GroupMark: A WWW Recommender System Combining Collaborative and Information Filtering

Duncan Pemberton and Tom Rodden

Department of Computing
University of Lancaster
Lancaster, LA1 4YR, England
Tel: +44 (0)1524 593479
Fax: +44 (0)1524 593608
Email: pemberto|tom@comp.lancs.ac.uk

Rob Procter

Division of Informatics
University of Edinburgh
Edinburgh, EH1 1HN, Scotland
Tel: +44 (0)131 650 2707
Fax: +44 (0)131 650 6513
Email: rnp@dai.ed.ac.uk

Abstract. The objective of the SELECT project is to help Internet users find the most reliable, valuable, important and interesting information quickly and easily, hence reducing information overload. In these ways, SELECT will make a positive contribution to the problem of helping users tailor their information environments to meet their individual needs. The approach adopted in SELECT is to develop a general architecture for information filtering and recommendation systems, and to use this to implement and evaluate different strategies and techniques. In this paper we describe GroupMark, a prototype of a SELECT-based social recommendation tool for the WWW that is based upon shared bookmarks. We focus in particular on how GroupMark seeks to combine content-based and collaborative filtering techniques, and on the user interface issues raised by recommendation tools: i.e., the mechanisms for controlling behaviour and the visualisation of results.

1. INTRODUCTION

As information spaces such as the WWW grow ever larger, the need for tools to help users find high quality reliable information quickly and easily becomes ever more acute. Also the need to tailor information environments to specific user requirements is important for users. One possible solution to these problems is to employ recommender systems. SELECT is a project funded under the EU Telematics Applications Programme whose overall goal is to develop a general architecture to support the development and evaluation of web-based recommender systems [Procter 99].

Recommender systems are examples of adaptive filters that use inferences drawn from users' known behaviour to recommend documents they have not yet seen. There are two basic strategies to solving the problem of how to generate recommendations. The first strategy, which is sometimes referred to as *information filtering* (IF) [Schafer 99], is to derive recommendations for a particular user from knowledge of that user's past behaviour alone. For example, an IF system would recommend that a user should visit a WWW page because it has established that this WWW page matches the user's declared interest profile, or is similar to WWW pages that the user has already seen and (possibly) expressed a liking for. The second strategy, which is usually referred to as *collaborative filtering* (CF), or social recommendation, derives recommendations using the behaviour of others, especially those that have displayed similar tastes and interests in the past. For example, a CF system would recommend a WWW page to one user because other users who are known to have similar tastes to that user (i.e., the first user's peer group) are also known to like that page. CF-based systems are generally accepted as being more powerful than IF-based systems because they are capable of finding relevant documents that may be quite different from those the user has already seen [Herlocker 99b, Resnick 97]. They also represent an important a way of bringing so-

cial affordances into digital information environments: as various researchers have argued (e.g., [Gross 98, Procter 97, Twidale 96]), their users will be impoverished if digital information environments follow the misguided perception that information seeking in the real world is a solitary activity. However, recent work suggests that approaches combining both IF and CF techniques can produce even better results than either technique employed by itself [Balabanovic 97, Schafer 99, Tuzhilin 99].

In this paper we outline GroupMark, a prototype recommendation system that combines IF and CF approaches for recommending WWW pages. We have implemented the GroupMark system on top of the SELECT architecture as a demonstrator and proof of concept. Specifically, we describe its operational principles and the mechanisms provided for users to interact with GroupMark so that they can more easily tailor its behaviour to suit their needs, and for visualising the results.

The behaviour of many CF systems is difficult for users to understand – and hence to control – because they are effectively black boxes encoding complex relationships between their inputs and outputs [Herlocker 99a]. Often, for example, the user has no control over what inputs are used to compute recommendations. One of our aims in developing the GroupMark prototype is to explore ways in which these relationships can be made more explicit and can be manipulated by users. As with more familiar kinds of decision aids, such as expert systems, providing recommender systems with the capacity to be accountable to their users for their behaviour may be important for user acceptance and trust. It is also important if users are to be able to refine their use of the system.

2. AN OUTLINE OF GROUPMARK

The aim of collaborative filtering is to recruit others to act as our filtering agents on the assumption that they are our peers, i.e., like us in tastes and judgement of quality. This we endeavour to ensure by comparing opinions over a set of known documents.

There are a number of CF systems for the WWW that are based upon sharing bookmarks (e.g., [Bouthors 99, Glance 98]) providing users with an adaptive filtering function on the bookmark pool. Bookmarks are not only the unit of sharing currency in such systems, they also can be used to provide evidence that users are similar and so are good candidates for sharing recommendations. Bookmarks can therefore be seen as providing reliable and high-value evidence of peoples' tastes, and unlike rating-based approaches, do not require users to perform additional explicit and perhaps fine-grained -- and thereby difficult and effortful -- judgements of value [Hill 95]. At its simplest, GroupMark can be used solely as an implicitly based (i.e., 'zero input') recommendation system, with no more effort required than the user would make anyway when creating a bookmark. To get more value out of GroupMark, users need to invest more effort, which is a sensitive issue for recommender systems [Kushmerick 2000]. However, our aim in the design of GroupMark has been to make the relationship between effort and return distinctive and clear to users. It is because this relationship is clear that users are able to control the behaviour of the system more effectively.

Figure 1: GroupMark Main Window.

For these several reasons it was decided to follow a bookmark-based approach for GroupMark. In GroupMark, peer recommender group membership is defined by the comparison of individual user bookmarking behaviour. So, membership of the peer recommender group may be defined as: "A set of users who are deemed to be similar because their bookmarks (partially) intersect." As can be seen from Figure 1, these peer groups in GroupMark can be associated with an individual user (such as **Tom's Group** above, i.e., the interests of this particular user **Pemberto** have been determined to closely match the interests of **Tom** as expressed through specific group profiles that he has defined and owns). Alternatively, peer groups can be associated with a particular interest, i.e., people interested in gardening, for example. The peer group itself can be open or closed in terms of group membership. The group's owner sets this preference. A closed group is only available for subscription if the user is invited to join by the owner. This is the case with **Tom's Group**. Open groups are

controlled by member interest similarity alone, i.e., if you fall within the matching criterion, you may choose to join.

2.1 Social Recommendations Through Shared Bookmarks

Many tools that are currently found on the WWW that deal with recommendations based on the action of sharing bookmarks do not sufficiently exploit the idea of extending the pooled bookmark concept to produce targeted recommendations. The collaborative filtering process can be simply performed by the filtering of bookmarks, their associated URLs and keywords based on the information content of the WWW pages to which they point.

The existing central goal of consolidated bookmark repositories (the more common of which are detailed below) is to simply make a user's personal bookmarks available when they move to another physical machine. This is essentially "teleporting" the browser's client side bookmarking mechanism onto a server to enable roaming. Such systems therefore bypass the WWW browser's in-built bookmark or 'favorites' list facilities. In some cases bookmark repository systems such as Groupfire (http://www.groupfire.com/) extend this concept further by making bookmarks shareable with other users of the mutual repository. Groupfire does this by allowing the user to select if they wish to make certain bookmarks public or keep them as a private resource not shareable by others. The CSCW3 system [Gross 98] provides a similar facility for bookmark sharing, plus numerous others, including a facility for making history lists persistent and shareable.

As we show in this paper, much more targeted information may be extracted from pooling bookmarks and using these to deduce personal and shared user interests. For example, such a tool can guide users to interesting, highly rated, and relevant resources. This is because we surmise that users will only bookmark pages that they consider useful and, at some later point, return to. Low quality or pages with little personal user interest will generally be ignored in terms of bookmarking actions within WWW browser software.

The Coolsync system (http://bookmarks.coolsync.com/) adds further features such as lists of top bookmarked sites, and a "ring" of users with similar interests, although these groups are not deduced from your bookmark list by the system and personally recommended to you as in the SELECT GroupMark system. To join a "ring", a user must be invited. In GroupMark we extend this concept by having public recommender groups where users are free to join if their interests match, or in someway overlap with those described by the group owner. GroupMark also supports private circles where the group is "closed" and users can only subscribe based on the action of the group owner inviting the user into the circle. Other systems such as BackFlip (http://www.backflip.com/) and BestBookmarks (http://www.bestbookmarks.com/) simply act as a mechanism to centrally browse, store, and share a user's published bookmark lists. In summary, these systems provide very little reasoning about personal user interests, and in no way try to match up such interests with those of others.

GroupMark uses pooled bookmarks to reason about user interests via recommender groups and their associated group profiles. The addition of profile definitions in GroupMark allows the group owner to accurately match other users' interests to his/her recommender group description. This helps to ensure that GroupMark will only recommend a particular user to an interest group where

there is a reasonable amount of certainty that it holds some interest to that user. The group owner controls the membership by defining a group membership profile that users must be able to match in order to join the group (see Figure 2). As we explain below, the group owner defines the profile through example, identifying a set of URLs that other users must also have bookmarked. If they share a greater subset of the bookmarks, then it is highly probable that they have a similar interest to the group's originator and hence the set of recommendations that are identified with that particular group. Group recommendations can themselves be ranked in terms of potential interest. For example, if there is a very high profile match then GroupMark can display an appropriate message such as, "Essential References".

We argue that this additional filtering via recommender groups is an essential part of a shared book-marking system, and, in particular, that it is critical to ensuring that users are only recommended pages where there is a high probability that they contain interesting content. In general, systems that allow the users to browse other users' public bookmark recommendations do not sufficiently filter the information according to interest in any way. This can result in the user wading through lists of potentially uninteresting information that also makes it difficult to extract any relevant information that may be there. In other words, these systems have focused on the realisation of collaborative filtering techniques and have failed to consider how their behaviour might be tempered and improved through the incorporation of information filtering.

The approach of defining profiles through example documents is similar to Zloof's [Zloof 75] concept of 'query by example', and to Dix and Patrick's [Dix 94] more recent refinement of 'query by browsing'. One of the advantages of the approach is that, by being able to control which inputs (the user's own recommendations as represented by her bookmarks) are used to generate GroupMark's outputs (the recommendations of other users), the user is able to understand and explore more easily the input-output relationship, and so can learn how to control GroupMark's behaviour more quickly.

2.2 Implementation

GroupMark is built on top of the SELECT architecture [Procter 99]. The system is implemented as a SELECT agent with each browser bookmark mapped onto a SELECT page rating. The ratings are stored in the SELECT database and the GroupMark agent is used to extract recommendations from the SELECT relational database tables according to the user's previous bookmark behaviour.

The agent based recommender approach taken in SELECT allows the development of various thin-client recommender systems to be built on top of SELECT. These agents may then implement different recommendation algorithms with GroupMark being an example of one of these.

3. CONTROLLING RECOMMENDER TOOL BEHAVIOUR

Our approach with GroupMark is to partially automate the filtering process by the specification of recommender group profiles to match user interests to provide an effective and efficient collaborative filter. A key issue for the design of any social recommendation tool is how users interact with it to control or tailor its behaviour to meet their particular needs. Specifically, the question is: what control might a user wish to be able to exercise over:

1. who provides the recommendations, and

2. which recommendations they receive.

The first is an issue to do with being able to control the CF component, i.e., the composition and size of the recommender group(s), or recommender neighbourhood [Herlocker 99b]. The second is an issue to do with being able to control the IF component, i.e., isolating some relevant sub-set of all the recommendations identified by the recommender group.

Recommender systems like GroupMark exemplify a variety of approaches to the issue of filter control. In Pharos, users have total control over the filter, but this is at the cost of users themselves having to do the work of identifying their peer recommender group membership, i.e., those users who are similar to themselves [Bouthors 99]. In contrast, users of the Knowledge Pump [Glamce 98] complete personal profiles that are then matched computationally to define recommender group membership.[1] GroupMark also uses a computational approach to defining recommender group membership, but it gives the user a finer degree of control over how this is defined. GroupMark resembles systems such as Siteseer [Rucker 97] in that it uses bookmarks not only as the shared resource, but also as the means for determining recommender group membership.

The fundamental question here is how to operationalise the representation of similarity through the filter user interface so that a balance between computationally and user-driven adaptation can be achieved, and users can control GroupMark's behaviour easily and effectively. In particular, we argue that similarity is a contingent, situated and emergent relationship between users and that users may be better served by recommender systems that afford a graceful transition between computational and people-based filtering mechanisms.

Both the CF and IF control issues are fundamentally a matter of filter thresholding. In the general case, filter thresholds may be about trust, i.e., as in: "Show me only documents recommended by people whose opinions I value". They may be about quality, i.e., as in: "Show me only the highest quality documents". They may also be about time, i.e., as in: "Show me only documents recommended in the last week." Filter thresholds might be about content, i.e., as in: "Show me only those documents that contain the following keywords." GroupMark supports setting filtering thresholds in each of the above forms. For simplicity, we will focus on recommender group neighbourhood thresholds for CF and relevance thresholds for IF.

Such filter thresholding specifications might be used singly or in combination. The quality threshold is already implicit in the use of shared bookmarks as the recommender resource. By default, it is assumed that bookmarking is a high quality indicator, although there are other possible interpretations or nuances. For example, documents may be bookmarked because of an expectation of repeated use.

3.1. The Similarity Threshold

The aim of collaborative filtering is to recruit others to act as our filtering agents on the assumption that they are our peers, i.e., like us in tastes and quality judgements. This we endeavour to ensure by comparing opinions over a set of known documents. In this case, the approach to defining peer re-

---

[1] In fact, Knowledge Pump users must bootstrap the system by providing an initial list of "advisors", people whose opinions users particularly trust [2].

commender group membership is by comparison of individual users' bookmarking behaviour. So, membership of the peer recommender group may be defined as: "A set of users who are similar to me because their bookmarks (partially) intersect with mine."

The similarity threshold is a way of expressing the degree of similarity that is the condition for peer recommender group membership, and hence the neighbourhood size. If user A, having a bookmark set denoted by A_b, creates a recommender group G_A, then $G_A = \{X: |A_b \cap X_b| \geq t\}$, where t is the similarity threshold, the number of bookmarks that users X_1, X_2, .. etc., must have in common with A for them to be classified as *similar* to A. The set of recommendations for A, $R_A = \{X_b: X \in G_A\}$.

Clearly, the higher the similarity threshold, the smaller the neighbourhood and the 'tighter' the recommender group will be. By providing the similarity threshold as a user-definable attribute, GroupMark enables the user to exercise control over the neighbourhood size and "spread of opinion" within her recommender group. This will, in turn, have an impact on the number and quality of recommendations in the recommendations pool, and so permits the user to experiment with trade-offs between the two.

3.2. Defining Similarity by Example

A single, globally computed measure of similarity is unlikely to be very useful as a filtering device. GroupMark provides the means to contextualise similarity by allowing users to specify by example which bookmarks are relevant to its computation. This enables a user to express the following: "Here is a set of documents that I like; create a recommender group whose members also like (some or all of) these same documents."

Returning to the illustration above, user A may identify a subset of A_b, the *similarity set* A_b', that is to be used in computing the similarity threshold: $G_A = \{X: |A_b' \cap X_b| \geq t\}$. Further, user A may also distinguish between *essential* members of A_b', A_b'' and *optional* members: $G_A = \{X: |A_b' \cap X_b| \geq t$ and $A_b'' \subseteq X_b\}$.

3.3. Relevance and Resemblance Thresholds

Having defined the membership criteria for the recommender group that user A wishes to create, i.e., the CF component of GroupMark, the question is how can user A identify those members of the recommended bookmarks set R_A that are relevant to her information needs. In other words, user A needs some means to specify, perhaps in advance, a relevant subset of R_A. This is where the IF component of GroupMark now comes into play. The GroupMark system provides users with two ways to filter out unwanted members of recommended bookmark sets. First, users can use keywords to define relevance thresholds. These take the form of a set of example keywords. As in the similarity threshold, the relevance threshold can be specified as a combination of *essential* matches and *optional* matches.

Second, adapting a concept used in cluster analysis, GroupMark allows the user to define a resemblance threshold. For example, if document d (where $d \in R_A$) has been bookmarked by every member of G_A then the resemblance of d to other recommendations in R_A is maximal. Conversely, if d has been bookmarked by only one member of G_A, its resemblance to other recommendations is

minimal. Of course, it is possible to imagine cases where high or low resemblance may be the attribute that the user is looking for.

Users of GroupMark may specify keywords manually, or use system-generated keywords. The latter are generated using the well-known term frequency times inverse document frequency (TFxIDF) algorithm [Salton 89]. The principle of TFxIDF is that words that are common in an individual document, but rare in within the whole corpus are good indicators of content.

4. THE GROUP PROFILE

The group profile is a user-defined template that allows a user to express the salient properties of a recommender group by example, and is the mechanism through which users instantiate the similarity, relevance and resemblance attributes defined above.

By associating profiles with groups rather than individual users, GroupMark provides a natural way for users to define multiple recommender groups to serve multiple information needs. The practical effect of this distinction is that people can now be members of different groups simultaneously, perhaps switching between groups according to the kinds of documents they are looking for. This seems more natural (why should someone who shares your tastes in sport also share your tastes in music?) and provides a way of introducing some contextual information. For example, users can form different recommender groups to serve different information needs. The downside is that these needs may have to be articulated before recommendations can be provided (though pre-defined public groups may be an answer here) and that users will have to explicitly switch groups when the subject matter of their information target changes.

GroupMark gives the individual user the option of making her group profiles public. A user can simply re-use another's profile as found, or can re-use it as a template for defining her own recommender group(s). In this way, by providing a way of moving from sharing information resources to sharing information finding strategies, GroupMark gives its users extra collaborative leverage for the creation of recommender groups. This is an important resource for aiding the transitioning between computationally and people-maintained filtering mechanisms.

4.1. The Group Profile Editor

The Group Profile editor is the mechanism through which the user defines and controls the behaviour of the CF and IF components of GroupMark. Figure 2 shows the user interface for creating a new recommender group profile, or editing an existing one.

The editor allows the owner of the group to change the attributes of the group such as the group name, description, if the group is an "open" public group that anybody may join if they match the group criterion such as the specified keywords and URLs. The group may also be closed where only the group owner can perform the join operation whereby other users register an interest in the group so that they can view and recommend sites within that particular group.

To make recommender group profile definitions more flexible, they can be specified in terms of optional and essential keywords, and URLs, that the user must possess in order to join the group. This means that the user must have bookmarked sites that contain all the essential keywords and URLs

specified in the profile definition in order to be recommended to join this group. We can control this matching behaviour by the match threshold selector, to ensure that users who match exactly the group profile will be recommended this group. If we set a lower threshold then more relaxed group membership rules are applied. A closer match to the optional URLs and keywords indicates a better match to the recommender group and hence the user should more closely consider group membership.

5. GROUPMARK USAGE SCENARIOS

GroupMark can be used to provide a service in a number of areas.

- At a basic level GroupMark may be used to provide bookmark portability, making webbrowser bookmarks available wherever the user chooses to browse the web. This functionality is independent of any recommendation functionality.
- To share selected bookmarks with other users.
- To locate sites related to a specific area through the GroupMark topic search facility.
- To display the top most popular sites recommended by the users' of GroupMark.
- To locate other users with specific personal interests and find sites that they personally recommend.
- To be given a list of top recommended sites of other users according to your GroupMark derived interest profile.
- To be recommended other groups based on your current bookmark set to help create your personal interest profile.

Before using GroupMark the user must execute a small Java program to register with the GroupMark service. This program searches the user's machine for bookmarks, submitting them to the GroupMark database along with the extracted keywords that define the content of the page. After registration, the user can re-execute the bookmark submission program to update their GroupMark bookmarks. The system will also remove old bookmarks that are no longer in use by the user.

After logging into GroupMark the user will be presented with the main menu (see Figure 1) where all of the GroupMark functionality is available. GroupMark then displays current user information such as the current number of bookmarks you are sharing (i.e. your number of public bookmarks), the number of private bookmarks you currently hold, and information about current group membership, etc. Below this GroupMark also recommends some groups for you to join based on your bookmarks to get you started and build up your interest profile. Finally, the main menu presents top rated sites of other GroupMark users who match your current interest profile. From the main menu after submission the user may "fine-tune" their bookmark submission by editing their captured bookmarks, changing the captured keywords for each site and modifying the availability of the bookmarks to other users.

If the user does not find any groups that particularly match their personal interests then they can choose to create a new recommender group. They will then become the group's owner and as an effect of this they will be able subscribe and unsubscribe other users to this group. The user will then go on to specify a group profile whereby they can control who is recommended to join this group (see Figure 2). If you are a group owner you cannot unsubscribe from any groups that you own.
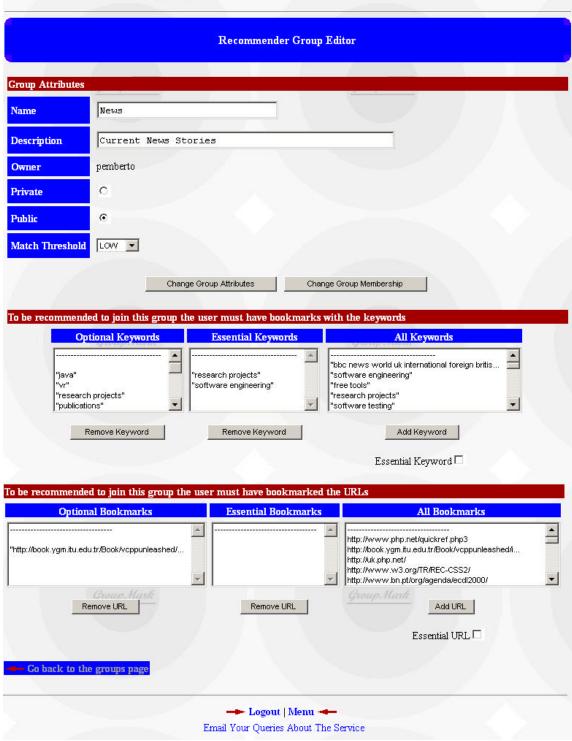
Figure 2: **The GroupMark Group Profile Editor.**

Once the user has joined and created any new groups they can then return to the main GroupMark page where new recommendations will be displayed. This provides a convenient jumping off point, or web portal to research your user interests based on the recommendations of other matched users.

6. SUMMARY, CURRENT AND FURTHER WORK

GroupMark is an example of a class of recommender system that combines CF and IF techniques together to produce higher quality recommendations. In addition, we believe that GroupMark's use of a 'recommendation by example' approach provides a simple, easy to understand, concrete but effective mechanism for users to interact, control and experiment with its behaviour. These two elements of GroupMark, we argue, represent a significant advance over other recommendation systems.

A prototype of GroupMark has been implemented and one part of our current work is focused on evaluating it with groups of users. In this, we will attempt to extend the evaluation of GroupMark beyond the more common objective performance metrics to incorporate an assessment of the value of the system as perceived by its users. As Belkin et al. [Belkin 94] observed, it is no longer sufficient to use standard precision-recall measures to evaluate information retrieval systems. It is equally important that such technologies are delivered to users in usable forms. For this evaluation to yield useful data, it is important that this evaluation be performed over an extended period of time and be able to measure subjective satisfaction as well as objective performance. As essentially social systems, collaborative filtering tools can be expected to exhibit 'network externality' properties, i.e., properties that are dynamic in time and non-linear, and for users' evaluations to reflect these properties. For example, users' perceptions may be quite different in the early stages when the amount of shared rating 'capital' is small. As this accumulates, the behaviour of both collaborative filtering systems and their users may change significantly.

Part of our future work will focus on devising improved visualization mechanisms for groups, their attributes and recommendations. For this we envisage providing some form of 3D representation of the bookmark information space, possibly VR based, or using a representation such as hyperbolic trees [Pirolli 2000].

ACKNOWLEDGEMENTS

REFERENCES

[Balabanovic 97] M. and Y. Shoham, Fab: Content-Based, Collaborative recommendation. *Communications of the ACM*, 40(3), 1997, pp. 66-72.

[Belkin 94] N. Belkin, C. Borgman, S. Dumais and M. Hancock-Beaulieu, Evaluating Interactive Retrieval Systems, in Proceedings of the *17th annual ACM-SIGIR Conference*, Dublin, ACM Press, 1994.

[Bouthors 99] V. Bouthors and O. Dedie, Pharos, a Collaborative Infrastructure for Web Knowledge Sharing, in S. Abiteboul and A-M. Vercoustre (eds), Proceedings of the *Third European Conference on Research and Advanced Technology for Digital Libraries*, Springer-Verlag, Berlin, 1999, pp. 215-233.

[Dix 94] A. Dix and A. Patrick, Query by Browsing, in P. Sawyer (ed), *Proceedings of IDS'94*: The 2nd International Workshop on user Interfaces for Databases. Lancaster, UK, Springer-Verlag, 1994, pp. 236-248

[Glance 98] N. Glance, D. Arregui and M. Dardenne, Knowledge Pump: Supporting the flow and use of knowledge, in U. Borghoff and R Pareschi (eds), Information Technology for Knowledge Management, Springer-Verlag, Berlin, 1998.

[Gross 98] T. Gross, Supporting Collaboration and Cooperation in Digital Information Environments, in workshop on Collaborative and Cooperative Information Seeking in Digital Environments, ACM Conference on Computer Supported Cooperative Work, 1998.

[Herlocker 99a] J. Herlocker, Explanations in Recommender Systems, CHI'99 Workshop on Interacting with Recommender Systems, May, Pittsburgh, 1999.

[Herlocker 99b] J. Herlocker, J. Konstan, A. Borchers and J. Riedl, An Algorithmic Framework for Performing Collaborative Filtering, in Proceedings of the *22nd annual ACM-SIGIR Conference*, Berkeley, CA. ACM Press, 1999.

[Hill 95] W. Hill, L. Stead, M. Rosensteian and G. Furnas, Recommending and Evaluating Choices in a Virtual Community of Use, in Proceedings of *CHI'95*, Denver. ACM Press, 1995.

[Kushmerick 2000] N. Kushmerick, J. Mckee and F. Toolan, Toward zero-input personalisation: Refereer-based page recommendation, poster presented at the Nineth International World Wide Web Conference, May, 2000, Amsterdam.

[Pirolli 2000] P. Pirolli, S. K. Card, and M. M. Van Der Wege, The Effect of Information Scent on Searching - Information Visualization of Large Tree Structures, in *Advanced Visual Interfaces 2000*, pp. 161-172.

[Procter 97] R. Procter, A. Goldenberg, A. McKinlay and E. Davenport, Enhancing Community and Collaboration in the Virtual Library, in C. Peters and C. Thanos (eds) Proceedings of the *First European Conference on Research and Advanced Technology for Digital Libraries*, Pisa, September 2nd-4th, 1997, Lecture Notes in Computer Science v. 1324, Springer, pp. 25-40.

[Procter 99] R. Procter, et al., SELECT: Social and Collaborative Filtering of Web Documents and News, in: A. Kobsa and C. Stephanidis (eds), Proceedings of the *Fifth ERCIM Workshop on User Interfaces for All*, Dagstuhl, Nov 28th-Dec 1st, 1999, pp. 23-37.

[Resnick 97] P. Resnick and H. Varian, Recommender Systems, *Communications of the ACM*, vol. 40(3). ACM Press, New York, 1997, pp. 56-58.

[Rucker 97] J. Rucker and M. Polanco, Siteseer: Personalised Navigation for the Web, *Communications of the ACM*, vol. 40(3). ACM Press, 1997, pp. 73-75.

[Salton 89] G. Salton, Automatic Text Processing, *The Transformation, Analysis and Retrieval of Information by Computer,* Addison-Wesley, Reading MA, 1989.

[Schafer 99] J.B. Schafer and N. Good, Talk to my Agents: Research Issues in Combining Information Filtering Agents with Collaborative Filtering, CHI'99 Workshop on Interacting with Recommender Systems, May, Pittsburgh, 1999.

[Tuzhilin 99] Tuzhilin and G. Adomavicius, Integrating User Behaviour and Collaboratiive Methods in Recommender Systems, CHI'99 Workshop on Interacting with Recommender Systems, May, Pittsburgh, 1999.

[Twidale 96] M.B. Twidale and D.M. Nichols, Collaborative browsing and visualisation of the search process, in *Aslib Proceedings*, vol. 48(7-8), 1996, pp. 177-82.

[Zloof 75] M. M. Zloof, Query by Example, in Proceedings of *AFIPS National Computer Conference,* AFIPS Press, 1975, New Jersey, pp. 431-438.