Want a Recommendation? Ask my Digital Personality!

Joost Broekens
Telematica Instituut, The Netherlands
Joost.Broekens@telin.nl

Abstract
Personalized content is an important goal for interactive television for many reasons. It enables better-targeted advertisements to individual users as well as television content that better matches the user's interest. Recommendation mechanisms often gather large amounts of user data of many users and the user is typically not in charge, and often not even aware of the gathering of this data. In this position paper we argue that within the context of iTV, this approach toward user modelling for recommendation is fundamentally wrong from a user perspective. We argue that it is better to approach recommendation from the perspective of a *private digital personality*. This is a representative agent of the user that can be used to ask, as in real life anyone would ask the user, what the user likes. This is a better understandable model for users, enables the user to be in full control of its preferences as well as who accesses these. As a result users will be much more likely to accept and see the benefits of personalized content.

1. Introduction

In an ideal world, personalization of content is useful and desired. It is useful for users as well as advertisers. It helps users to find more interesting content (e.g., television shows, products, Web2.0 content) at the right moment in time. It helps advertisers to better target products towards individual users. In essence one could argue that in this ideal world advertisement has shifted from a necessary evil to a much-wanted service. Users will get product advertisements that are interesting to them: advertisements are consumed and perceived as content. As a result, users are more likely to buy advertised products, lowering advertisement cost and thus increasing the amount of money that can be asked per user for viewing an ad. This means that even niche content can benefit from advertisement by selling advertisement slots. This opens up new funding opportunities for content production. All in all, personalization is a useful addition to interactive television.
However, in this ideal world, users don't care about privacy, and malicious use of personal data does not exist. Unfortunately, this world does not exist. Users do and should care a great deal about how and by whom their personal data and preferences are used (see also, Resnick & Varian, 1997). This means that personalization potential is critically dependent on a user's willingness to give away personal information, not on how technically advanced recommendation algorithms might get. Currently, many recommendation techniques are based upon gathering user-behaviour related data for a large set of users. For example, to recommend products based on collaborative filtering techniques, i.e., recommending content to one user based on that user's similarity to other users (Goldberg et al., 1992), this kind of data is critical. Virtually all recommendation services store large amounts of user-behaviour related data, and the user has no insight in, let alone control over this data. This is going to present a problem in the near future, especially for interactive TV (iTV). If the goal of iTV is that it is to be launched for a wide public, then there are many persons who do not trust the idea of having a "big brother" they do not know gathering their behavioural data and storing that data somewhere they do not have access to. Even if mechanisms are available to secure the data, personalization and recommendation have to be marketed as being in control of the user. Almost everyone has behaviour he or she does not want others to know about, including their iTV service provider. Users will want to be in charge of who accesses what personal data.
Here we argue that unless personalization and recommendation is approached and marketed as a *private digital personality* at whom recommendation requests can be addressed, many users will not accept the technology.

2. Recommendation and Digital Personalities

First, we have to consider the main goal of recommendation. Recommendation aims at providing more interesting content to a user delivered in a timelier manner. This helps users to find interesting content as well as helps to reduce information overload (van Setten, 2005; for a longer list of recommender usage see Herlocker et al., 2004). Recommendation is not a tool to sell more products per se, although it can be used to do so when products are viewed as content items. This

means that a user must have a genuine interest in the product.

In the real world, you would just ask the user if he or she wants the product. The user can then (a) decide if he/she tells you anything at all, and (b) decide how much he/she likes the product.

In the current recommendation world, no-one asks the user anything. A service provider (e.g., Amazon) magically knows what the user likes based on previous buying and clicking behaviour. In the case of Amazon, this is not really a problem, as the user also trusted buying products at Amazon in the first place, and the data gathered only consists of the buying and clicking behaviour of that user at the Amazon website. However, if Amazon decides to offer its recommendation engine as a service to other web stores, or if Amazon can request recommendations from a user's iTV provider, this will generate privacy-related problems. In other words, recommendation cannot be a generic service, while from an iTV point of view this is needed. ITV is a multi-medial experience across a potentially wide set of content domains. It is a very good channel to gather vast amounts of user data to construct detailed user preference models. These models can be used for recommendation to third parties, providing an interesting market opportunity. If this channel can not be used, because a user does not accept the fact that data is gathered, this would be a huge loss.

As mentioned above, users are familiar with being asked what they like. Their friends and family members do this all the time. Telephone-based marketing campaigns also do so, and the user still has the freedom not to give away the information. In many societies, persons see themselves as active individuals with a private set of goals, ideas, intentions and responsibilities. This agent-based model does not map well to the current state of the art recommender systems: i.e., monolithic multi-user data sets containing private information of everyone. The problem is not so much that this data can not be secured and shielded, the problem is the user's perception of the model: all user data together, single point of failure, big brother, insurance companies potentially using the data, cross-user comparisons, cross-media comparison, insight in my private life, etc…

The perception of recommendation should map better to how users see themselves. One way to do this is to approach recommendation with the concept of a *private digital personality*. It is a digital representative of the user that gathers the user's behaviour, infers preferences, alerts the user to interesting stuff, etc. Key is that the user and the digital personality are the same entity, the digital personality in the virtual world, the user in the real world. From a user perspective this has several advantages. First, the digital personality would not do something the user would not do: i.e., it would not give away private information to an unknown party; it would not allow comparisons between the user and other persons except the user's friends, etc…As a result, the user is in charge, not the user's iTV provider, visited web shops, or any other party. Second, a user can decide to not have a digital personality, but still benefit from other iTV functions. Third, when recommendations are asked by other parties this is visible to the user.

From a content provider point of view, such a digital personality is an interesting service to use: it gathers a large amount of data about the user and it becomes interesting to have access to that data. As content can be better targeted, providers might offer discounts to users that allow their digital personality to answer recommendation requests. Furthermore, recommendation techniques that compare users and as a thus need to access multiple profiles are still feasible, safe and acceptable using Identity Management (IdM) techniques (e.g., Liberty Alliance. http://www.projectliberty.org/) that ensure the users privacy and/or anonymity.

3. Conclusion

We have argued that users are more likely to accept their behavioural data being gathered for content recommendation if recommendation is seen as a service that is provided by the user for the user. We called this service a *private digital personality*. Key elements in the argument are:(a) if recommendation techniques in iTV are to be broadly accepted, users must perceive recommendation as something safe and something they are in charge of, and (b) it is better to match the concept of personalisation to the way individuals see themselves: i.e., acting agents with private goals, intentions and beliefs.

References

Goldberg, D., Nichols, D., Oki, B. M., and Terry, D. (1992). Using collaborative filtering to weave an information tapestry. *Communications of the ACM, 35* (2), 61-70.

Herlocker, J., L., Konstan, J. A., Terveen, L., G., & Riedl, J. T. (2004). Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems, 22* (1), 5-53.

Resnick, P., & Varian, H. R. (1997). Recommender Systems. *Communications of the ACM, 40*(3), 56-58.

van Setten, M. (2005). *Supporting People in Finding Information. Hybrid recommender systems and goal-based structuring.* Doctoral dissertation, University of Twente, The Netherlands.