IntersectionExplorer: the Flexibility of Multiple Perspectives

Bruno Cardoso
Dept. of Computer Science
KU Leuven
Celestijnenlaan 200A
3001 Heverlee, Belgium 3001
bruno.cardoso@cs.kuleuven.be

Peter Brusilovsky
School of Information Sciences,
University of Pittsburgh
Pittsburgh, PA, USA
peterb@pitt.edu

Katrien Verbert
Dept. of Computer Science
KU Leuven
Celestijnenlaan 200A
3001 Heverlee, Belgium 3001
katrien.verbert@cs.kuleuven.be

ABSTRACT

Recommender systems are currently an ubiquitous presence on the web, helping us find relevant items in the ever-growing plethora of information available. However, there is not a one-size fits-all for recommender systems, and flexibility and control are crucial for enabling the possibility of adapting the recommender system to different user preferences. In this paper, we present the results of a study designed to assess user interaction with IntersectionExplorer (IEx), a multi-perspective tool for exploring conference paper recommendations. The study was conducted at the Digital Humanities 2016 Conference, an event with a rather large, heterogeneous, and not technology-oriented audience. The results obtained indicate that the IEx multi-perspective approach lends enough flexibility to accommodate different user preferences. When contrasting these results with a previous study conducted at a conference with a highly technological audience, it becomes apparent that the flexibility of IEx is key to empower users with different profiles to customize their approach to finding relevant recommendations.

CCS CONCEPTS

•Information systems →Information systems applications;

KEYWORDS

Recommender Systems, User Interfaces, User Study

ACM Reference format:
Bruno Cardoso, Peter Brusilovsky, and Katrien Verbert. 2016. Intersection-Explorer: the Flexibility of Multiple Perspectives. In *Proceedings of Joint Workshop on Interfaces and Human Decision Making for Recommender Systems, Como, Italy, August 27, 2017,* 4 pages.
DOI:

1 INTRODUCTION

Recommender systems are nowadays a common fixture in many environments like the web, where they play a pivotal role in helping us find our way through the ever more dense information jungle [7]. However, there is evidence that user trust tends to be lost when recommendations fail, particularly when users can not understand the rationale for those recommendations - the "black box" issue. There are, of course, many ways of addressing this problem, ranging from textual explanations to more elaborate, visual approaches like TasteWeights [2].

Joint Workshop on Interfaces and Human Decision Making for Recommender Systems, Como, Italy

DOI:

In addition to the "black box" problem, other factors have an impact in how recommender systems perform with users (e.g., the "cold start" issue), and research indicates that the nature of the system itself and that of its users may also condition recommendation acceptance. Indeed, as Guy el al. [6] have noticed "*for some users, recommendations based on people work better, while for others, recommendations based on tags are more effective*". Addressing this need for flexibility in accommodating user's preferences and expectations (among other requirements), we developed and presented Intersection Explorer (IEx) in previous work [13].

IEx is a tool for exploring conference papers that proposes a different way of interacting with recommendations - through the exploration of multiple, intertwining *perspectives of relevance*. In this work we define "perspective of relevance" as an umbrella term encompassing the source and nature of recommendations. We identify three types of perspective, each one occupying its own place in IEx's user interface (UI): (1) the *perspective of personalized relevance*; (2) the *perspective of social relevance* and (3) the *perspective of content relevance*. The first of these perspectives is composed by sets of papers that have been suggested by different recommendation engines: since recommender systems leverage previous knowledge about the user to provide suggestions that would likely fit his/her interests and goals, their suggestions are relevant mainly because they are personalized. The perspective of social relevance is composed by sets of papers that have been marked as relevant by other users of the system: if another user is perceived as like-minded, a collection of his/her items of interest may likely be considered as a set worth exploring. Finally, the perspective of content relevance is composed of sets of papers tagged by the community with the same keywords applied by the user. Since these keywords are usually drawn or derived from the contents or the experience of people with an item, they provide insightful glances about the contents of the tagged items. A key feature of IEx is the seamless way it allows users to combine sets from these three perspectives, making no distinction between them in terms of interaction or UI representation.

This approach lends IEx enough flexibility to allow its users to explore and combine recommendations based on human-generated data and produced by automatic agents in a seamless manner, all carrying the same potential weight and relevance. In order to understand if users do indeed leverage IEx's adaptability potential, we conducted a user study at the 2016 edition of the Digital Humanities (DH2016), a conference with a heterogeneous and not technology-oriented audience. We discuss the results of this study in this work and contrast our findings with those of a previous study [13] conducted with participants sampled from the audience of a technology-oriented event, the European Conference on Technology Enhanced Learning (EC-TEL2015).

2 RELATED WORK

Social recommendation based on people and tags has been researched extensively (e.g., [10]). For instance, SFViz (Social Friends Visualization) [5] visualizes social connections between users and their interests in order to increase awareness of others and thereby help people find potential friends with similar interests.

We can also find research focused on hybrid recommenders, i.e., systems involving different recommendation techniques in synergy. An interesting reflection on this approach was made by Guy et al. [6], who found that a hybrid people-tag-based recommender has a slightly higher accuracy than a tag or people-only approach. Other advantages are also mentioned in their work, such as "*low proportion of expected items, high diversity of item types, richer explanations*" and, as previously stated, *"the simple fact that for some users, recommendations based on people work better, while for others, recommendations based on tags are more effective"* [6]. Although we also combine different user-generated data sources in IEx, we do not merge them automatically into a hybrid recommender system. Instead, we empower users to select which users and tags they are interested in and also - akin to the idea of enabling users to switch between recommenders presented by Ekstrand et al. [4] - to choose which automatic recommendation agents' suggestions they want to explore.

Regarding visualization-based approaches, TasteWeights is a system designed to allow its users to control the influence of friends' and peers' profiles and behaviors on the recommendation processes and, like IEx, it features a UI for presenting and interacting with recommendations. The recommendation process is adapted at runtime by user-entered preference and relevance feedback. This idea can be traced back to the work of Schafer et al. [12] concerning meta-recommendation systems, where users are provided with personalized control over the generation of recommendations by altering the importance of specific factors on a scale from 1 to 5. In the same line, SetFusion [11] is another example that allows users to fine-tune the weights of a hybrid recommender system, representing relationships between recommendations through Venn diagrams. IEx extends these concepts by focusing on the visualization of relationships between perspectives of relevance, including human-generated data such as user bookmarks and community tags in addition to recommender outputs in a scalable, set-based visualization, the UpSet [8]. The UpSet is a visualization technique dedicated to the analysis of sets, their intersections, and aggregates of intersections. Set intersections are visualized in a matrix layout that enables the effective representation of associated data, such as the number of elements in set aggregates and intersections (see Figure 1, *Set Exploration View* callout).

3 INTERSECTIONEXPLORER (IEX)

As previously stated, IEx is a platform that allows for multi-perspective exploration of recommendations. An overview of its user interface is shown in Figure 1. IEx uses a simplified version of UpSet [8], a matrix-based visualization technique to represent sets and overlaps between sets. It is separated in three connected views (Figure 1, top green callouts).

The **Set Selection View** allows the user to select sets of recommendations from three different perspectives: the *Perspective of Personalized Relevance*, the *Perspective of Social Relevance* and the *Perspective of Content Relevance* (Figure 1, labels *a*, *b* and *c*, respectively). The Perspective of Personalized Relevance lists the papers suggested by different recommendation engines, the Perspective of Social Relevance is composed of papers that have been bookmarked by other users of the system and, finally, the Perspective of Content Relevance shows sets of papers labelled by the community with a specific tag. While the first perspective is clearly associated to automatic processes, the last two are based on human-generated data meaning that, in a sense, IEx's users play the role of "human recommenders".

In the **Set Exploration View** the user can explore all possible combinations between the sets selected in the Set Selection View. Sets of papers are represented as columns (the current user is highlighted in blue) and set combinations are depicted as rows (e.g., Figure 1, label *d*), where intersecting sets are represented as filled circles. The horizontal bar next to circle rows represents the relative (the row itself) and the absolute (the number by the row) amount of papers in the selected intersection. For example, the row selected in Figure 1 (the fourth row) indicates that there are 5 papers in common between the suggestions of the bookmark-based agent and papers bookmarked by the user named "User 1".

The **Intersection Exploration View** allows the user to explore the details and bookmark the papers contained in the selected intersection (Figure 1, label *e*). In the example of Figure 1, the user is exploring the 5 papers contained in the intersection represented by the fourth row of the Set Exploration View.

4 USER STUDY

4.1 Setup and Demographics

To provide IEx with data, we have deployed it on top of Conference Navigator 3 (CN3) [3]. CN3 is a social, personalized web-based system that supports academic conference attendees and suggests talks using different recommendation engines. In IEx's UI these engines' recommendations are metaphorized as "agents" and compose the Perspective of Personalized Relevance (Figure 1, label *a*). The engines are: (1) the *top-10* agent that suggests the 10 papers that have been bookmarked the most; (2) the *tag-based* agent that matches the tags assigned to papers by the current user to those of other users (using the Okapi BM25 algorithm [9]); (3) the *bookmark-based* agent models the user interest profile as a vector of terms with weights based on the TF-IDF statistic [1] using the contents of the papers bookmarked by the user; (4) the *external bookmark* recommender engine, that combines both the contents of the papers bookmarked by the user in CN3 and other social bookmarking systems like Mendeley, CiteUlike, or BibSonomy [14]; and finally, (5) the *bibliography* recommender engine uses the content of papers previously published by the user [14].

The CN3 also supplies IEx with data regarding other user's bookmarks and community-tagged papers, which respectively compose the latter's perspectives of social- and content-relevance. To address the well-known cold start problem, we requested participants

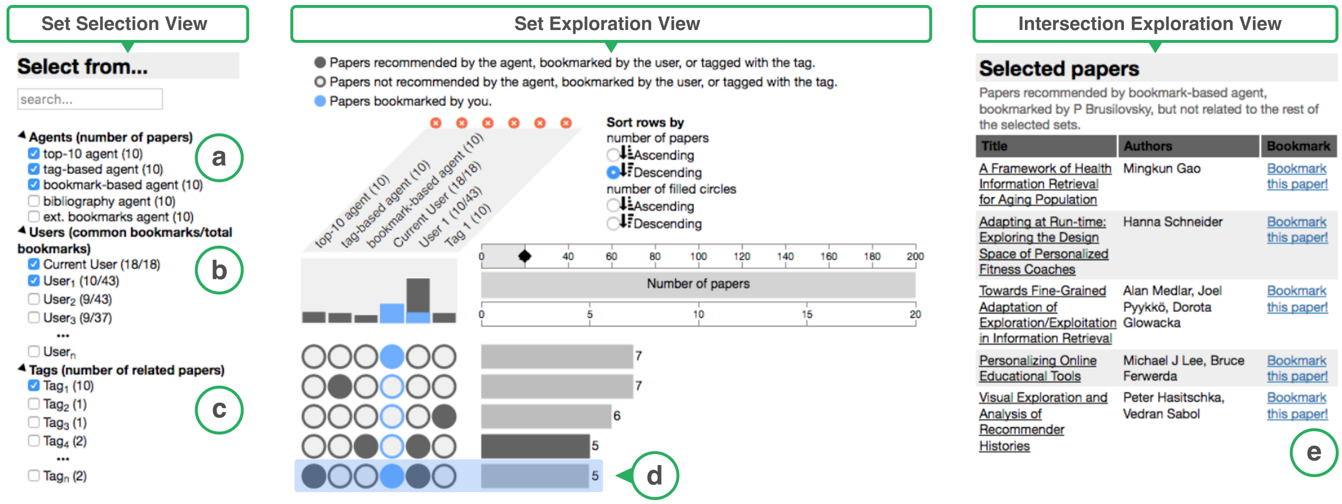

Figure 1: IEx's user interface, composed of three views (identified by the top green callouts): the *Set Selection View* lists the (a) recommendations of automatic agents (*Perspective of Personalized Relevance*), (b) the bookmarks of other users (*Perspective of Social Relevance*) and (c) papers tagged by the community (*Perspective of Content Relevance*); the *Set Exploration View* allows users to explore the (d) intersections between the selected sets of papers as rows; and, finally, the *Intersection Exploration View* displays the items (e) of the intersections clicked in the Set Exploration View, thereby allowing users to explore and bookmark the suggested papers.

to bookmark and tag a minimum of five papers from the conference proceedings its CN3 proceedings page.

In order to understand how flexible IEx's multi-perspective approach is, we conducted a user study at the DH2016 Conference, an event with a rather large, heterogeneous, and not technology-oriented audience, mainly composed of researchers from the areas of social sciences and humanities. We recruited 37 participants through direct invitation out of the DH2016 attendees, 11 female and averaging 38 years (SD: 10). For background, our previous EC-TEL2015 study had 20 participants, 3 female, averaging 32.9 years old (SD: 6.32).

Before starting the tests, all participants received the same presentation that introduced IEx, explained its functionality and covered its essential concepts. All participants were asked to perform the same task: to freely explore the DH2016's papers through IEx, and bookmark five relevant papers.

We collected data about participants' actions, like paper bookmarking actions and visualizations. To provide some definitions, we consider that a set of papers is "explored" when the user clicks on its respective row (Figure 1, *d*); that papers are "visualized" when they are listed in the Intersection Exploration View (Figure 1, *e*); and that a paper is "bookmarked" when the user clicks on the "Bookmark this paper" link that is adjacent to each visualized paper. In order to simplify our analysis, we define the metric *precision* as the fraction of papers that were visualized and bookmarked, across all users (e.g., if the user was to bookmark one paper out of five he/she visualizes, that would yield a precision of 1/5, or 0.2).

4.2 Results

In Table 1, we can see the results of participant interactions with agents, namely *single agents* (exploring the suggestions of a single agent), *multiple agents* (exploring the overlapping suggestions of more than one agent) and *augmented agents* (exploring the overlaps between the suggestion of agents and sets of papers from other perspectives). It is noticeable that in our DH2016 study single and augmented agents were explored the most, with comparable precision scores, while participants of our first study mainly explored the suggestions of multiple agents.

Table 1: Results for participant interaction with automatic recommendation agents (results of our first study between parentheses).

Agents	Bookmarks	Papers Viewed	Precision	Explorations
Single	41 (5)	196 (93)	**0.21** (0.05)	**31** (26)
Multiple	1 (15)	7 (166)	**0.14** (0.09)	**4** (40)
Augmented	15 (8)	63 (50)	**0.24** (0.16)	**37** (27)

Table 2 displays the results of single perspective explorations, i.e., explorations of the overlaps between one or more sets of papers from the same perspective. It is noteworthy how, in both studies, the perspective of content relevance yielded a noticeably higher precision than the other two perspectives.

Finally, Table 3 presents the results of perspective involvement in explorations. We consider that a perspective is involved when the user is exploring a combination containing at least one set of papers from that perspective. Once again, participants of our two studies were most likely to make a bookmark when the perspective of content-relevance was involved, i.e., when one ore more sets of tagged papers were combined with other sets.

Table 2: Interaction results for single-perspective explorations (results of our first study between parentheses).

	Bookmarks	Papers Viewed	Precision	Explorations
Agents	42 (20)	203 (259)	**0.21** (0.08)	**35** (66)
Users	49 (14)	335 (107)	**0.15** (0.13)	**41** (30)
Tags	44 (11)	94 (28)	**0.47** (0.39)	**80** (19)

Table 3: Precision scores for perspective involvement in explorations, across participants. The black square (■) represents perspective involvement (results of our first study between parentheses).

		Bookmarks	Papers Viewed	Precision	Explorations
Agents	■	57 (59)	267 (398)	**0.21** (0.15)	**73** (156)
		96 (25)	1383 (145)	**0.07** (0.17)	**133** (59)
Users	■	66 (45)	408 (239)	**0.16** (0.19)	**86** (119)
		87 (39)	1242 (304)	**0.07** (0.13)	**120** (96)
Tags	■	48 (25)	110 (71)	**0.44** (0.35)	**94** (56)
		105 (59)	1540 (472)	**0.07** (0.13)	**112** (159)

5 DISCUSSION

The results of our studies allow us to conclude positively about the flexibility of IEx's approach to accommodate different user preferences. Indeed, after the definition of "precision" that we make in this work (see section 4.1), our results indicate that the perspective of content-relevance (composed by sets of tagged papers) is the one accounting for the higher precision (see Table 2). This may be explained in light of the nature of this perspective, since well-applied tags provide accurate insights into the contents of the labeled items, and conference papers are interesting to readers mainly because of their content. Also, we found that there is a tendentially higher precision when sets of tagged papers are involved in explorations (see Table 3). Since this involvement implies that all explored papers are also community-tagged papers, this finding provides support to our previous observation.

Another interesting result reports to participant interaction with automatic recommendation agents (see Table 1). It is noticeable that while participants of our first study were mainly interested in the suggestions of multiple agents, those of our second study were not (respectively 40 vs. 4 explorations). In turn, while the precision was higher in our first study for augmented agents, the precision was the highest in our DH2016 study for single and augmented agent explorations. These findings suggest that IEx use data reflects the nature of its users, i.e., technology-oriented users prefer to explore the overlaps of automatic processes while less technology-oriented people were more interested in complementing the recommendations of automatic agents with sets of suggestions based on human-generated data - or, in other words, in having a human perspective over machine-produced recommendations.

These results can be extrapolated to conclude about the control that IEx lends to its users. Indeed, our platform seems to be flexible enough to allow them to select and explore the perspectives they judge the most productive and, what is perhaps more interesting, to mix them freely and discover new and customized approaches that fit best with their personal objectives.

6 CONCLUSIONS AND FUTURE WORK

Our results indicate that IEx's multi-perspective approach is a promising way of presenting recommendations to its users, flexible enough to adapt and allow them to follow their own path to trustworthy recommendations. For the future, it would be interesting to further challenge IEx in domains of application other than the recommendation of conference papers, and also with different audiences. While the UpSet is an effective way of presenting intersections between sets, its focus on information entails domain agnosticism. Therefore, different, multi-perspective visualizations may also be considered to bring IEx closer to its users.

ACKNOWLEDGEMENTS

The research has been partially financed by the KU Leuven Research Council (grant agreement no. C24/16/017).

REFERENCES

[1] Ricardo Baeza-Yates, Berthier Ribeiro-Neto, et al. 1999. *Modern information retrieval*. Vol. 463. ACM press New York.
[2] Svetlin Bostandjiev, John O'Donovan, and Tobias Höllerer. 2012. TasteWeights: a visual interactive hybrid recommender system. In *Proc. RecSys'12*. ACM, 35–42.
[3] Peter Brusilovsky, Jung Sun Oh, Claudia López, Denis Parra, and Wei Jeng. 2016. Linking information and people in a social system for academic conferences. *New Review of Hypermedia and Multimedia* (2016), 1–31.
[4] Michael D. Ekstrand, Daniel Kluver, F. Maxwell Harper, and Joseph A. Konstan. 2015. Letting Users Choose Recommender Algorithms: An Experimental Study. In *Proceedings of the 9th ACM Conference on Recommender Systems (RecSys '15)*. ACM, New York, NY, USA, 11–18. https://doi.org/10.1145/2792838.2800195
[5] Liang Gou, Fang You, Jun Guo, Luqi Wu, and Xiaolong Luke Zhang. 2011. Sfviz: interest-based friends exploration and recommendation in social networks. In *Proceedings of the 2011 Visual Information Communication-International Symposium*. ACM, 15.
[6] Ido Guy, Naama Zwerdling, Inbal Ronen, David Carmel, and Erel Uziel. 2010. Social media recommendation based on people and tags. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*. ACM, 194–201.
[7] Matevž Kunaver and Tomaž Požrl. 2017. Diversity in recommender systems–A survey. *Knowledge-Based Systems* (2017).
[8] Alexander Lex, Nils Gehlenborg, Hendrik Strobelt, Romain Vuillemot, and Hanspeter Pfister. 2014. UpSet: visualization of intersecting sets. *Visualization and Computer Graphics, IEEE Transactions on* 20, 12 (2014), 1983–1992.
[9] Christopher D Manning, Prabhakar Raghavan, Hinrich Schütze, et al. 2008. *Introduction to information retrieval*. Vol. 1. Cambridge university press Cambridge.
[10] Michael G Noll and Christoph Meinel. 2007. Web search personalization via social bookmarking and tagging. In *The semantic web*. Springer, 367–380.
[11] Denis Parra and Peter Brusilovsky. 2015. User-controllable personalization: A case study with SetFusion. *International Journal of Human-Computer Studies* 78 (2015), 43–67.
[12] J Ben Schafer, Joseph A Konstan, and John Riedl. 2002. Meta-recommendation systems: user-controlled integration of diverse recommendations. In *Proceedings of the eleventh international conference on Information and knowledge management*. ACM, 43–51.
[13] Katrien Verbert, Karsten Seipp, Chen He, Denis Parra, Chirayu Wongchokprasitti, and Peter Brusilovsky. 2016. Scalable exploration of relevance prospects to support decision making. In *Proceedings of the Joint Workshop on Interfaces and Human Decision Making for Recommender Systems co-located with ACM Conference on Recommender Systems (RecSys 2016)*. CEUR-WS, 28–35.
[14] Chirayu Wongchokprasitti. 2015. *Using external sources to improve research talk recommendation in small communities*. Ph.D. Dissertation. University of Pittsburgh.