Explaining Recommendations through User Groups

Jarmo Laaksolahti
Social Computing Group
Swedish Institute of Computer Science
Isafjordsgatan 22
Box 1263, SE-164 29 Kista
Sweden
+46 8 633 1524
jarmo@sics.se

ABSTRACT

Recommender systems are a specific type of information filtering tools that have emerged in recent years. Until now research in recommender systems has been focused on developing algorithms for collaborative filtering, little effort has gone into considering how users understand recommendations. We describe a system that visualizes the structure of the user population thereby making it easier to understand what recommendations are based on and thus anticipate the effect of your own actions

Keywords

Recommender systems, collaborative filtering, user groups, visualization

1. INTRODUCTION

A specific type of information filtering tool that has emerged in recent years are recommender systems. Recommender systems are aiming to help users find relevant information within a domain by giving recommendations as to what the user should look at. The recommendations are based on what users with similar interests previously have found to be relevant information.

Until now research in collaborative filtering has been focused on developing algorithms for collaborative filtering and tuning the performance of collaborative filtering systems in various ways. The underlying research question has been: How do we make collaborative filtering work? Relatively little effort has yet gone into considering how users actually interact with recommender systems, how users understand recommendations or in fact how users actually want to use recommender systems. We believe that these are important aspects of recommender system design that deserve further investigation, since in the end it is the way in which a user perceives and uses a system that decides if it is successful or not.

2. EXPLAINING RECOMENDATIONS

In most current recommender systems users are required to first rate a set of items previously known to them before they can start getting recommendations. The recommended items are usually presented as a simple list or similar structure. Usually the items are ordered so that items that are more likely to appeal to the user are at the top of the list. Sometimes the 'rating', a number indicating how much the system believes the user will like a particular item, is also shown. A typical example of this kind of system is MovieLens (http://movielens.umn.edu) that recommends movies.

What is lacking in this kind of system is an explanation of how the system has come up with a recommendation or why the system believes it would be interesting to a user.

There are at least two ways of tackling this challenge: using statistical measures explaining the mathematics and/or process behind recommendations or visualizing the structure of the underlying user population and/or collection of items in the recommender system.

Herlocker (Herlocker, 1999) who has explored the first approach, divide methods belonging to this group into three broad categories: Data explorative (explaining the data that underlie a recommendation), process explorative (explaining the mathematical process behind a recommendation) and argumentative (using logical argument techniques to support a recommendation).

The second approach is less well explored although recently some researchers have started to look at this possibility. For example Tatemura (Tatemura, 2000) describes a system which uses what he calls ‘virtual reviewers’ to get recommendations that are based on a subset of the user population which can be dynamically chosen at runtime. The question is however if this can truly be considered a visualisation of the underlying user population since the ‘virtual reviewers’ are entities created by the user and not necessarily reflected in the actual structure of the user population.

What is striking when looking at most recommender systems is that although recommendations are based on users activities (ratings as well as other activities), those activities and the consequences of those activities are not visible.

3. AN ON-LINE FOOD STORE

We have implemented an on-line food store (Svensson et al, 2000) that visualizes the user population. In this store users shop groceries by choosing recipes they like. The recipes are then transformed to a shopping list. The recipe space is divided into groups that manifest a certain type of recipes or rather a group of people that like a certain type of recipes. Users can create their own groups or rely on the system editor to analyse usage logs to find patterns that can be the basis for new groups.

Recipes within a group are ranked according to the *profile* for the group, which in turn is affected by the recipes chosen in that group.

The recipe groups are visualized in a map giving an overview of the group space (see Figure 1). By visualising the user groups users get the possibility to browse the different groups, find the one(s) s/he likes the most, and hopefully get a better sense of how the system works for several reasons:

1. Users can see which recipes have contributed to the group profile.
2. Users can see which user groups are represented in the system, thereby giving them a chance to decide where they best belong.
3. Users can see the effects of their own actions. Recipes get a different ranking in the ranking list.

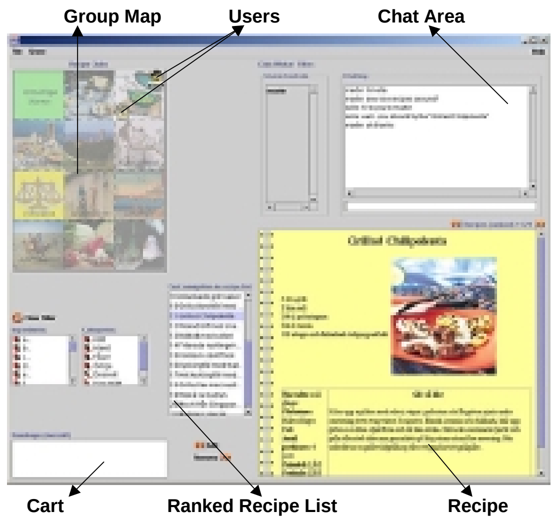

Figure 1. The on-line food store

We believe that seeing which groups are available and which recipes are popular in them can help the users build a better mental model of the recommender system and the information that underlies its’ operation. The aggregated user information is in this way also integrated into the users decision process and can hopefully give a higher level of trust in the recommender system by making it easier to make the right choices from the start.

An initial small-scale study indicates that users understood the effects of their actions to a high degree when using the system. When asked if they believed that their choice of recipes affected the recipe groups in any way, over 50% answered they believed the ranking of the recipes was affected.

4. REFERENCES

Herlocker, John. (1999) Explanations in Recommender Systems. *Workshop on Interacting with Recommender Systems, CHI99*.

Tatemura, Junichi. (2000) Virtual Reviewers for Collaborative Exploration of Movie Reviews. *Proceedings of IUI2000***,** 272-275, ACM Press.

Svensson, M, Laaksolahti, J, Höök, K, Waern, A. (2000) An On-Line Food Store. *Proceedings of IUI2000***,** 260-263, ACM Press