its recommendation engine to generate a set of TV programme recommendations for a target user based on their profiled interests using a combination of collaborative and case-based strategies, and presents these recommendations in the form of a personalized programme guide. So far research in the area of personalisation and recommendation within the DTV domain has focused almost exclusively on the use of explicit profiling. Consequently, the key question we wish to answer relates to the type of implicit interest indicators that can be identified within the DTV domain and the extent to which they can accurately reflect a user's true preferences.

2 Implicit Ratings in the DTV Domain

In this research we use the Fischlar video library system [Lee *et ai,* 2000] developed by the Centre for Digital Video Processing at Dublin City University (www.cdvp.dcu.ie). Fischlar has been fully operational for a number of years and attracts hundreds of users on a regular basis. It is a video library system that allows users to record, playback and browse TV shows through a normal Web browser. Users browse personalized television schedules, provided by the PTVPlus personalization engine [Smyth and Cotter, 2000], to select programmes to be recorded by the system. TV shows are captured in MPEG-1 format and processed in order to index their content for storage and to extract key-frames for content-based browsing. Fischlar users can select programmes to record up to 2 days in advance of broadcast and can access previously recorded programmes for playback or browsing.

As Fischlar takes advantage of the PTVPlus personalized TV listings service, there exists the ability to capture explicit ratings feedback for its user profiles (Figure 1) using the rating icons beside each programme's description. However, unlike PTVPlus, Fischlar users can interact directly with programming content, and this provides new profiling opportunities. In particular, the record, playback and browsing actions of users are suitable as implicit interest indicators; e.g., selecting a programme to be recorded is surely a strong potential indicator of interest. Thus, Fischlar's profiling capabilities have been enhanced to capture implicit as well as the explicit profiles, by logging all record, playback and browsing actions (Figure 2).

USER #1762

Channels	BBC1.BBC2,..
+Progs	Friends, ER, ...
Keywords	Comedy, Soap
Times	PnmeTime

Figure 1: Explicit Profile Example.

USER #1762

Behaviour	Play
+Progs	Scrubs, Cheers,

Figure 2: Implicit Profile Example.

3 Experimental Evaluation

Our aim is to evaluate the usefulness of Fischlar's implicit user profiles relative to its explicit profiles. In particular we wish to determine whether or not these implicit profiles are any more or less accurate at predicting user preferences than the explicit profiles gathered directly from user ratings.

3.1 Test Data

Our experiments draw on 650 live-user profiles from the Fischlar system. In particular we make use of 5 different types of profile:

1. *Explicit* profiles contain explicit ratings that Fischlar users have provided;
2. *Play* profiles contain only those programmes that a user has played back;
3. *Record* profiles contain only those programmes that a user has recorded;
4. *Browse* profiles contains only those profiles that a user has browsed using Fischlar's key-frame browser;
5. *Combined* profiles are a combination of play, record and browse profiles.

3.2 Algorithms

In addition we make use of two different programme recommenders:

1. *CF* is a traditional collaborative filtering recommender as described in [Konstan *et al.,* 1997; O'Sullivan *et al.,* 2002b]; it finds similar users by calculating program overlap between the target user and all others, and then recommends to the target user programs from the most similar users that they have yet not rated.
2. *Sim* is a similarity-based recommender that draws on similarity information mined from user profiles as described in [O'Sullivan *et al.,* 2002b]. It uses data mining techniques to find relationship between items and then harnesses this information both in profile similarity and recommendation ranking (ranking all possible recommendable programmes) to recommend programmes to the user.

Due to space constraints, it is not possible to discuss these recommenders in detail; the interested reader is referred to [O'Sullivan *et al.,* 2002b; 2002a]. Suffice it to say that they each draw on a fundamentally different recommendation strategy, be it collaborative as in the case of CF, or content-based, as in Sim.

3.3 Method

First each profile is randomly split into a training profile and a test profile; this is repeated for a range of different split-ratios for each profile. The training profiles are used as the basis for recommendation; in each case we make 50 programme recommendations for each user. The quality of the resulting recommendations is evaluated with reference to the corresponding test profiles by computing the percentage of recommended programmes that appear in the target test profile. So, for example, we use the training play profiles to generate recommendations that are evaluated with respect to the test play profiles; similarly for record and browse profiles. In this way we can compare the quality of the recommendations generated using implicit play, record, browse and combined indicators to the quality of recommendations generated using explicit ratings profiles.

3.4 Results

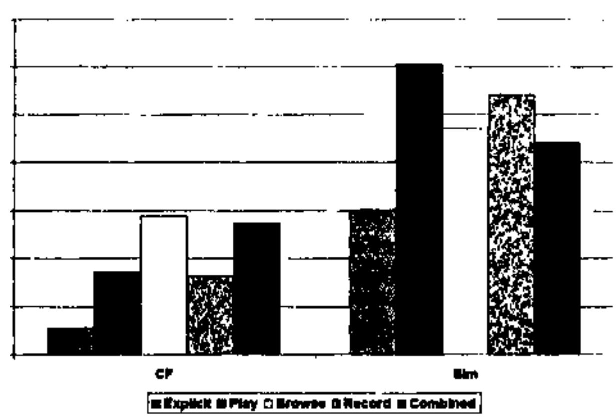

Figure 3: Recommendation accuracy for explicit and implicit (play, record, browse, combined) profile groups.

The results are presented in Figure 3 as an accuracy bargraph for the various different groups of explicit and implicit profiles (play, record, browse and combined). In each grouping we present the accuracy results obtained by the CF and Sim recommenders. For example, we find that the explicit profiles deliver recommendation accuracies of 2.5% and

15%, for CF and Sim recommenders, respectively; clearly the enhanced Sim rccommcndcr offers significant benefits over the standard CF recommender as previously discussed in [CTSullivan *et ai,* 2002b]. Each of the implicit profiles generate consistently higher recommendation accuracies for both CF and Sim recommenders. For instance, in the case of the play profiles, the recommendation accuracy is 9% and 30% for the CF and Sim recommenders, which is a significant increase on the explicit profile accuracy, especially for the Sim rccommcndcr. In fact, the implicit profiles yield recommendation accuracies in excess of 20% for the Sim recommender compared to the 15% obtained for the explicit profiles.

These results indicate that playback, recording, browsing and combined behaviours in Fischlar serve as competent interest indicators when it comes to profiling user preferences. In each case we find an increase in the quality of recommendations made from profiles containing implicit information. Moreover, the scale of this increase is largest in the case of the Sim recommender, due to the fact that this recommender directly mines the profiles in order to generate programme similarity knowledge as the basis for recommendations. The availability of the higher quality implicit profiles helps to improve the similarity rules used by Sim when compared to the similarity rules mined from explicit profiles.

4 Conclusions

We have evaluated the use of implicit interest indicators as the basis for user profiling in the DTV domain. Research in more traditional domains, such as Web browsing or Usenet News, indicates that some implicit interest indicators (e.g., read-time and mouse movements) are capable of serving as reliable alternatives to explicit profile information such as user ratings. Our research indicates that the same is also true in the DTV domain. We have compared the quality of recommendations generated from explicit and implicit profiles and found, through the use of a deployed video library system, that in each case the implicit profiles are capable of generating recommendations that are as good as those generated by the explicit profiles.

Acknowledgements

The support of the Informatics Research Initiative of Enterprise Ireland is gratefully acknowledged. The support and advice of the Centre for Digital Video Processing at Dublin City University is also gratefully acknowledged.

References

[Balabanovic and Shoham, 1997] M. Balabanovic and Y. Shoham. Fab: Content-Based, Collaborative Recommendation. *Communications of the ACM,* 40(3):66-72, March 1997.

[Claypool et al.,2001] M. Claypool, P. Le, M. Wased, and D. Brown. Implicit Interest Indicators. In *Intelligent User Interfaces,* pages 33-40, 2001.

[Goecks and Shavlik, 1999] J. Goecks and J Shavlik. Automatically Labelling Web Pages based on Normal User Interactions. In *Proceedings of the UCAI Workshop on Machine Learning for Information Filtering,* Stockholm, Sweden, July 1999.

[Goldberg *el ai,* 1992] D. Goldberg, D. Nichols, B. Oki, and D. Terry. Using Collaborative Filtering to Weave an Information Tapestry. *Communications of the ACM,* 35(12):61-70, December 1992.

[Good *et al,* 1999] N. Good, J. Schafer, J. Konstan, A. Borchers, B. Sarwar, J. Herlockcr, and J. Reidl. Combining Collaborative Filtering with Personal Agents for Better Recommendations. In *Proceedings of the Conference of the American Association of Artificial Intelligence (AAA 199),* Orlando, Florida, USA, July 1999.

[Hammond *et al,* 1996] K. Hammond, R. Burke, and K. Schmitt. A Case-Based Approach to Knowledge Navigation. In D.B Leake, editor, *Case-Based Reasoning Experiences Lessons and Future Directions,* pages 125-136. MIT Press, 1996.

[Konstan *et al,* 1997] J.A. Konstan, B.N. Miller, D. Maltz, J.L. Herlockcr, L.R. Gordon, and J. Ricdl. Grouplcns: Applying Collaborative Filtering to Usenet News. *Communications of the ACM,* 40(3):77-87, 1997.

[Lee *et al,* 2000] H. Lee, A. Smeaton, C. OToole, N. Murphy, S. Marlow, and N. O'Connor. The Fischlar Digital Video Recording, Analysis and Browsing System. In *Proceedings of RIAG'2000,* pages 1390-1399, 2000.

[Morita and Shinoda, 1994] M. Morita and Y. Shinoda. Information Filtering Based on User Behaviour Analysis and Best Match Text Retrieval. In *Proceedings of the 17th ACM Annual International Conference on Research and Development in Information Retrieval,* pages 272-281, 1994.

[CTSullivan et al., 2002a] D. CTSullivan, D. Wilson, and B. Smyth. Improving Case-Based Recommendation: A Collaborative Filtering Approach. In S. Craw and A. Precce, editors, *Proceedings of the 6th European Conference on Case Based Reasoning (ECCBR-02),* pages 278-291,2002.

LO'Sullivan et al., 2002b] D. CTSullivan, D. Wilson, and B. Smyth. Using Collaborative Filtering Data in Case-Based Recommendation. In *Proceedings of the 15th International FLAIRS Conference,* pages 121-125, 2002.

[Rafter et *al,* 2000] R. Rafter, K. Bradley, and B. Smyth. Personalised Retrieval for Online Recruitment Services. In *Proceedings of IRSG'2000,* 2000.

[Resnick *et ai,* 1994] P. Resnick, N. lacovou, M. Suchak, P. Bergstrom, and J. Riedl. Grouplens: An Open Architecture for Collaborative Filtering of Netnews. In *Proceedings of ACM CSCW'94,* pages 175-186, 1994.

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

[Smyth and Cotter, 2001] Barry Smyth and Paul Cotter. Personalized Electronic Programme Guides. *Artificial Intelligence Magazine,* 21 (2), 2001.