Visualizing Recommendations to Support Exploration, Transparency and Controllability

Katrien Verbert
Affiliation
Address
e-mail address
Optional phone number

Denis Parra
Affiliation
Address
e-mail address
Optional phone number

Peter Brusilovsky
Affiliation
Address
e-mail address
Optional phone number

Erik Duval
Affiliation
Address
e-mail address
Optional phone number

ABSTRACT
Research on recommender systems has traditionally focused on the development of algorithms to improve accuracy of recommendations. So far, little research has been done to enable user interaction with such systems as a basis to better incorporate current needs and to support exploration and control by end users. In this paper, we present our research on the use of information visualization techniques to interact with recommender systems. More specifically, we investigated how information visualization can improve user understanding of the typically black-box rationale behind recommendations in order to increase their perceived relevancy and meaning and to support exploration and user involvement in the recommendation process. We evaluated our approach in the context of academic conferences. Users can explore articles, users, tags, and recommendations and the relationship between these entities. Based on evaluation results of user studies at two conferences, we obtained interesting insights to enhance user interfaces that integrate recommendation technology. More specifically, effectiveness of recommendations and probability of item selection both increase when users are able to explore and interrelate multiple entities – i.e. items bookmarked by users, recommendations and tags.

Author Keywords
User interfaces for recommender systems, information visualization, user studies.

ACM Classification Keywords
H.5.2. Information interfaces and presentation (e.g., HCI): User interfaces. H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms
Human Factors; Design; Experimentation.

INTRODUCTION
Interactive information visualization and recommendation techniques have both been researched as ways to help people deal with abundance of information. The main advantage of interactive visualization is that 2 or 3 dimensional representation allows the user to more easily see multiple aspects of data while being in control when exploring information. The main advantage of the traditional recommendation approach is that offering a clear list of items ranked by perceived interest, it reduces cognitive overload associated with exploring a rich set of items.

In this paper, we present our research on the combination of both approaches. We investigated how graphical representations and the ability to combine a personalized prospect offered by a recommender engine with other valuable prospects can improve user trust in the results offered by the black-box recommender engines and increase user ability to find interesting items.

Our work has been motivated by the presence of multiple *relevance prospects* in modern personalized social tagging systems. An important feature pioneered by social tagging systems and later used in other kinds of social system is the ability to explore different *community relevance prospects* by examining items bookmarked by a specific user or items associated by various users with a specific tag. Items bookmarked by a specific users offer a social relevance prospect: if this user is known and trustable or appears to be like-minded (bookmarked a number of items known as interesting) a collection of his or her bookmarks is perceived as an interesting and relevant set that is worth to explore for more useful items. Similarly, items marked by a specific tag offer a content relevance prospect. Items related to a tag of interest or a tag that was used to mark many known interesting items are also perceived as potentially relevant and worth to explore. In this context, a ranked list of recommended items offered by a specific recommender engine can be considered as yet another *personalized* relevance prospect.

The problem that we address is that existing personalized social systems do not allow their users to explore and combine multiple relevance prospects. Only one prospect can be explored at any given time – a list of recommended items, a list of items bookmarked by a specific user or a list of items marked with a specific tag. We believe exploring a single prospect is not sufficient for since none of the prospect could be fully reliable and trustable by the users (that includes recommendations generated by black-box engines). In this context, the ability to combine prospects might offer a more reliable and attractive way to explore. For example, knowing that a specific item has been not only recommended by an engine, but also bookmarked by two trustable users can remarkably increase user trust in the quality of this item.

To solve the aforementioned problem and to allow users an approach to explore and combine multiple relevance prospects, we suggest a specific interactive visualization approach. This visualization embodies suggestions offered by various recommender systems *as recommender agents* that can be perceived as being analogous to human users and as a result, exportable in parallel with the relevance prospects offered by users and tags in social systems. We believe that using interactive visualization can increase the transparency of the recommendation process and allow the users to be in control of their exploration.

A special issue on interfaces for recommender systems [25] illustrates the interest and importance of intelligent interfaces for recommender systems. Such interfaces are researched to provide new capabilities to the users of the recommender system to search, browse, and understand the results of the recommendations. In our work, we focus on the use of information visualization techniques to support such new capabilities. The research contribution of this work is threefold:

1) First, we present a novel and synergetic approach to combine multiple relevance prospects that includes personalized relevance as offered by different recommenders and social relevance as offered in social bookmarking systems as a basis to provide transparency and to support data exploration search. In this approach, recommender systems are presented as agents and their interrelationship can be explored (i.e. a user can explore which items are suggested by multiple recommender agents). In parallel, real users and their bookmarks are shown and users can explore both interrelationships between users as well as interrelationships between agents and users. To our knowledge, this combination of agents and real users has not yet been explored as a means to support exploratory search and controllability, and to increase trust and acceptance of recommendations.
2) Second, we present a user interface which serves to both explain the provenance of recommendations in a transparent way and to support exploratory search. Users can browse bookmarks of other users, tags and suggestions of recommender agents as a basis to find relevant items.
3) Third, we have evaluated the usefulness of this interactive interface with 20 users to gain insights in the influence on user satisfaction when users can control and combine entities involved in the recommendation process.

This paper is organized as follows: first we present related work in the area of user interfaces for recommender systems. Then, we introduce TalkExplorer, an interactive visualization of users, tags, talks at a conference and recommendations for conference attendees. The evaluation of this visualization at conferences is presented. Finally, we discuss the results of this case study, lessons learnt and future research opportunities.

BACKGROUND AND RELATED WORK

Recommender systems

Recommender algorithms can be broadly categorized in three areas:

1. *Collaborative filtering* recognizes commonalities between users or between items on the basis of explicit or implicit relevance indications [11] such as ratings [2] and tags [23]. Implicit data used by recommender systems include actions like reading [19] or watching TV series [12]. A standard user-based collaborative filtering algorithm first identifies similar users based on their overlapping interactions or similar ratings of common items. It then makes recommendations based on preferences of these similar users. A standard item-based recommendation algorithm analyzes similarities between items and then uses these similar items to identify the set of items to be recommended. Collaborative filtering is the most widely implemented and most mature technology [5].
2. *Content-based filtering* matches descriptions of items to descriptions of users [24]. This approach bases predictions on information about individual users and items, and ignores contributions from other users.
3. *Hybrid recommender systems* combine recommendation techniques, to gain better performance with fewer drawbacks [5].

Although these algorithms have been implemented and validated on a large scale in several application areas [20], there are important challenges that need to be addressed before recommendation can realize its full potential.

1. Collaborative recommendation techniques often suffer from *cold start* issues, i.e. they cannot make effective recommendations for new users or for new items that have no explicit or implicit relevance indicators yet [6].
2. It is *difficult to explain* the rationale behind recommendations to end users [10]: the complexity of recommendation algorithms often prevents users from comprehending recommended results and can lead to *trust* issues when recommendations fail.
3. Allowing users to *control* the way they can sort lists of recommendations [15] or the neighbors' contribution in

a social recommender [16] has shown a positive effect in user satisfaction. However, there are several ways that users can control elements in the interface: which are the most effective for the user experience?

The design and development of user interfaces for recommender systems has gained increased interest. Such interfaces are researched to provide new capabilities to the users of the recommender system to search, browse, and understand the results of the recommendations [25]. Among others, explaining recommendations to provide transparency and increase trust has been researched extensively over the last decade [26]. In most cases, such explanations are presented in plain text and indicate why a specific item is suggested to a user – such as *"Because you have selected or highly rated: Movie A"*.

In addition to supporting transparency, we are particularly interested to enable interaction with recommender systems as a basis to support exploration and controllability. In recent years, some research has been done to visualize recommendations to enable such new capabilities. We elaborate on existing work in this area in the next section.

Visualizing recommendations

Some research has been done to visualize recommendations to enable user interaction with recommender systems. Most existing work in this area focuses on visualization of collaborative filtering recommender systems. PeerChooser [21] is a visual interactive recommender that uses a graph-based representation to show relations between users and recommended items of a collaborative filtering recommender system. Similarly, SmallWorlds [8] allows exploration of relationships between recommended items and similar friends, and recommended items and more distant friends in multiple layers of similarity. These systems enable users to explore such relationships as a basis to provide transparency and to support the user to find new items that may be relevant. Pharos [28] is a social map-based recommender system that visualizes a summary of social network activity of different communities. The system uses topic modeling [3] to provide new users with an overview of the site in order to alleviate the cold start problem of collaborative filtering recommenders.

Some systems focus specifically on tags that are used by social recommenders. SFViz (Social Friends Visualization) [9] visualizes social connections among users and user interests as a basis to increase awareness in a social network and to help people find potential friends with similar interests. This system uses a Radial Space-Filling (RSF) technique [7] to visualize a tag tree and a circle layout with edge bundling to show a social network.

FaceTag [27] is a tool that helps users see the relationship between user-generated tags and recommended facets in a classification scheme. Kammerer et al. [13] designed a tag-based search browser to recommend relevant tags for further search. Research on this stream only focuses on information and meta-information concerning items, and ignores the users who contributed such information and relationships among those users [9].

More recently, TasteWeights [4] has been introduced as a system that allows users to control the importance of friends and peers in social systems to obtain recommendations. Similar to our work, TasteWeights introduces the concept of an interface for hybrid recommender systems. The system elicits preference data and relevance feedback from users at run-time and uses these data to adapt recommendations to the current needs of the user. To our knowledge, this is one of the first systems that enables interaction with a hybrid recommender system and that can be adjusted by end users to control the output of the systems. In our work, we extend this concept of visualizing and combining the output of multiple recommenders as a basis to support exploration and controllability. More specifically, users can explore suggestions of different recommender agents (i.e. a content-based and a tag-based social recommender), bookmarks of other users, tags, and interrelationships between these entities. Thus, we do not specifically focus on explaining collaborative filtering results – but rather on combining and interrelating suggestions of recommenders and real users to increase the relevancy and meaning of recommended items. We present our approach in the next section.

TALKEXPLORER

TalkExplorer is an interactive visualization tool that enables users to explore and bookmark most potentially interesting research papers and talks at a conference using recommender agents and social data (tags, bookmarks, and connections to other users). The visualization was built as a component of a conference support system *Conference Navigator 3*. We first present Conference Navigator 3 and the recommendation functionalities that it provides. Then, we present the objectives of visualizing recommendations and details on the design and development of TalkExplorer. Evaluation results are presented in the next section.

Conference Navigator

Conference Navigator 3 (CN3) is a social personalized system that aims to support users in several activities related to academic conferences [22]. At the time of writing, 18 conferences have been supported by CN3. Among different features, Conference Navigator provides a conference schedule, a list of the conference talks and details of each talk. (illustrated in Figure 1). It also provides information about people related to the conference such as the list of attendees, and the list of authors. Users can add papers to create a personal schedule, they can add tags (free user keywords) to each talk, and they can also connect with other CN3 users by following them (unidirectional relationship) or connecting with them (bidirectional relationship). Social information collected by CN3 is extensively used to help the users in finding most interesting papers. For example, in the page called "Top

Items", CN3 summarizes the most popular articles, the most active people in each conference, their institutions, and also the most popular tags associated to the papers.

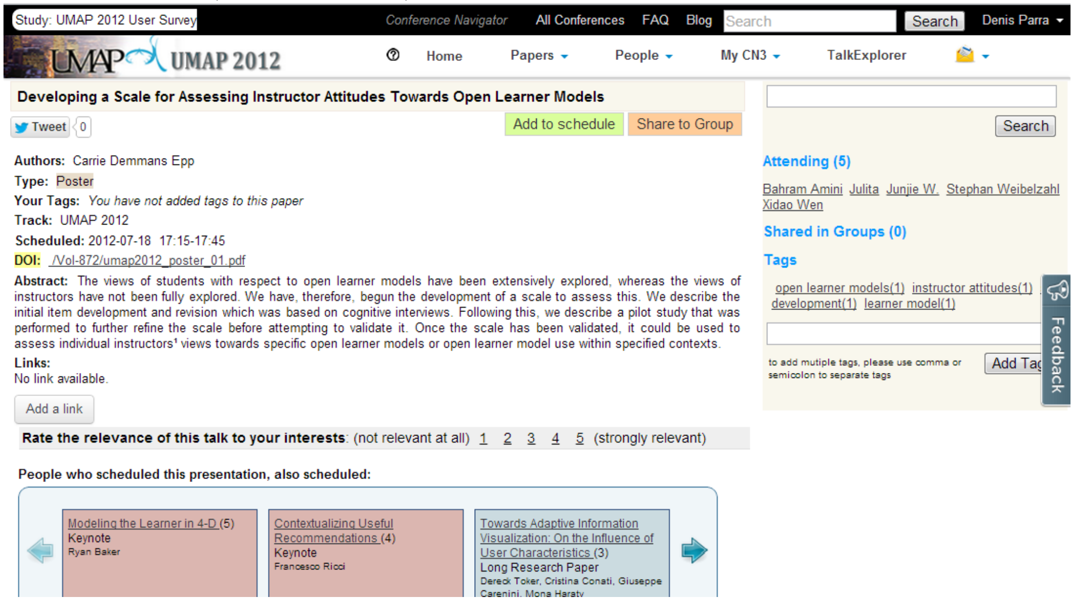

Figure 1: screenshot of page details in CN3

When visiting talk information, as shown on Figure 1, users can also see who *scheduled* each talk during the conference and which tags were assigned to this talk. This social information is also used to provide links to similar papers ("People who scheduled this presentation, also scheduled:") mimicking the well-known Amazon.com's suggestions [17]. Similarly, when visiting a user page, other users can see which talks she or he is planning to attend (given that personal privacy settings provide access to this information). Finally, a click on a tag will allow a user to see all talks marked with this tag by the community of users.

A social information access supported by CN3 interface is complemented with the system ability to offer personal recommendation of talks and papers. CN3 supports two kinds of recommendations that are offered through two separate ranked lists: content-based and tag-based. The content-based engine constructs the user interest profile as a vectors of terms with weights based on TF-IDF [1] using the content data of the papers that the user has scheduled. Then, it recommends papers that match the profile of interests. The tag-based recommender engine makes use of the tags (user-defined keywords) that users associate to conference talks. By matching user tags (tags applied by a target user) with item tags (tags assigned to different talks by the community of users) using the Okapi BM25 algorithm [18], the engine identifies relevant talks and suggests them to the active user.

Visualizing recommendations, tags and users

In the original CN3, ranked links produced by the content-based and tag-based recommenders are presented in separate pages and can be used by users to find new talks to attend at a conference. In addition, users can explore bookmarks and tags of other users as a basis to find new items. In this paper, we are particularly interested in assessing the potential influence of perceived relevancy and meaning of recommended items when we enable end users to explore and combine multiple relevance prospects, most importantly personalized and social prospects.

TalkExplorer is an interactive visualization developed on top of data collected by CN3. The interface serves to both explain the provenance of recommendations in a transparent way and to support exploration and control by end users. More specifically, users can browse and interrelate bookmarks of other users, tags and suggestions of recommender agents as a basis to find relevant items.

The visualization is implemented as a Java applet and uses the Aduna clustermap visualization library[1]. This software library enables to visualize sets of categorized objects and their interrelationships. The library has been used in related research to explore the interaction of users, resources and tags in social tagging systems [14]. In this research, the library was deployed on top of delicious.com data to explore bookmarks of users and to support exploratory search. We adapted this initial version to visualize the interactions of users, tags, and *recommender agents* in terms of conference talks that they have in common – as illustrated in Figure 2. The objective was to explore new ways to enable end users to interact with recommender systems and items that are suggested to them. More specifically, we wanted to explore a novel and synergetic approach to combine outputs of multiple recommenders and

---

[1] http://www.aduna-software.com/technology/clustermap

bookmarks of users they know or with whom they share common interests as a basis to enable controllability by end users and to increase perceived relevancy and meaning of recommendations.

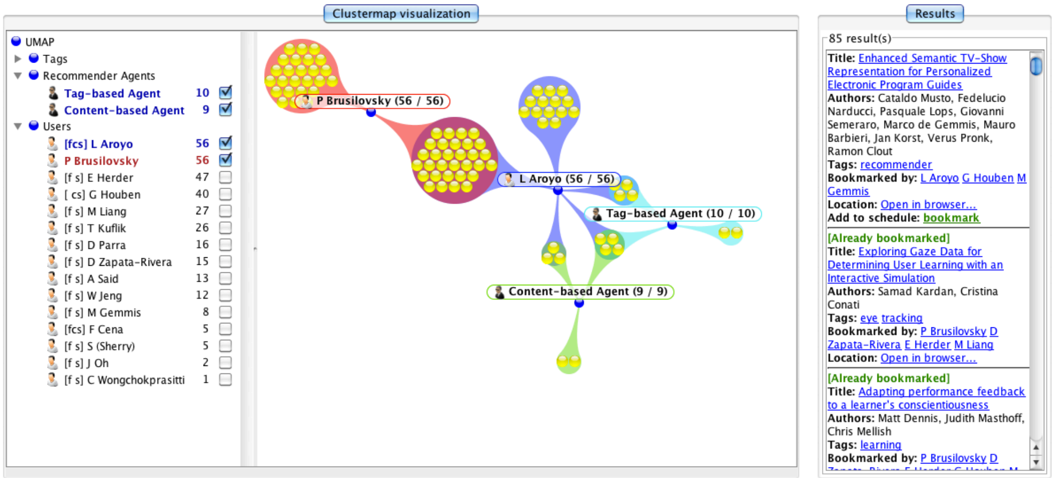

Figure 2: TalkExplorer

Recommender systems are presented as agents and their interrelationship can be explored. In parallel, real users and their bookmarks are shown and users can explore both interrelationships between users as well as interrelationships between agents and users (i.e. which other users have bookmarked talks that are recommended to them). In addition, relationships with tags can be explored to identify relevant items. We are researching whether visualizing these relationships can help users to find relevant talks to attend at a conference, and whether this visualization can provide transparency and increase trust.

As shown in Figure 2, fifteen users from the neighborhood of the active user are shown and users can explore items that these users have bookmarked as a basis to find new items. The selection of these users is based on uni- and bidirectional connections between users - i.e. users that the active user follows (user name preceded by [f] in Figure 2) and connections (user name preceded by [c]), respectively. In addition, users that have similar interests based on common bookmarks are shown in this neighborhood (preceded by [s]). In addition, we deployed a content-based and a tag-based recommender agent that present the output of two CN3 recommendation algorithms. Suggestions of these recommender agents are shown in parallel to bookmarks of users.

TalkExplorer allows users to explore the different entities of the conference by means of three principal components, as shown in Figure 2. On the left side, the entity selection panel allows users to select tags, users and recommender agents that are added and displayed in the canvas area. This canvas area, at the center of the applet, shows a clustermap visualization - i.e., different clusters of talks linked by connected components. The labeled circles in this canvas area represent either real users, recommender agents or tags. Yellow circles represent individual talks, and the bubbles that involve them represent clusters of talks.

In Figure 2, two users are shown (P Brusilovsky and L Aroyo), as well as suggestions of the tag-based and content-based recommender agent. The cluster map visualization enables users to explore relationships between items that were suggested to them by these recommender agents and bookmarks of users on the screen. For instance, a user can see which other users have bookmarked a talk that is suggested to them by a recommender agent by exploring the intersection of the agent and a specific user on the screen. Users can arrange the different entities displayed in the canvas by dragging them with the mouse.

Finally, the rightmost panel shows the detailed list of talks. This can be a list of all the talks presented in the canvas area, or a subset of them related to the selected entity. For example, if a user clicks on a specific CN3 user in the canvas area, the papers that the selected user has bookmarked are presented in the list. If a user clicks on a cluster (for example, the cluster showing talks that were bookmarked by a user and a specific agent) the list of these talks is presented.

EVALUATION

We evaluated TalkExplorer at two conferences where CN3 was used as the main conference assistance system. . One evaluation was conducted at the ACM Hypertext 2012 conference in June 2012 (HT 2012). For this evaluation, the visualization was deployed on top of data that was collected from HT 2012 attendees. Another evaluation was conducted at the User Modeling, Adaptation, and Personalization conference in July 2012 (UMAP 2012). Both evaluations were performed with attendees of respective conferences using real conference data (i.e., using actual talks and schedules and collected data on talks of the conference participants have bookmarked, tagged or rated). As explained in the previous section, in the process of evaluation the users were asked to explore conference talks using full-featured Aduna Clustermap visualization provided by TalkExplorer. The visualization provided access to the content-based and tag-based recommender agents and allowed to explore talks bookmarked by related users or tagged with user-employed tags.

Participants

In the ACM Hypertext evaluation, fourteen users participated in a controlled experiment at the conference. We inquired about the number of Hypertext conferences participants have attended, as well as their knowledge and expertise in recommendation and visualization techniques, respectively. The latter were rated on a five point Likert scale. On average, participants attended 1.5 conferences in the past (std. deviation 0.732). Most of the participants have knowledge about or expertise with visualization techniques (average 4.285, std. deviation 0.7). In addition, familiarity with recommendation techniques is high – although less extensive than with visualization techniques (average 3.7, std. dev. 0.8).

Seven participants of the UMAP conference participated in the second study of our visualization. They had a high familiarity with visualization techniques (mean 4.2, std. deviation 0.76) and a relatively high familiarity with recommendation techniques (mean 3.7, std. deviation 0.95). On average, participants attended 2 UMAP conferences (std. deviation 1.5).

Tasks

We asked users to complete three tasks:

1. In the first task, they were asked to find a new relevant talk to attend by exploring talks that users in their neighborhood bookmarked (Task 1 – T1)
2. In the second task, subjects had to find a new relevant talk by exploring the content-based and tag-based recommender agents (Task 2 – T2).
3. In the third task, they were asked to find a new relevant talk by exploring the tags (Task 3 – T3).

Data collection

Data was collected in two ways. The think aloud protocol was used during the experiment to facilitate the collection of relevant feedback from participants. We recorded the screen and voice of participants using Camtasia Studio[2] during the experiment. Afterwards, participants were asked to fill out a survey inquiring about their needs at a conference and the respective usefulness of the visualization to address these needs.

Results

To assess the value of interactive multi-prospect visualization offered by TalkExplorer, we have analyzed the way in which users *explore* and *use* the visualization. In the remainder of this section, we refer to selectable users, agents and tags as *entities* in the visualization. Papers or talks associated with these entities are referred to as *items*. We refer to *intersections* of entities when multiple entities were selected at the same time and their common items were explored. When a user found a relevant item, we said indistinctively that she *bookmarked* or *scheduled* the items.

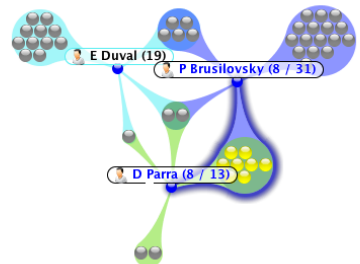

Figure 3: Exploring an intersection: items bookmarked by 2 other users but not yet bookmarked by the active user

We measured the *effectiveness* and yield (*average precision*) of different combinations of entities to gain insight in the relative success rate of different combinations of entities to find relevant items.

Effectiveness measures how frequently a specific combination type produced a display that was used to bookmark at least one interesting item. It is calculated as the number of cases where the exploration of this combination type resulted in a bookmark divided by the total number of times this combination type was explored. For instance, a combination of an active user with one specific user with whom the active user is related focused on the set of items bookmarked by this user and not yet bookmarked by the active user was explored 75 times by all participants. 23 of these sets were used to bookmark a new item. Thus, the effectiveness of exploring the set of items of a specific user is 23/75=31%. The number of item sets explored and the item sets used to bookmark a relevant talk, as well as the *effectiveness*, are presented in Figure 5.

[2] http://www.techsmith.com/camtasia.html

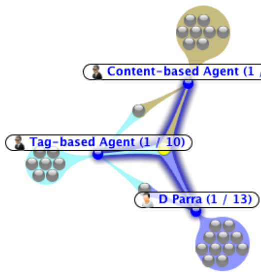

Figure 4: talk in intersection of agents and one other user

In addition, we counted the number of items in the sets where the selection was made to check *yield* of different kinds of sets. The yield of a specific combination type was measured by summing up the total number of selections made from each combination type divided by the total number of items that were shown in the combinations where the selection was made. In other words, yield measures a chance of a random item shown in a specific combination type to be useful (bookmarked). Yield results are presented in Figure 6.

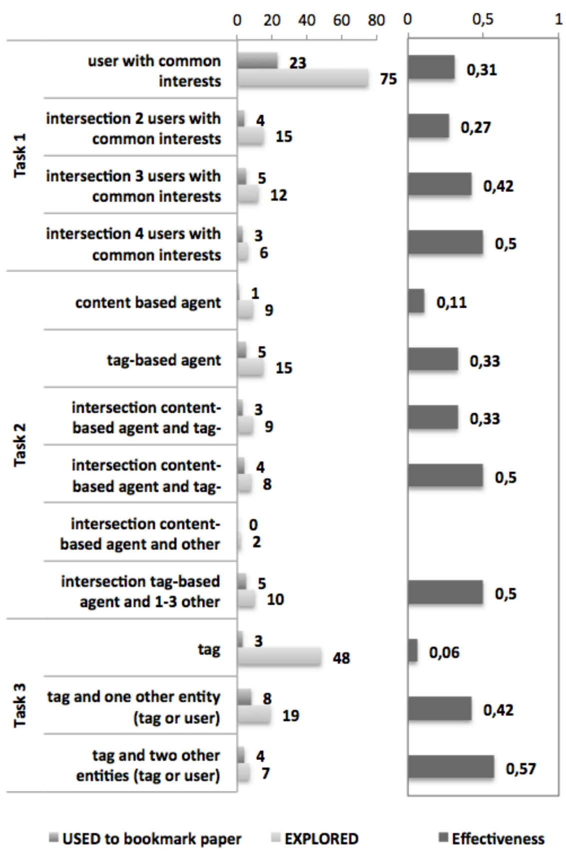

Figure 5: summary of actions explored, used to bookmark a paper, and effectiveness of those actions.

Figure 3 presents an example where the active user, E Duval, used the intersection of two other users as a basis to find a relevant item. In this example, E Duval used the set of 8 items in the intersection of two other users (P Brusilovsky and D Parra) to find an item. The yield indicates the number of selections made by users from a specific set of entities divided by the sum of the number of items in this set (8 in the example presented in Figure 3). For the intersection of 2 users, there were 2 selection out of 2 items and 2 selections out of 1 item. We have yield or probability of selection of (1+1+1+1)/(2+2+1+1)=0.66.

Task 1

In the first task (T1), users were asked to find a relevant talk by exploring bookmarks of users in their neighborhood. Results are presented in Figure 5. The set of items of one specific user with whom the active user is related was explored 75 times by all participants. 23 of these sets were used to bookmark a new item. The sum of items in these 23 sets is 276. Thus, the effectiveness of 23/75 = 31% (first top bars in Figure 5) and the yield or average precision is 23/276 = 8% (first bar in Figure 6).

Fifteen users explored intersections of two related users focusing on talks that they have not yet bookmarked (as illustrated in Figure 3). This kind of set was used to bookmark a talk 4 times (effectiveness = 27%) and the sum of items in the used sets was 6 (yield = 66%).

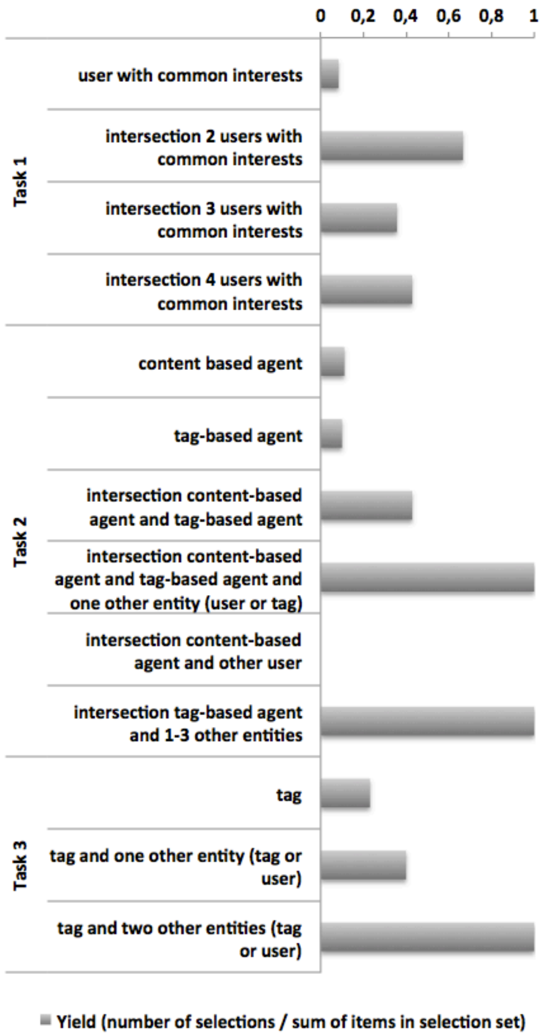

Figure 6: summary of yield results

Talks in the intersection of three or four other users were explored 12 and 6 times and used 5 and 3 times, respectively (effectiveness of 42% and 50%). The sum of items in the selection set was 14 and 7, respectively (yield of 37% and 43%). As we can see from this data, the general trend is clear: the sets that allow users to explore the overlap of several prospects are both more efficient and have higher yield (precision). Moreover, there is a general tendency for both efficiency and probability of selection to increase when more entities are used in the selection process. Small fluctuations within the general trend can be explained by the small sample.

Task 2

In the second task (T2), users were asked to find a relevant talk by exploring the output of recommender agents (a content-based and a tag-based agent). Results are presented in Figure 5 and Figure 6, the middle set with 6 possible actions.

One out of nine users found a relevant talk by exploring suggestions by the content-based agent that were not related to any other entities on the screen (effectiveness=11%, yield=11%). Five out of 15 users found a relevant talk by exploring suggestions of the tag-based agent. Three out of nine users found relevant items by exploring the intersection of agents (i.e. talks that were suggested to them by both the content-based and the tag-based agent). Four out of eight users found relevant items by exploring the intersection of the agents with another entity. Figure 4 presents such an example.

The set of items of the content-based agent in combination with another user was explored twice, but not used to find a relevant item. The tag-based agent in combination with one or more entities was successful in 50% of the cases. The results presented in Figure 5 indicate the same trend: a higher number of entities (prospects) involved in the intersection again increases the effectiveness and the yield of the resulting set.

Task 3

In the third task (T3), we asked users to find interesting talks by exploring tags that were added by users to talks. Results are presented in Figure 5 and Figure 6, bottom set of 3 actions.

As the data shows, using a single tag prospect (i.e., exploring items related to one selected tag) results in the lowest effectiveness registered in the study – as only three users were able to find a relevant item (effectiveness 6%). The sum of the number of items in the set when a selection was made was 13 (yield 3/13=23%).

Combining a tag prospect with prospect of one or more additional entities was more effective. 19 users explored the combination of a tag with one other entity and 8 users used this intersection to bookmark an item (effectiveness=42%, yield=40%). A tag in relation to two other entities was even more effective (effectiveness=57%, yield=100%). in 57% of the cases. Some users indicated that they particularly liked this functionality – as this allows them to retrieve specifically items of their topic of interest from users they know or who have a high reputation in the field.

Summary results

Effectiveness and average precision (yield) results are summarized in Figure 7. Overall, these results indicate that effectiveness of an explored set increases once more prospects are integrated, i.e., more entities are involved in the exploration process, a pattern clearly seen in Figure 7.b).

Similar trends are observed when we look at yield. These results indicate that the probability of selecting an item from a set of entities increases if more entities are overlapped.

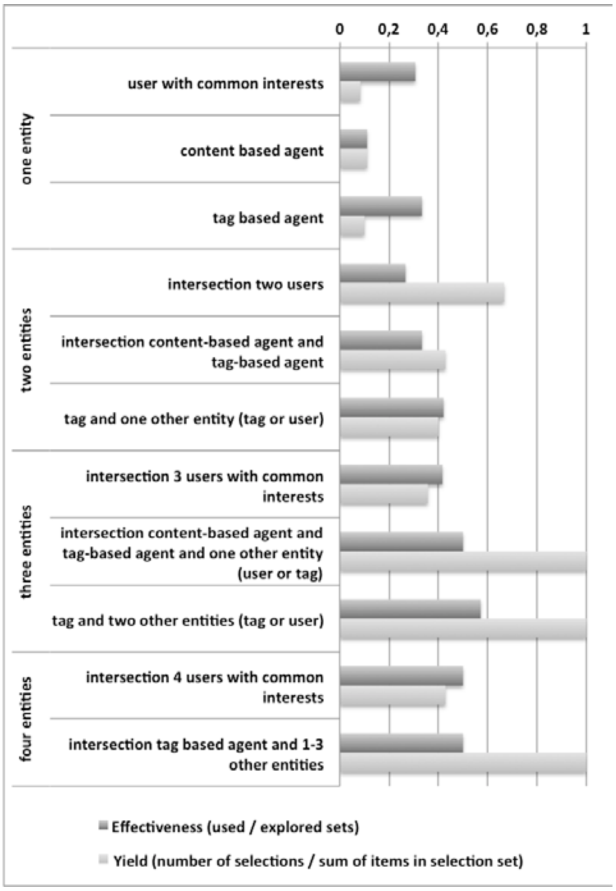

a)

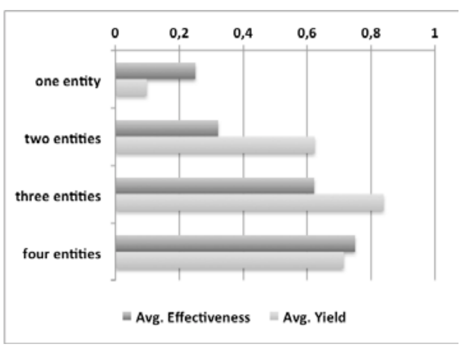

(b)

Figure 7: summary effectiveness and yield results

In Figure 7.a), we can see that yield increases from 8% when the set of items of one user is used, 11% when the set of items of one recommender agent is used and 23% when the set of items of one tag is used to 66% when the intersection of two users is used, 43% when the intersection of agents is used and 40% when a tag is used in combination with another tag or user.

These results illustrate that enabling end users to explore interrelationships between two prospects (sets of items of in the overlap of two entities) increases the probability of finding a relevant item. Except for a few cases, this probability keeps increasing if more entities are involved in the selection process (up to 100% when the intersection of agents is combined with items of a user in the neighborhood and when items in the set of a tag are combined with two other entities). The results also allow us to make several interesting observations.

First, it is interesting to note that the least effective kind of set is the set of items related to exactly one tag (6% effectiveness). Incidentally, is the only option to use tags for item exploration offered in many tag-based system. As shown by our data the systems that don't allow exploring items related to combinations of several tags are not doing good service to their users.

In contrast, exploring a prospect of a single related user in relation to the target user is a relatively effective approach – almost 1/3 of explored combinations produced bookmarks. It shows that the prospects of human users are much more valuable (and trustable) for users than prospects offered by tags.

A single recommender falls somewhat between single tags and single users by its effectiveness. It might be related to the fact that the exact mechanism of a recommender agent selection is not as clear for the user and as a result, a set of items offered by a single agent might not be as trustable as a set of items collected by a related user. Yet, when one or more additional prospects is added to the prospect offered by a recommender agent (i.e., another agent, another entity, or both) it produces best combinations in our study – those with both high effectiveness and high yield.

Overall, these results illustrate the added value of enabling users to combine social and personalization features of recommenders as a basis to increase the relevancy and meaning of suggested items.

Questionnaire results

To collect additional feedback, we used a questionnaire to inquire about needs and issues that users have when attending a conference and the extent to which our visualization addresses those needs and issues. This questionnaire was used to collect some preliminary additional feedback and assembled only a few questions we considered important at this stage. We elaborate in the next section on the use of standardized questionnaires in future studies.

Figure 8 presents the results of our questionnaire on a five point Likert scale. The first column of bar charts presents answers to questions that inquire about the importance of issues and needs at a conference. The second column presents to which extent TalkExplorer addresses these needs and issues. These results indicate that conference attendees perceive 'finding relevant talks at a conference' as an important need (median 5) and that TalkExplorer addresses this need in a good way (median 4). 'Being aware which talks my friends are attending' is also perceived as important (first column) and most participants indicate that this need is addressed by TalkExplorer (second column). 'Being aware which talks people I follow are attending' is considered less important and is also not so well addressed by our tool (median 3). 'Finding other users who have similar interests' is important for many participants (median 4) and is addressed well by TalkExplorer (median 4). 'Gaining insight into why talks are recommended' is important for many, but not all, participants. This need is addressed well by TalkExplorer according to most, but again not all, participants (median 4).

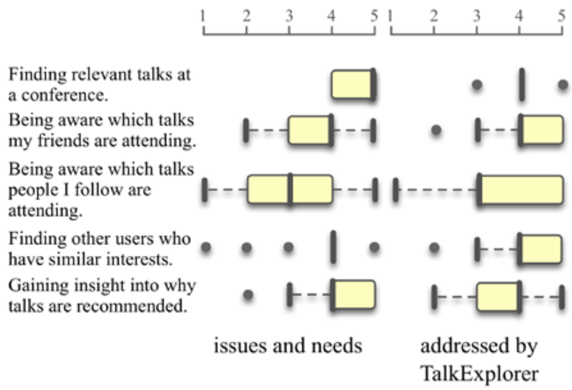

Figure 8: questionnaire issues and needs

In addition, we inquired about perceived usefulness of the visualization to support transparency and exploration. These results are presented in Figure 9 and indicate that participants perceive the visualization as useful because it gives more insight than a plain list of talks (median 4). In addition, most participants liked our idea of adding agents in parallel to real users as a means to find interesting talks (median 4) – among others to compare relevancy of recommendations (median 5).

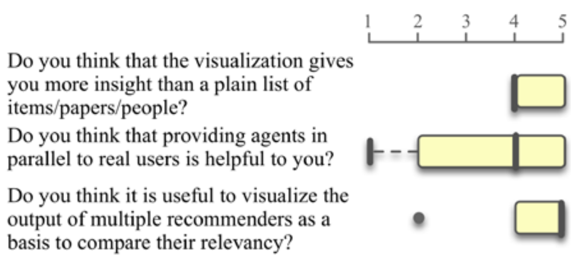

Figure 9: perceived usefulness

Additional feedback

During the study, many participants also gave suggestions to improve the usability of the visualization. Among others, users requested sorting functionality to more easily be able to identify tags. In addition, agents were not easy to locate by some participants. These agents were presented in the list of users with a different icon – and not as a separate node in the tree on the left side panel as in the current version that is shown in Figure 2. We added functionality to sort users and tags by frequency and by name in the next version that was deployed for the UMAP 2012 conference. In addition, agents are presented in a separate node as presented in Figure 2.

While users liked exploring bookmarks of users in their neighborhood, some users remarked that not all these users are known to them and they would like to see a profile page of these users. In addition, many users asked why some users are shown by default in the visualization, and how these users are selected. People recommendation as opposed to recommendation of papers and talks was another suggestion from a participant that would be useful to explore.

DISCUSSION

Evaluation results presented in the previous section indicate that the perceived utility of visualizing and combining bookmarks of users, suggestions of recommender agents and tags is perceived as useful to increase the relevancy and meaning of recommendations. Results from our questionnaire are generally positive and indicate that participants value exploring our visualization to gain insight into why talks are recommended. In addition, they indicate that such a visualization gives more insight than a typical ranked list of recommendations.

Interaction patterns indicate that users often explore relationships between entities to find relevant items. Although items of one specific user were explored most often to find a relevant item (see Figure 5, first row), the effectiveness and probability of selecting an item from the set of items is lower than with intersections of multiple entities. Also remarkable is that the effectiveness and yield of individual recommender agents and tags is low. Interrelating these entities with other entities on the screen increases their perceived relevancy and meaning in a significant way (results are summarized in Figure 7).

While these results illustrate the usefulness of visualizing and combining recommendations, tags and users, there are several limitations to this study that should be articulated and addressed in follow up studies. First, we asked users to explicitly explore users in their neighborhood, recommender agents and tags in three separate tasks. While results of these tasks give some interesting insights in the usefulness of these entities and the way users interacted with additional entities during these tasks, we cannot draw strong conclusions about the relative effectiveness of tags, users and agents in this way. First, the order of the tasks may have had an influence on the effectiveness of these entities. Second, we explicitly asked to explore these entities. In a follow up study, we are capturing interactions of users with the visualization in an open setting where users are free to explore various entities. Such a study and analysis of interaction patterns will yield more accurate data with respect to relative effectiveness of tags, users and agents. Second, the questionnaire we used was a preliminary set of questions we assembled to gather initial feedback. In a follow up study, we plan to conduct more elaborate surveys based on standardized questionnaires used to assess the accuracy, diversity and novelty of recommendations. Such studies are necessary to gain insight in other potential benefits that this interface can offer to end users.

CONCLUSION AND FUTURE WORK

In this paper we have presented and discussed the results of two studies involving conference attendees that select relevant talks by making use of the TalkExplorer visualization tool, embedded in CN3. TalkExplorer allowed users to explore items (talks) by combining different entities (users, tags, recommender agents). After analyzing the users' behavior and their answers in a survey, we highlight three results.

First, for the best of our knowledge, we haven't find in the recommender systems literature neither studies that represent recommender algorithms as agents nor systems that let users interact and combine the output of recommendation algorithms. Our results indicate that this can be a significant contribution to the area of user interaction in recommender systems, and we will expand our research in this area to other domains –beyond conferences and paper suggestions- and to investigate different ways that users can control and combine the output of several recommendation methods rather than providing a system's optimized blend of algorithms.

Second, our study confirms previous results on trust-based and social-based recommender systems regarding the positive influence on user satisfaction when they can control and inspect the recommendations [21, 15, 16]. However, it is still an open question understanding the effects of personal characteristics (such as expertise, and the users' visual processing fit), and how different ways that a user can control the interface might influence her performance and user satisfaction. We are preparing ongoing research in this direction too.

Finally, users show a better performance finding relevant items – in terms of the number of actions needed to discover relevant items - by foraging for additional evidence. In TalkExplorer, users accomplish that by intersecting the preferred items of several users, items associated with tags, and agent recommendations. This opens research opportunities in studying further the benefits on users with more interactive recommender systems interfaces in terms of user trust in the systems, system's transparency and user satisfaction.

ACKNOWLEDGMENTS

The authors would like to thank Joris Klerkx for making available the source code of his framework, Aduna clustermap. In addition, we thank all the participants of the user studies for their participation and useful feedback.

REFERENCES

1. Baeza-Yates, R.A., Ribeiro-Neto, B. Modern Information Retrieval. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, 1999
2. Bennett, J. and Lanning, S.. The Netflix Prize. *Communications of the ACM*, 52, 10 (2007), 8 pages. Citeseer.
3. Blei, D.M., Ng, A.Y. and Jordan, M.I. Latent dirichlet allocation. *The Journal of Machine Learning Research*, 3, (2003), 993-1022
4. Bostandjiev, S., O'Donovan, J. and Höllerer, T. TasteWeights: a visual interactive hybrid recommender system. In *Proceedings of the sixth ACM conference on Recommender systems* (RecSys '12). ACM, New York, NY, USA, 2012, 35-42.
5. Burke, R. Hybrid Recommender Systems: Survey and Experiments. *User Modeling and User Adapted Interaction*, 12, 4 (2002), 331–370.
6. Burke, R. Evaluating the dynamic properties of recommendation algorithms. In Proceedings of RecSys10: fourth ACM conference on Recommender systems, pages 225–228, New York, NY, USA, 2010.
7. Chuah M. Dynamic aggregation with circular visual designs. In *Proceedings of the IEEE Symposium on Information Visualization* (1998), pp. 35–43
8. Gretarsson, B., O'Donovan, J., Bostandjiev, S., Hall, C. and Höllerer, T. SmallWorlds: Visualizing Social Recommendations. *Comput. Graph. Forum*, 29, 3 (2010), 833-842.
9. Gou, L., You, F., Guo, J., Wu, L., and Zhang, X. SFViz: interest-based friends exploration and recommendation in social networks. In *Proceedings of the 2011 Visual Information Communication - International Symposium* (VINCI '11). ACM, New York, NY, USA.
10. Herlocker, J.L., Konstan, J.A. and Riedl, J. Explaining collaborative filtering recommendations. In *Proceedings of CSCW00: ACM conference on Computer supported cooperative work*, 2000, 241–250
11. Herlocker, J.L., Konstan, J.A., Terveen, L.G. and Riedl, J.T. Evaluating collaborative filtering recommender systems. ACM Trans. Inf. Syst., 22, 1 (2004), 5–53
12. Hu, Y., Koren, Y. and Volinsky, C. Collaborative filtering for implicit feedback datasets. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining* (2008), 263–272. IEEE Computer Society, Washington, DC
13. Kammerer, Y., Nairn, R., Pirolli, P., & Chi, E. H. (2009). Signpost from the masses: learning effects in an exploratory social tag search browser. In *Proc. of CHI '09*, (2009) 625-634.

14. Klerkx, J. and Duval, E.. Visualising social bookmarks. *Journal Of Digital Information*, 10, 2, 2009, 40 pages. Multimedia Research Group, U. of Southampton.

15. Knijnenburg, B.P., Reijmer, N.J.M. and Willemsen, M.C. Each to his own: how different users call for different interaction methods in recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems* (RecSys '11) (2011), 141-148

16. Knijnenburg, B.P., Bostandjiev, S., O'Donovan, J. and Kobsa, A. Inspectability and control in social recommenders. In *Proceedings of the sixth ACM conference on Recommender systems* (RecSys '12) (2012), 43-50.

17. Linden, G., Smith, B, and York, J. Amazon.com recommendations: item-to-item collaborative filtering. *IEEE Internet Computing*, 7, 1 (Jan/Feb 2003), 76- 80

18. Manning, C.D., Raghavan, P., Schtze, H. Introduction to Information Retrieval, Cambridge University Press, New York, NY, 2008

19. Morita, M. and Shinoda, Y. Information Filtering Based on User Behavior Analysis and Best Match Text Retrieval. In *SIGIR 1994: Proceedings of the 17th Annual International ACM SIGIR Conference* (1994), 272–281. Springer-Verlag New York, Inc., New York

20. Nageswara, R.K. and Talwar, V.G. Application domain and functional classification of recommender systems a survey. *DESIDOC Journal of Library and Information Technology*, 28, 3 (2008), 17–35

21. O'Donovan, J., Smyth, B., Gretarsson, B., Bostandjiev, S., and Höllerer, T. PeerChooser: visual interactive recommendation. In *Proceedings of the twenty-sixth conference on Human factors in computing systems* (CHI '08). ACM, NY, USA, 2008, 1085-1088.

22. Parra, D., Jeng, W., Brusilovsky, P., Lopez, C., Sahebi, S. Conference Navigator 3: An Online Social Conference Support System. Poster at UMAP 2012. Montreal, Canada.

23. Parra, D. and Brusilovsky, P. Collaborative filtering for social tagging systems: an experiment with CiteULike. In *Proceedings of the third ACM conference on Recommender systems* (RecSys '09). ACM, New York, NY, USA (2009), 237-240.

24. Pazzani, M.J. and Billsus, D. Content-Based Recommendation Systems. In P. Brusilovsky, A. Kobsa, and W. Nejdl, editors, *The Adaptive Web*, volume 4321 of LNCS (2007), 325–341. Springer-Verlag, Springer.

25. Riedl, J. and Dourish, P. Introduction to the special section on recommender systems. *ACM Trans. Comput.-Hum. Interact.* 12, 3 (Sept.2005), 371-373.

26. Tintarev, N. and Masthoff, J. Designing and Evaluating Explanations for Recommender Systems. *Recommender Systems Handbook*, (2011), 479-510

27. Tonkin, E., Corrado, E. M., Moulaison, H. L., Kipp, M. E. I., Resmini, A., Pfeiffer, H. D. & Zhang, Q.. Collaborative and Social Tagging Networks. *Ariadne*, 54. (2008).

28. Zhao, S., Zhou, M.X., Yuan, Q., Zhang, X., Zheng, W., and Fu, R. Who is talking about what: social map-based recommendation for content-centric social websites. In *Proceedings of the fourth ACM conference on Recommender systems* (RecSys '10). ACM, New York, NY, USA, 2010, 143-150.