

# The Effect of Different Set-based Visualizations on User Exploration of Recommendations

Katrien Verbert<sup>1</sup>, Denis Parra<sup>2</sup>, Peter Brusilovsky<sup>3</sup>

<sup>1</sup>Departement Computerwetenschappen, KU Leuven, Celestijnenlaan 200A, B-3001 Leuven, Belgium  
katrien.verbert@cs.kuleuven.be

<sup>2</sup>Department of Computer Science, Pontificia Universidad Católica de Chile, Santiago, Chile  
dparra@ing.puc.cl

<sup>3</sup>School of Information Sciences, University of Pittsburgh, 135 North Bellefield Avenue,  
Pittsburgh, PA 15260, USA, peterb@pitt.edu

## ABSTRACT

When recommendations fail, trust in a recommender system often decreases, particularly when the system acts like a “black box”. To deal with this issue, it is important to support exploration of recommendations by explicitly exposing relationships that can provide explanations. As an example, a graph-based visualization can help to explain collaborative filtering results by representing relationships among items and users. In our work, we focus on the use of visualization techniques to support exploration of *multiple* relevance prospects - such as relationships between different recommendation methods, socially connected users and tags. More specifically, we researched how users explore relationships between such multiple relevance prospects with two set-based visualization techniques: a clustermap and a Venn diagram. A comparative analysis of user studies with these two approaches indicates that, although effectiveness of recommendations increases with the use of a clustermap, the approach is too complex for a non-technical audience. A Venn diagram representation is more intuitive and users are more likely to explore relationships that help them find relevant items.

## 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

The design and development of user interfaces for recommender systems has gained increased interest. Such interfaces are researched to provide new capabilities to search, browse, and understand recommendations [8]. Among others, explaining recommendations to provide transparency and to increase trust has been researched extensively [11]. Several approaches have been presented that represent relationships between users and items as a basis to support exploration and transparency [5][3][12].

Most of these existing approaches enable users to explore relationships between two entities, such as relationships between *users* and recommended *items*.

In our work, we focus on the use of set-based visualization techniques to support exploration of *multiple* relevance prospects. In contrast to existing approaches, we enable end-users to interrelate multiple dimensions to support exploration and transparency of recommendations.

We have developed two visual interfaces for exploring relationships between multiple relevance prospects of recommendations. A first user interface (TalkExplorer) uses a clustermap visualization technique that enables users to explore relationships between diverse *recommendations*, *users* and *tags*. A second interface (SetFusion) uses a Venn diagram to support exploration of multidimensional relationships.

The original work on TalkExplorer [12] and SetFusion [7] has been performed independently, with no intention to compare the results of our studies with these sufficiently different systems. At the same time, an extensive set of data collected in several user studies opened an interesting opportunity to uncover the participation puzzle that we observed when comparing the results of two TalkExplorer studies. These results indicate that effectiveness of recommendations increases in a significant way when users are able to interrelate multiple entities. However, when deployed in an open setting, users do not explore such intersections often when a clustermap is used.

In earlier work, we hypothesized that the likely reason for this phenomenon is the complexity of the TalkExplorer interface, especially the challenge of understanding the complex visualization of overlapping sets. However, at that time we were not able to provide any arguments in favor of these hypotheses.

In this paper, we re-assess this hypothesis. The presence of the SetFusion study that was performed in the same system, with similar kind of data and using a similar approach, enables us to compare how users interact with the visualizations. SetFusion explores exactly the kind of interface that we believe could increase exploration of

overlaps: Venn Diagrams are known to be both straightforward and standard to visualize set overlaps. We re-process the data of our user studies and analyze how users interact with both interfaces to re-assess their value for exploring recommendations.

This paper is organized as follows: first we present related work in the area of visualizing recommendations and set-based visualization. Then, we introduce TalkExplorer, an interactive clustermap visualization of recommendations, as well as SetFusion, an interactive Venn diagram representation of recommendations. Results of user studies conducted with both interfaces are presented next. Then, we present a comparative analysis of how users interact with these visualizations. Finally, we discuss these results, as well as future research opportunities.

## RELATED WORK

Most existing work in the area of visualizing recommendations focuses on interaction with collaborative filtering recommender systems. PeerChooser [5] is a visual interactive recommender that uses a graph-based representation to show relationships between users and recommended items of a collaborative filtering recommender system. Similarly, SmallWorlds [3] allows exploration of relationships between recommended items and similar 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 relevant items.

Some systems focus specifically on tags that are used by social recommenders. SFViz (Social Friends Visualization) [4] 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 to visualize a tag tree and a circle layout with edge bundling to show a social network.

More recently, TasteWeights [2] 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. The idea can be traced back to work of Schafer et al. [9] on meta-recommendation systems. These meta-recommenders provide users with personalized control over the generation of recommendations by indicating how important specific factors are – such as genre of a movie and film length, on a scale from 1 (not important) to 5 (must have). In our work, we extend this concept by visualizing relationships to relevance prospects in order to enhance exploration by end-users of the item space and to increase perceived relevance and meaning of items. More specifically, we use a set-based visualization approach to

represent relationships of items to specific relevance factors or prospects. Thus, in addition to enabling end-users to specify which prospects are relevant, we enable them to see how recommendations are related to these prospects with set-based visualization techniques.

Relevance or set-based visualization applies an approach to spatially organize recommendation results. Relevance-based visualization has been originally developed in the field of information retrieval for visualization of search results. For example, for a query that uses three terms, it will create seven set areas to show which results are relevant to each of the three terms, each of two pairs of these terms, and all three terms at the same time. The classic example of set-based relevance visualization is InfoCrystal [10]. The Aduna clustermap visualization [1] approach also belongs to this category offering a more complex visualization paradigm and a better level of interactivity. A strong point of set-based approach is a clear representation to which of the query terms each document is relevant along with grouping documents by this aspect.

The novelty of the approach suggested in our paper is twofold. First, we are using a set-based relevance approach not just with keywords or tags where relevance approaches are usually applied, but with a diverse set or relevance-bearing entities (tags, users, recommendation agents). To the best of our knowledge, this is the first attempt to visually represent recommendations with set-based visualization techniques. The major difference and innovation of our work is that we allow end-users to combine *multiple* relevance prospects in order to increase the perceived relevance and meaning of recommendations. Second, we present two different techniques to visually present these sets: a clustermap visualization, implemented in TalkExplorer [12], and a Venn diagram, implemented in SetFusion [7]. Although the interactive hybrid recommender interface TasteWeights [2] and meta-recommendation systems [9] also allow users to consider three potential sources of relevance to make recommendations, TalkExplorer allows more flexible exploration by visually presenting relationships to relevance prospects with a clustermap, and SetFusion uses a completely different visualization paradigm, relying on a Venn diagram. We present results of user studies with these visualizations that assess the impact of the interfaces on the effectiveness of recommendations, as well as a comparative analysis of how users interact with these representations.

## TALKEXPLORER AND SETFUSION

TalkExplorer and SetFusion represent two attempts to implement a visual interactive interface to explore recommendations of research talks at academic conferences. Both visualization interfaces were implemented and released as components of the conference support system *Conference Navigator 3 (CN3)* [6]. Each of the interfaces was developed to explore a range of ideas related to visualization, interactive access, transparency,

etc. One of the core ideas essential for the purpose of this paper was integration of several aspects of relevance within the same visualization. We believed that a talk might be perceived by users as relevant for a range of reasons that we call aspects (for example, it could be recommended by one of the recommender engines or bookmarked by a socially connected user). We also believed that talks that are relevant in more than one aspect could be more valuable to the users and that displaying multiple aspects of relevance visually is important for the users in the process of talk exploration. Following these beliefs, TalkExplorer and SetFusion offered two different approaches to visualize talk relevance in a way that helps to identify talks that are relevant for the users in two, three, and even more aspects. Both systems use different versions of set-based visualizations to achieve this goal. The user studies that we ran with both interfaces included specific provisions that enabled us to examine the value of displaying several aspects of relevance. The next sections explain the details of both visualization approaches and results of their evaluation that are relevant for this paper.

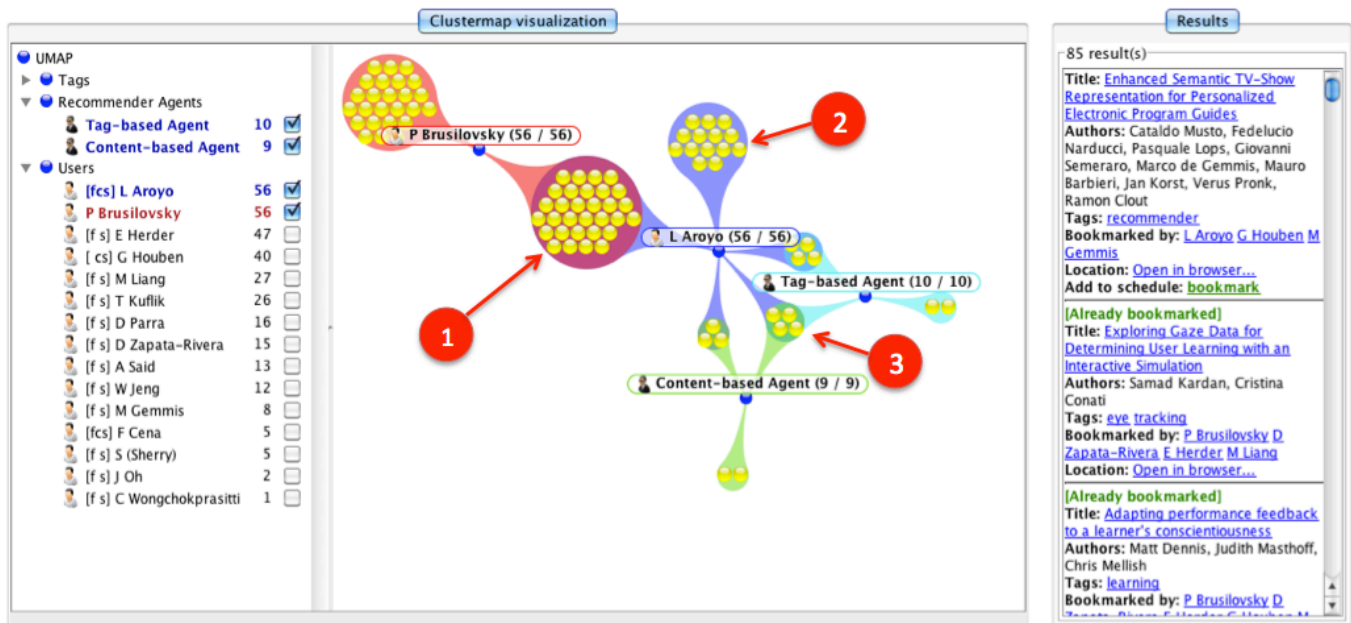
### VISUALIZING RELATIONSHIPS IN TALKEXPLORER

The key idea of TalkExplorer is to enable users to explore talks recommended by two recommender engines (presented in the interface as recommender agents) along with talks that were bookmarked or tagged by other system users. The visualization is implemented as a Java applet and

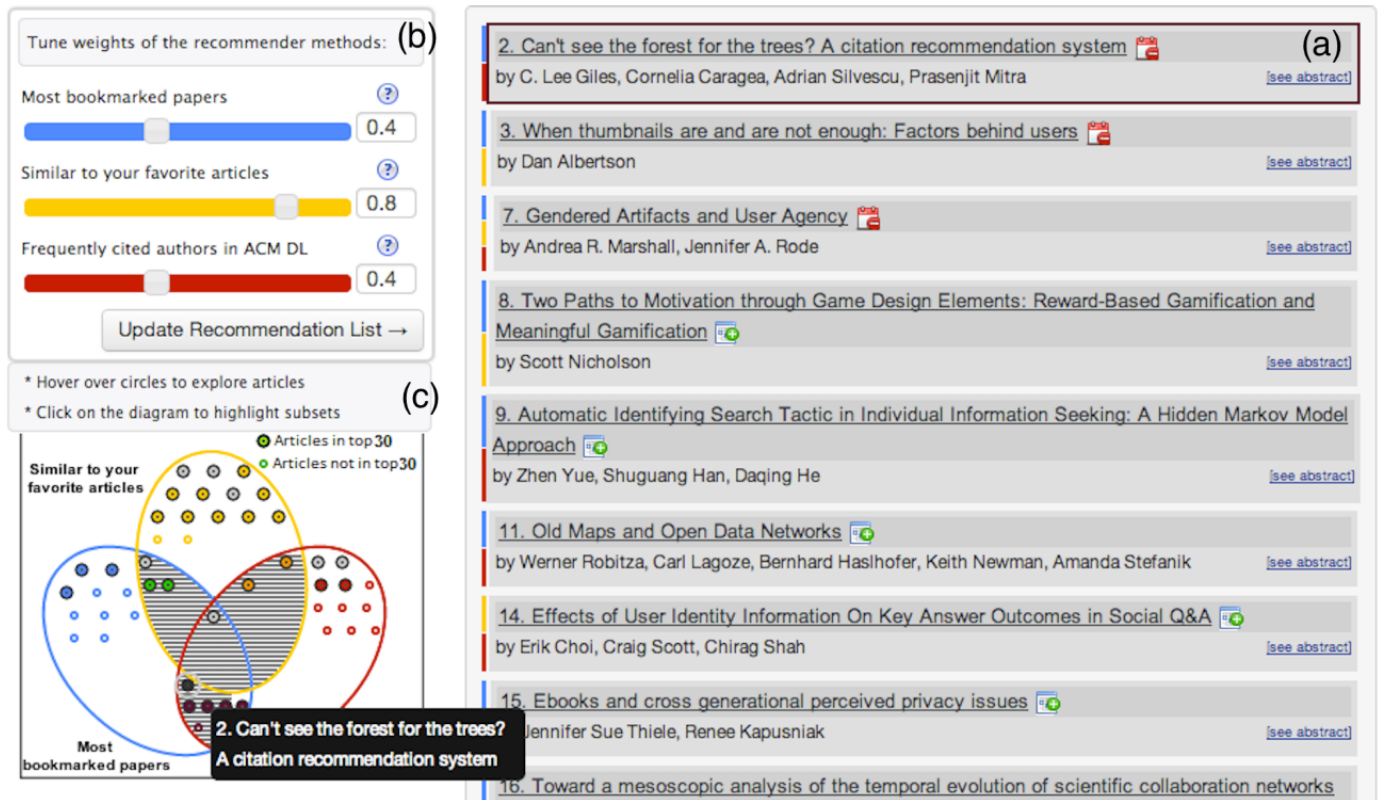
uses the Aduna clustermap visualization library [1]. This software library visualizes sets of categorized objects and their interrelationships.

Recommender systems are presented as *agents* and their interrelationships 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 by one or more agents). 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.

TalkExplorer allows users to explore the different entities of the conference by means of three principal components, as shown in Figure 1. 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 screen, 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.



**Figure 1:** Screenshot of TalkExplorer. Labeled numbers indicate clusters of talks (yellow circles) which are the result of intersecting talks bookmarked or tagged by real users, or suggested by recommender agents.



**Figure 2:** Screenshot of SetFusion displaying (a) a filtered list of papers recommended, (b) sliders, and (c) the Venn diagram

In Figure 1, two users are shown (P Brusilovsky and L Aroyo), as well as suggestions of the tag-based and content-based recommender agent. The clustermap 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 by a recommender agent by exploring the intersection of the agent and a specific user. In the example presented in Figure 1, the active user (P Brusilovsky) can explore which of the talks he has bookmarked are also bookmarked by user L Aroyo (label 1), which additional talks are bookmarked by L Aroyo but not recommended by an agent (label 2) and which talks are recommended to him by both the content-based and tag-based agent and are also bookmarked by L Aroyo (label 3) - to further filter out the potentially more relevant recommendations.

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. If a user clicks on a cluster (for example, the cluster showing talks that were bookmarked by L Aroyo and a specific agent) the list of these talks and their details are presented.

#### **VISUAL HYBRID RECOMMENDATION IN SETFUSION**

SetFusion is inspired by the same set-based approach than TalkExplorer, i.e., allowing users to choose items by combination of multiple prospects of relevance. The main

difference is that SetFusion uses a Venn diagram rather than a clustermap with links to show the intersections (fusions).

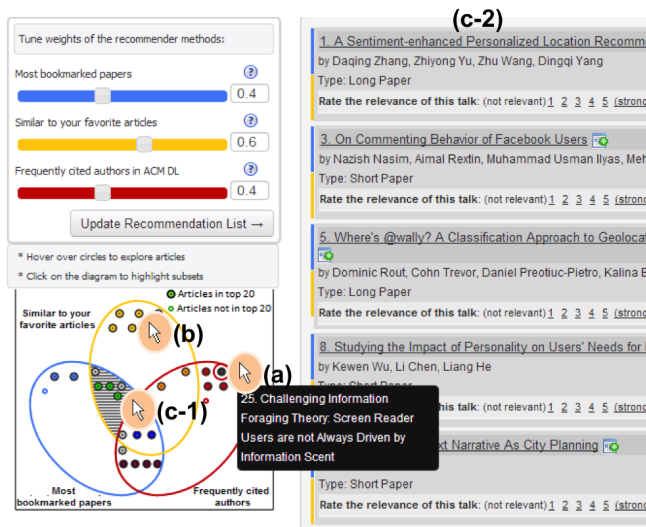
Another difference is the type of entities used as relevance prospects in order to support decision-making. While TalkExplorer uses tags, recommender agents and users, SetFusion mixes three recommendation methods, turning it into a hybrid recommender. The methods that SetFusion allows the user to combine are:

- *Most bookmarked papers:* this method recommends papers based on their popularity, i.e., papers that receive more bookmarks are ranked at the top.
- *Similar to your favorite articles:* this is a content-based recommendation method that considers the papers already bookmarked by the user to create a bag-of-words user profile. With this profile, the method matches the most similar non-bookmarked papers by cosine similarity. In order to make this method more effective, we tuned it using 10-fold cross validation and the final parameters considered filtering out terms with frequency less than three, appearing on less than two documents, and with a minimum length of four letters.
- *Frequently cited authors in the ACM Digital Library:* In this method, we recommended papers based on the popularity of their authors. Papers with authors that have been frequently cited in the ACM digital library are ranked at the top.

In SetFusion, users are provided with certain level of control over these methods: they can tune the importance of each prospect of relevance by adjusting their weight through sliders (Figure 2.b), an interaction method inspired by TasteWeights [2]. Despite these differences, the list of recommended items in SetFusion (Figure 2.a) can be filtered in a similar way to TalkExplorer, by clicking on the ellipse areas or their intersections (Figure 2.c).

Finally, users can interact with the Venn diagram as an inspection and filtering mechanism:

- Hover over the circle:* Each small circle represents a talk, and hovering over one of them displays a dialog with the title of the talk (Figure 3.a).
- Click on a circle:* By clicking in a small circle, the user will highlight the same element in the list of talks at the right side (Figure 3.b).
- Click on a Venn diagram area:* Users can also click on the area surrounded by the big ellipses with white background, and by clicking on such an area, the visualization will become shaded as in Figure 3.c-1 and it will filter the list on the right side to the selected items (Figure 3.c-2).



**Figure 3:** User interactions available on the Venn diagram of the SetFusion interface.

### USER STUDIES OF TALKEXPLORER

We have conducted two user studies with TalkExplorer. In the first study, we conducted a controlled experiment with users at two conferences (ACM Hypertext 2012 and UMAP 2012). The number of participants was 21. Users were asked to perform three tasks (exploring users, exploring agents and exploring tags). We recorded the screen and captured think aloud data. This controlled experiment enables to gain first insights into the relative effectiveness of each of these entities and to collect user feedback. Users had high familiarity with visualization techniques (mean 4.2, std. deviation 0.7) and a relatively high familiarity with

recommendation techniques (mean 3.7, std. deviation 0.9). Details of this study have been reported in [12].

In the second study (N=18), we have deployed TalkExplorer again at two conferences and asked users to explore the visualization without any specific tasks. Users were free to interact with the visualization and were not required to use any specific components or controls. With this second study, we expected to gain insight into the usefulness of the visualization in an open setting. We wanted to find out how users explore and use the visualization without guidance and what attracts their interests. The analysis of interaction patterns yields less biased data, as users were not constrained to three separate and fixed tasks. In addition, the study was conducted at two conferences in the Technology Enhanced Learning field (EC-TEL 2012 and LAK 2013). Conference attendees have less technical knowledge than participants of the UMAP and Hypertext conferences of the first study. Most of the participants have again knowledge visualization techniques (average 4.23, std. deviation 0.79), but familiarity with recommendation techniques was less high (average 3.15, std. dev. 1.23).

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, displayed in clusters, were explored.

We measured the *effectiveness* 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, the set of items of single entity (i.e. a user, a tag or a recommender agent) was explored 147 times by participants of study 1. Thirty-two of these sets were used to bookmark a new item. Thus, the effectiveness of exploring the set of items of a specific user is  $32/147=22\%$ .

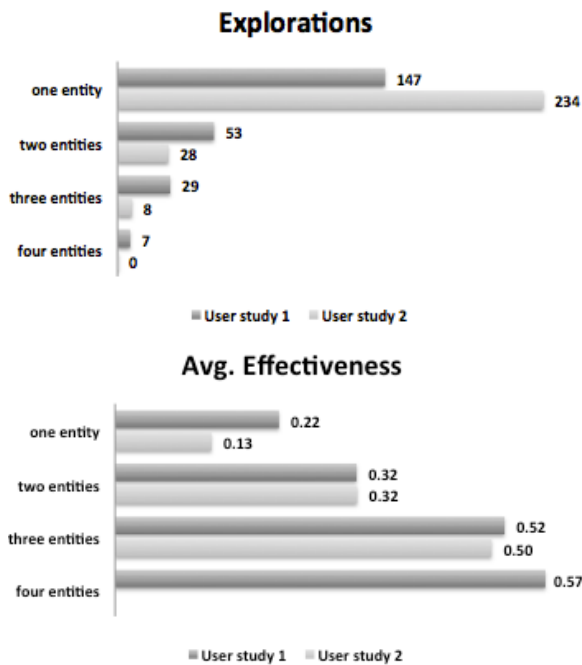
Effectiveness results are summarized in Figure 4. Overall, these results indicate that effectiveness of an explored set increases once more entities are integrated. More specifically, effectiveness increases from 22% (user study 1) and 13% (user study 2) when a single entity is used to 52% (user study 1) and 50% (user study 2) when three entities are used. Effectiveness is significantly higher when multiple entities are used in both studies (p-value 0.003 in study 1, 0.0009 in study 2). These results illustrate that enabling users to explore interrelationships between



prospects (sets of items in the overlap of entities) increases the probability of finding a relevant item.

Whereas both user studies demonstrated the clear value of multi-prospect visualization, we can't ignore one interesting difference. Despite the clear value offered by the intersection areas, the number of times that intersections were explored is lower in the second user study: items in the intersection of two entities were explored 28 times in the second user study (versus 53 times in the first user study) and items in the intersection of three entities were explored eight times (versus 29 in the first user study). Items in the intersection of four entities were not explored in the second study. The data are summarized in Figure 4.

Particularly the visualization of intersections of three or four entities seems to be non-intuitive or complex for end-users, as they do not tend to explore these intersections. In the first study, users explored these combinations more often and were more positive about the usefulness of this concept.



**Figure 4:** Summary results user study 1 and user study 2. To summarize, results of both studies illustrate the usefulness of visualizing multiple prospects. Users are interested to explore users, agents and tags and indicate that these multiple prospects are useful as a basis to find relevant talks. Exploring intersections increases effectiveness, but these intersections are not used often in an open setting.

A likely reason is the complexity of the TalkExplorer interface. A more intuitive way for exploring such overlapping sets are Venn diagrams, which are known to be both straightforward and standard to represent sets and set overlaps. In this paper, we are interested to explore whether it will help if we show overlaps in a more traditional and

easy to understand way. SetFusion explores exactly the kind of interface that we believe could increase explorations of overlaps. We present user study results of SetFusion in the next section.

### USER STUDIES OF SETFUSION

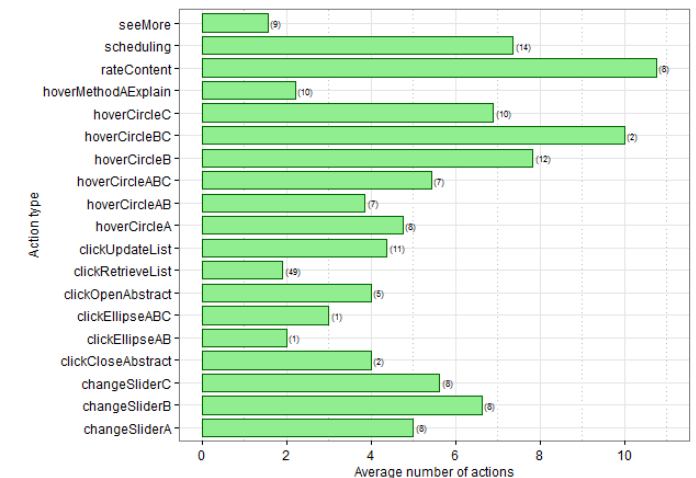
In order to test whether the more intuitive representation of the Venn diagram had an effect on increasing CN3 users' engagement and effectiveness with the interface, we conducted a field study using SetFusion to recommend papers during the UMAP 2013 conference. In this study, users were free to access and explore the visualization.

The analysis of user participation and engagement data (Table 1) shows a good effectiveness of the interface in turning user exploration into bookmarked papers. The fraction of users who tried the SetFusion interface among those having a chance to use it was over 50% (50/95).

Metric	SF UMAP13
# Users exposed to recommendations	95
# Users who used recommender page	50
# Users who bookmarked	14
# Talks explored (user avg.)	14.9
# Talks bookmarked / user avg.	103 / 7.36
# People returning to recommender page	14 (28%)
Average time spent in page (seconds)	353.8

**Table 1:** Participation and engagement metrics in the SetFusion interface at UMAP13 conference.

The average number of each type of action in SetFusion during the UMAP 2013 field study is summarized in Figure 5. In parenthesis, the amount of users for each action is shown.



**Figure 5:** Average number of each type of action in SetFusion during the UMAP 2013 field study. In parenthesis, the amount of users performing those actions is shown.

These users explored 14.9 papers on average and bookmarked 7.36 papers, indicating a good level of effectiveness of the interface. Among the users that tried the SetFusion interface, 28% (14 users) bookmarked at least one paper. The same percentage of users came back to SetFusion page for a second time or more. If we consider the total time that users spent on the page among one or more sessions, users spent on average around 6 minutes (353.8 seconds) on the interface. More detailed study results are reported in [7].

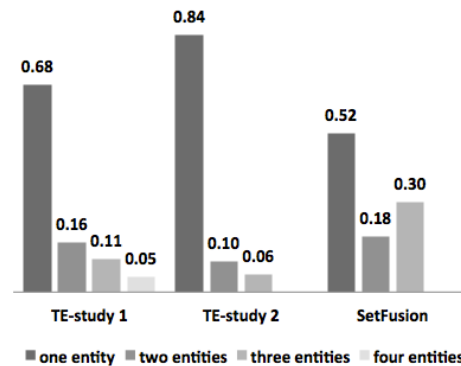
### META-ANALYSIS

The original work on TalkExplorer and SetFusion has been performed independently, with no intention to compare the results of our studies with these sufficiently different systems. At the same time, an extensive set of data collected in the mentioned studies opened an interesting opportunity to uncover the participation puzzle that we observed when comparing the results of two TalkExplorer studies. As presented above, results of our TalkExplorer studies indicate that effectiveness of recommendations increases in a significant way when users are able to interrelate multiple entities (see Figure 4). However, when deployed in an open setting, users do not explore such intersections often when a clustermap is used. The original paper that presents our work on TalkExplorer hypothesized that the likely reason for this phenomenon is the complexity of the TalkExplorer interface, especially the challenge of understanding the complex visualization of overlapping sets. While the Aduna visualization approach is very powerful and makes it possible to present multiple subsets created by overlapping three, four, five and more sets, understanding the picture is a real challenge. We suggested that this leads to the lower use of overlaps in the second study where the users were not specifically requested to do it. We also speculated that the “free” usage of overlaps could be increased when the users get more experience or when a simpler and more traditional visualization such as Venn diagrams will be used. However, at that time we were not able to provide any arguments in favor of these hypotheses.

The presence of the SetFusion study that was performed in the same system, with similar kind of data and using a similar approach, enabled us to re-assess this hypothesis. Indeed, SetFusion explored exactly the kind of interface that we believed could increase the usage of talks that are relevant for more than one prospect. Venn Diagrams are known as both a straightforward and a standard way (i.e., used in high school math classes) to visualize set overlaps. In this context, by re-processing the data of SetFusion study, we could provide some ground behind our complexity hypothesis. Below we present our attempt to re-process the data of the SetFusion study and present it in comparison with the data of the TalkExplorer study.

Figure 6 compares the number of times that sets were explored in all the presented user studies. TE-study 1 is the

first (controlled) user study that we conducted with TalkExplorer. TE-study 2 is the second study with TalkExplorer that was conducted in an open setting: i.e. users were free to explore the visualization. Sets of a single entity were explored most in both studies: 147 times or 68% on average in the first study and 234 times or 84% on average in the second study. Sets representing items in the intersection of two entities were explored less often: 53 times or 16% in study 1, 28 times or 10% in study 2. Whereas items in intersections of three entities were still explored relatively often in study 1 (29 times or 11%), exploration of such sets was rare in the second user study: users explored intersections of three entities only eight times (6% on average). Intersections of four entities were not explored in the second study.



**Figure 6:** comparative analysis avg. number of explorations

Results of SetFusion draw a different picture. With a traditional Venn diagram, users explored items of a single entity in 52% of interactions. 18% of the interactions were explorations of intersections of two entities and 30% were explorations of intersections of three entities. It means that the use of two-entity overlap was higher than in the first TalkExplorer study where the users were specifically asked to do so. The use of three-entity overlap was almost three times higher than in the first controlled TalkExplorer study and five times more than in the second “free” study (TE-study 2).

Thus, our results indicate that with a more intuitive representation, the use of multiple relevance prospects is high even in a free exploration context where the users are not specifically required to use overlaps. There is no real difference between explorations of a single entity (52%) versus multiple entities (18%+30%=48%). Items in the intersection of three entities were explored more often than items in the intersection of two entities – which is an interesting result as such combinations were most effective for finding relevant items in our TalkExplorer studies.

The Venn diagram visualization therefore seems more promising than the clustermap visualization. As multiple entities increase effectiveness of recommendations, the approach would help users to explore those sets that help them find the more relevant items. A drawback of the

approach is that it is typically limited to three entities, whereas a clustermap enables to interrelate more than three entities. Despite this functionality, users did not explore such intersections in our second TalkExplorer study.

In summary, as results of our TalkExplorer study indicate that effectiveness of recommendations increases when multiple entities are interrelated, the Venn diagram approach is likely to better support our hypotheses. The data of the SetFusion study indicates that the approach is more intuitive for users – especially for interrelating multiple entities.

### CONCLUSION AND FUTURE WORK

In this paper, we have presented two approaches that enable end-users to explore recommendations. Both approaches allow end-users to combine *multiple* relevance prospects in order to increase the perceived relevance and meaning of recommendations. The first approach uses a clustermap representation and has been implemented in TalkExplorer. The second approach uses a Venn diagram and has been implemented in SetFusion.

In our user studies of TalkExplorer, we were able to show that effectiveness of recommendations increases significantly when multiple entities are interrelated. However, the clustermap visualization of TalkExplorer seems too complex to use. Users do not tend to explore those intersections that will help them find the more relevant items in an open setting. To make the power of overlaps work in a realistic context, the interface should be easy to understand. Venn diagrams are likely to be a good candidate, as they are known to be straightforward and a standard way for representing set overlaps. By re-processing the data of our SetFusion study that embodies exactly this kind of representation, we were able to show that users explore these intersections frequently. As indicated above, this exploration of overlaps is key, as it helps users to find the items that are likely to be more relevant to them.

In follow up studies, we will leverage this evidence and research more intuitive ways to support exploration of intersections. A follow up study will also include multiple agents (so far, only two agents were shown to the user) and assess the added value of our visualization on top of larger data collections.

### ACKNOWLEDGMENTS

We thank all the participants of the user studies for their participation and useful feedback. Research of Katrien Verbert was supported by postdoctoral fellowship grant of the Research Foundation – Flanders (FWO).

### REFERENCES

1. Aduna clustermap. <http://www.aduna-software.com/technology/clustermap>
2. 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.
3. Gretarsson, B., O'Donovan, J., Bostandjiev, S., Hall, C. and Höllerer, T. SmallWorlds: Visualizing Social Recommendations. *Comput. Graph. Forum*, 29, 3 (2010), 833-842.
4. 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 (2011), 10 pages.
5. 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.
6. Parra, D., Jeng, W., Brusilovsky, P., Lopez, C. and Sahebi, S. Conference Navigator 3: An Online Social Conference Support System. Poster at UMAP 2012. Montreal, Canada.
7. Parra, D., Brusilovsky, P., and Trattner, C. See what you want to see: visual user-driven approach for hybrid recommendation. In *Proceedings of IUI '14*. ACM, New York, NY, USA (2014) 235-240.
8. Riedl, J. and Dourish, P. Introduction to the special section on recommender systems. *ACM Trans. Comput.-Hum. Interact.* 12, 3 (Sept.2005), 371-373.
9. Schafer, J.B., Konstan, J. A., and Riedl, J. Meta-recommendation systems: user-controlled integration of diverse recommendations. In *Proceedings of the eleventh international conference on Information and knowledge management*, ACM, New York, NY, USA (2002), 43-51.
10. Spoerri, A. InfoCrystal: A visual tool for information retrieval & management. In *Proceedings of the second international conference on Information and knowledge management*, ACM (1993), 11-20.
11. Tintarev, N. and Masthoff, J. Designing and Evaluating Explanations for Recommender Systems. *Recommender Systems Handbook*, (2011), 479-510
12. Verbert, K., Parra, D., Brusilovsky, P., and Duval, E. (2013, March). Visualizing recommendations to support exploration, transparency and controllability. In *Proceedings of IUI'13*, ACM (2013), 351-362
13. 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.