A Recipe Based On-line Food Store

Martin Svensson
Stockholm University/KTH
Department of Computer and Systems Sciences
Box 1263
SE-164 29 Kista
+46 8 6331500
martins@sics.se

Jarmo Laaksolahti, Kristina Höök, Annika Waern
Swedish Institute of Computer Science
Box 1263
SE-164 29 Kista
+46 8 6331500
{jarmo, annika, kia}@sics.se

ABSTRACT

Recent research in the area of information retrieval hypothesizes that people benefit from social clues, so called social navigation, when they try to navigate information spaces [7]. We have designed an on-line grocery store building upon those ideas manifested in several different ways. The most central feature is that the system uses a combination of content-based and collaborative filtering as the basis for recipe recommendations. This filtering process can in turn be controlled by editors, whose role is to control the content of the "recipe clubs". Other types of social clues are also present, such as displaying how many users that have chosen a recipe. Finally, the system shows information about other users currently present in the system, and allows users to get in direct contact through chat.

Keywords

Recommender system, social navigation, user groups, collaborative filtering, content-based filtering, on-line shopping.

1. INTRODUCTION

We find two problems interesting: one is the problem of how to aid users to navigate large information spaces, and the other is how to not only aid them, but also turn the experience into a pleasurable one?

Users today have to deal with large information spaces – on the web, through email, ICQ-messages, on-line calendars, through phones, mobile phones, etc. Not only do we have to deal with passive information spaces, but also vast numbers of different services available over the net. We know that some users find it more difficult than others to navigate and use these spaces [2]. The question is how to aid all users to find the most relevant information/service/functionality?

Since their inception, computers have been considered primarily as *tools*, making tasks more efficient. In recent years however, computers have entered the field of entertainment and art, not aiming at supporting "serious" work, but to create pleasurable, fun, social or aesthetic experiences in their own right [8]. Chat environments, games, digital art, and hypertext narratives are all examples of this shift from *work* to *pleasure*. The computer is a "place" where these two dimensions of life come together and possibly interact. Recent studies have aptly pointed out the tight connection between cognition and more affective dimensions [9]. Being in, and navigating through, information space is of course influenced by "rational" decisions and reasoning, but this side is more or less always accompanied by a general overall experience of the browsing, with distinct emotional and affective features. Frustration, anxiety, or "flow" are just some examples of such features.

With these considerations in mind, how can we design systems that better support users, than the ones that we see today? One way to tackle the problems mentioned above is to introduce the notion of *social navigation* [3, 7, 13]. In social navigation we let users (instead of the system) help each other in various ways. For instance, imagine surfing a web site and seeing all the people in it and being able to communicate with them. This will create a sense of not being alone in the space, in effect, reducing the feeling of being lost and the anxiety users feel when navigating large information spaces, such as, the WWW. Another example of social navigation could be in a more indirect way, for example, following trails of people in a web site or getting information filtered based on what other people (similar to you) think.

By making other users' actions visible we can take advantage of the work they have done to find their way around and to solve problems. By information space, we mean anything from the interface to a normal application to large hypermedia spaces such as the World Wide Web or virtual reality environments. Users' actions can be made visible in various ways: through *direct social navigation* (talking to or seeing individual users act), *indirect social navigation* (seeing the aggregated user behaviour as in recommender system advice), or *read wear* (seeing how an object has been used by other users through its texture).

Social navigation will not necessarily mean more *efficient* interfaces. What is gained by social navigation might not be, and maybe should not be, time and efficiency, but instead it might contribute to other factors. Maybe a better question to ask is how do we know that we have created a good navigational experience? Will it be a matter of more aesthetic or emotional factors, such as feelings of flow or having a delightful experience, as opposed to the efficiency measurements usually taken for the prevailing tool-based usability evaluations?

Social navigation seem to be a natural approach to the design of an information space; yet we still have not seen many practical solutions that allow users to behave socially, interfaces that allow

for the accumulation of social trails, or the aggregation of user behaviours. In this paper we present an on-line food store that draws upon the ideas of social navigation. The aim is to present a store that makes the shopping experience socially richer, but also to make it more efficient in terms of getting the food you want.

2. ON-LINE FOOD SHOPPING

The existing food stores on-line are all "dead" spaces where users fill in how many milk packages, etc. they want sent to their doorstep. It is clear that on-line stores are still viewed as tools, in the traditional computer science way, and not as places where people shop, meet friends, socialise etc. Taking into account these human aspects of shopping becomes especially important considering people that do not have the opportunity to go shopping regularly, e.g. people with low mobility. In a study by Richmond [11] on shopping in a VR environment, it was found that the users also want to be able to access the social aspects of a physical store, they want to socialise with other people and have a multi-user experience.

How can the ideas of social navigation be made central and be used to inform design? One trail that we can follow is to recommend recipes using collaborative filtering techniques [10]. Recipes are interesting accumulated pieces of knowledge in this context. Through which recipes we cook from we convey a lot about our personality, which culture we belong to, our habits, etc. Making recommendations on what food to buy based on recommending recipes is an interesting functionality in itself. Imagine that we on top of that add accumulation of user behaviour so that we understand which groups are most likely to choose which recipes. We have designed such a system that works as follows. As a (by the system) known user logs onto the system, it will present a recommended recipe. This recipe is the most downloaded recipe at that point in time for the category of users that this user belongs to. The user can add the recipe to her shopping basket, which in turn adds the ingredients from the recipe to the list of items that will be delivered. The user can then ask for the next-best recipes that fit her category of users - much along the same lines as Amazon.com recommendations ("other people who bought this book also bought these books"). The recommended recipe will be chosen on the basis of three different characteristics that the user can manipulate: user groups, the category of food (Italian, Thai, etc.), and any particular ingredient that should be included (shrimps, beef, etc).

The rest of this paper will describe our on-line food store in more detail. We start off by taking a closer look at how our recommender works, and then move on to discuss additional social clues available in the system. The system presented in the paper has been implemented as an on-line service. It is currently under evaluation through a series of experiments, focusing in particular on the effects of social navigation on usage.

3. THE RECOMMENDER SYSTEM

Users can select recipes by restricting search in three different ways: by selecting specific ingredients (e.g. pork), ingredient categories (e.g. meat), and recipes liked by a particular user group (e.g. meat-lovers). The user groups known to the system are called "clubs": they are formed by and given their names by "club owners", and recommend recipes that users/members of the club like. Information on categories and ingredients already exist in the recipe store, e.g. it can be extracted from the recipe database. There are several advantages to use both filtering techniques, the most obvious one being when *bootstrapping* the system. What usually happens when a user starts using a collaborative filtering system is that she gets poor recommendations, since the system does not know anything about her. As the user goes along, rating pieces of information, the system becomes better. However, with our approach a user can immediately get good recommendations, since she has the ability filter recipes by categories and ingredients as well as selecting among the available recipe clubs. The other main advantage is the amount of information that is available to us when we try to cluster users. Since we do not only look at recipe names (but also categories and ingredients) when clustering users it should be both easier and faster to find clusters of users, that is, we can find users with similar tastes even if they do not have a single recipe in common[1]. For example, if two users consistently choose recipes without red meat they will be clustered in the same group even if none of the chosen recipes are the same for the those two users.

3.1 Labelling User Groups

A common problem with existing recommender systems, such as, GroupLens [5], Firefly [12], and Phoakes [14], is that they give little or no feedback to a user on what user group a user belongs to, or what user groups a recommendation is built upon. Since recommender systems make their recommendations based on what other similar users have done in the past we believe that this is a very important piece of information that should be provided to the user. The problem is of course the rather complex task of automating the "labelling" of user groups. For instance, it would be extremely difficult for the Firefly system to label a cluster of users as "reggae lovers with a flavour of ska". However, if this could be done we would get a much richer recommender system.

Labelling of user groups not only tells something about the user's own group, but also gives information about other user groups. This will allow a user not only to navigate from the highest to the lowest ranked piece of information, but also between groups of users. In the recipe domain this seems like a sensible idea; a recipe that is rather low ranked because the user is classified as a "meat lover" can still be the recipe to choose since it is highly ranked for "thai food lovers". In this way, a user can try out being a "thai food lover" for a while. It is of course the case that a user's group is in no way static; if a user consistently chooses recipes based on what vegetarians like, s/he will gradually move towards the vegetarian user group.

3.2 The Role of the Editor

Our solution to the labelling problem is to put an editor back into the loop. We have two types of editors, one that investigates actual user groups and the so called "user editor". The first editor will look at the clusters of users (based on which recipes they have chosen) and "name" them with fuzzy names that convey somewhat of their content: "vegetarians", "light food eaters", "spice lovers", etc. To enhance this process, we base the system on a recommender algorithm that gives an *explicit* representation of user preferences as user and recipe profiles that change over time. It should thus be rather simple for an editor to find

[1] We are presently in the process of evaluating the recommender algorithm used in the system, this being one of the central evaluation criteria.

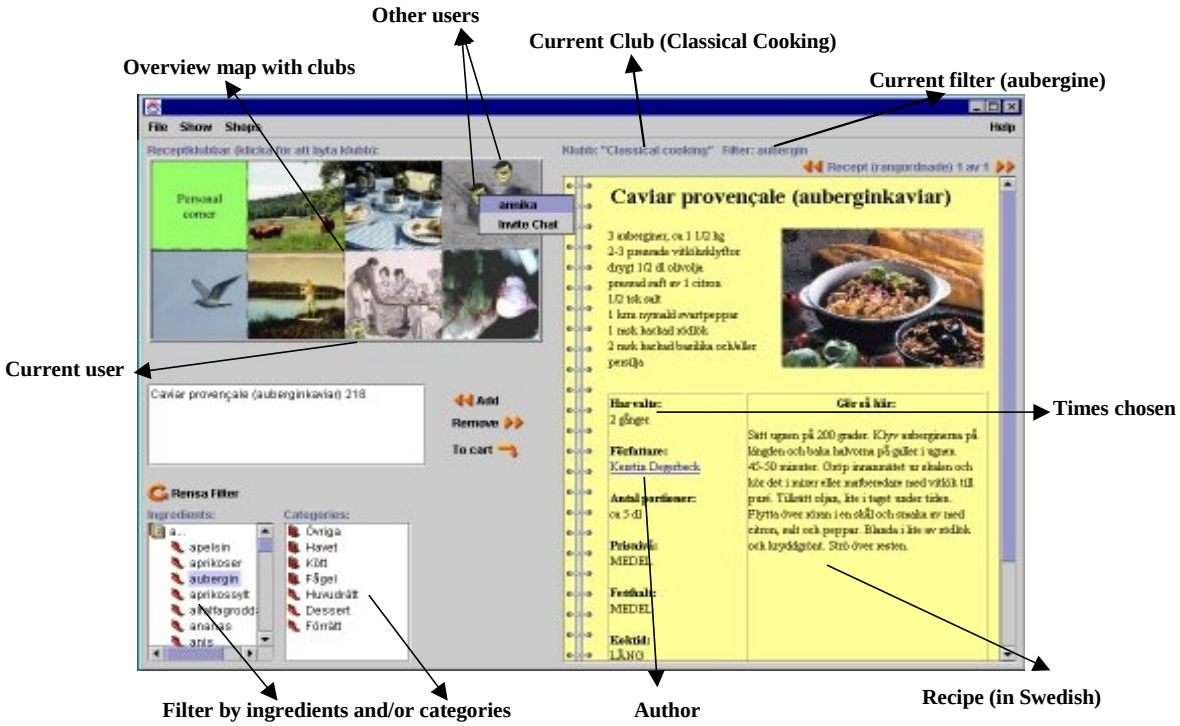

Figure 1. The on-line shop interface

similarities between users or recipes, and get an intuitive impression of why they are similar. The second type of editor is a user of the system, she can at any time create a new recipe club with certain constraints that she finds interesting, for instance, "Annika and her friends club". These clubs, obviously, do not have to reflect actual clusters of users. However, if a group of users choose recipes from one of these clubs on a regular basis their user profiles will converge.

In the designed system, we have many editors: each club owner is an editor for her clubs. There can also be a staff of editors for the recipe system as a whole, who can look for interesting clusters of users and initiate new clubs, supervise the quality of clubs, and remove clubs that are no longer in use. In the system we also introduce a special purpose club called the "personal club". The personal club reflects an individual user and is only affected by that particular user's actions. The idea is that a user can always go to her own club and get recommendations based only on her own actions, i.e. not influenced from any specialised club.

To bootstrap our system we have created a set of initial recipe clubs (user groups), for example, "The Italian Club" and "Classical Cooking". The idea is that users will start by using these clubs and as they move along create new clubs that more accurately reflect their interests, i.e. we want the users to dynamically change the system.

Our hypothesis is that visual recipe clubs will provide the users with more insight into the social trails of their own actions as well as other users' actions that have lead to the recommendations they finally get. It also provides some insight into the inner workings of the recommender system.

4. ADDITIONAL SOCIAL CLUES

The system provides additional social clues to recipes. This section summarises these briefly.

Macaulay [6] investigated how journalists find information. What was interesting is that they do not search for information but rather for sources of information, i.e. the source is more important than the information itself. We believe that the same holds for the recipe domain. That is, it is often the case that one consults a trusted source to get new recipes, for example, a cookbook, a family member, etc. We therefore put links on every recipe to other recipes by the same author (source).

A form of read wear [4] is also added to each recipe. This is the information about how many users that previously has selected a particular recipe.

A key issue in social navigation is awareness of each other or as Ackerman and Starr [1] put it: "people attract people". Based on the user groups that are known to the system we create an overview map of the on-line store. Actually users can choose which user groups (recipe clubs) that should be visual. In a scenario were there are a hundred recipe clubs this can be crucial. Whenever a user enters the system she will appear as an icon in the overview map within her own personal user group. Users then have the ability to talk to each other via chat and move to other recipe clubs by clicking the overview map, see Figure 1. We believe that this functionality will serve two purposes. First of all, seeing other users will make a user more relaxed and can also make her stay longer in the store since there are other people there. From our initial pilot study of the system we see that people actually follow each other, i.e. people tend to move to the recipe clubs that other people are in, or as one of our subjects put it: "I

visited the recipe clubs where there were a lot of people". Secondly, when users share a very special interest (e.g. allergies) the likelihood of them taking advice from one another will increase significantly.

5. CONCLUDING REMARKS

In this paper we have presented a new approach to building an online store based on recipe recommendations. The key feature that we want to stress is the use of names for different clusters of users. When the user understands what sort of user the system classifies her as, it becomes easier to understand why some information gets high ratings and some information gets low ratings. Also when the recommender system allows a user to navigate among clusters of users new and interesting ways of choosing recipes will arise. Finally we are changing the way people do on-line shopping of food from shopping groceries to shopping recipes.

6. REFERENCES

[1] Ackerman, M, Starr, B. Social activity indicators interface components for CSCW systems, *User interface and software technology*, 159, 1995.

[2] Nils Dahlbäck, Kristina Höök, and Marie Sjölinder. Spatial Cognition in the Mind and in the World - the case of hypermedia navigation, The Eighteenth Annual Meeting of the Cognitive Science Society, University of California, San Diego, July, 1996.

[3] Dourish, P, Chalmers, M. Running Out of Space: Models of Information Navigation, short paper, *HCI'94*, Glasgow, August 1994.

[4] Hill, W, Hollan, J, Wroblewski, D, McCandless, T. Edit wear and read wear, *Human factors in computing systems*, (1992), 3-9.

[5] Konstan, J, Miller, B, Maltz, D, Herlocker, J, Gordon, L, Riedl, J. GroupLens applying collaborative filtering to Usenet news, *Commun. ACM* 40(3), (1997), 77-87.

[6] Macaulay, C. Information Navigation in The Palimpsest, In deliverable 2.1.1 of the PERSONA project, SICS, 1998.

[7] Munro, A.J., Höök, K. & Benyon, D.R. (eds.). Social Navigation of Information Space, Springer Verlag, 1999.

[8] Murray, J. Hamlet on the Holodeck. The Future of Narrative in Cyberspace, New York: The Free Press, 1997.

[9] Picard, R. *Affective Computing*, Cambridge: The MIT Press, 1997.

[10] Resnick, P, and Varian, H. Recommender Systems, *Communications of the ACM*, 40(3), 1997.

[11] Richmond, A. Enticing online shoppers to buy - A human behavior study, Computer Networks and ISDN Systems, 28, (1996), pp. 1469-1480.

[12] Shardanand, U, Maes, P. Social information filtering algorithms for automating "word of mouth", Human factors in computing systems, (1995), 210-217.

[13] Svensson, M. Social Navigation, In Dahlbäck (ed.) Exploring Navigation: Towards a Framework for Design and Evaluation in Electronic Spaces, SICS, 1998.

[14] Terveen, L. G., Hill, W, Amento, B, McDonald, D, Creter, J. Building Task-Specific Interfaces to High Volume Conversational Data, *Human factors in computing systems*, (1997), 226-233.