

3-20-2019

Exploring User-Controlled Hybrid Recommendation in Conference Contexts

Chun-Hua Tsai

Behnam Rahdari

Peter Brusilovsky

Follow this and additional works at: <https://digitalcommons.unomaha.edu/isqafacpub>

Please take our feedback survey at: https://unomaha.az1.qualtrics.com/jfe/form/SV_8cchtFmpDyGfBLE

Exploring User-Controlled Hybrid Recommendation in Conference Contexts

Chun-Hua Tsai
University of Pittsburgh
Pittsburgh, USA
cht77@pitt.edu

Behnam Rahdari
University of Pittsburgh
Pittsburgh, USA
ber58@pitt.edu

Peter Brusilovsky
University of Pittsburgh
Pittsburgh, USA
peterb@pitt.edu

ABSTRACT

A hybrid recommender system fuses multiple data sources to deliver recommendations. One challenge of this approach is to match the changing user preferences with a list of static recommendations. In this paper, we present two user-controllable hybrid recommender interfaces, *Relevance Tuner* (for people recommendation) and *Paper Tuner* (for paper recommendation), which offer a set of sliders to tune the multiple relevance sources on the final recommendation ranking on-the-fly. We deployed the user interfaces to a real-world international academic conference with a field study. The result of the log analysis showed the conference attendees did adopt the interface in exploring the hybrid recommendations. The finding provided evidence in supporting the proposed controllable interface can be deployed to a broader set of conference context.

CCS CONCEPTS

• **Human-centered computing** → **Web-based interaction**; **User interface design**; **Empirical studies in interaction design**.

KEYWORDS

Hybrid Recommendation; User-Driven Fusion; User Interface Design; Conference Contexts

ACM Reference Format:

Chun-Hua Tsai, Behnam Rahdari, and Peter Brusilovsky. 2019. Exploring User-Controlled Hybrid Recommendation in Conference Contexts. In *Joint Proceedings of the ACM IUI 2019 Workshops, Los Angeles, USA, March 20, 2019*, 6 pages.

1 INTRODUCTION

Hybrid recommender systems [5] have been gradually becoming more and more popular due to their ability to combine strong features of different recommender approaches. One promising hybridization design is the paralleled hybrid recommender [5], which fuse recommendation results produced by diverse types of existing recommender algorithms as well as multiple kinds of traces left by modern internet users, i.e., browsing trails, bookmarks, ratings, created social links, etc. In this paper, we will refer to each contributing data source or approach that can generate a list of recommendation ranked by relevance to the target user as a *relevance source*. Each of these sources could be used to build a profile of user interests and deliver valuable recommendations.

Typically, paralleled hybrid recommender fuses multiple relevance sources by assigning static weights to different sources. The optimal weights are trained or learned using ground truth data (i.e., known ratings). The problem with this approach is that users might seek recommendations for different reasons and in different contexts. The individual sources in a hybrid recommender might become more or less valuable depending on each case. As a result, while the “optimal” static fusion could provide the best ranking with high algorithm accuracy, it might be sub-optimal for the users in some specific cases.

The problem of optimal source fusion has been originally explored in the domain of information retrieval where it was demonstrated that the user might be in a better position to decide which weight should be assigned to each relevance source in each case [2]. The idea of user-controlled personalization has been further explored in recommender systems domain by O’Donovan et al., Schafer et al. [13, 15]. More recently, Bostandjiev et al. [4] introduced sliders as an approach to engage the user into tuning various parameters of a recommendation approach. Following that, the use of sliders as a way to support user-controlled fusion has been explored in the domain of recommender systems [14] and information retrieval [7] brings additional evidence in favor of using sliders for user-controlled personalization.

In our past work, we explored sliders as a tool for the user-controlled hybrid recommendation in a research conference context where it was applied to suggest meeting with the most relevant attendees [16]. A controlled user study demonstrated the benefits of this approach. However, it remains unclear whether conference attendees would adopt this approach outside of a controlled study context where the use of sliders was strongly recommended. This paper expands our study of user-controlled hybrid recommendation in the same conference context by adding two new aspects. First, we attempted to extend this approach by applying it to the context of attendee, author, and paper recommendation. Second, instead of performing another controlled study, we assessed the new implementation in an uncontrolled field study by releasing the updated system to attendees of the EC-TEL 2018 conference.

The results of our log analysis showed the conference attendees did adopt the proposed controllable interface in browsing the recommendations. The finding supported the effectiveness of the proposed user interface and its applicability in a broader set of conference context. In the following sections, we review a few like-minded research projects, explain the design of the user-controllable recommender interface and how it can be applied for recommending both papers and people, and review research evidence obtained from this field study.

2 RELATED WORK

Nowadays, it became easier to leverage large amounts of user data to enhance personalization in online applications. A recommender system can create *user models* that utilize users' web browsing trails, item ratings, demographic information or connected social networks for providing personalized recommendations in different contexts. An effective user model can predict the *relevance* of each recommendable item for the user [12]. In search of better performance (i.e., algorithm accuracy), multiple data sources or recommendation techniques could be fused using different hybridization strategies, e.g., paralleled hybrid recommender fuses multiple relevance sources by training a classifier for determining the relevance sources' weighting [5]. The approaches have been widely adopted in many real-world online applications. However, user interests and information needs might not be constant, which makes it is hard to predict user preferences in every situation. That is, it is difficult to find a "one-fit-for-all" weights for a paralleled hybrid recommender in all cases. To overcome this limitation, one promising solution is to offer some form of *user control* so the users can interact with the system based on their current situation.

Bringing user control to a hybrid recommender system allows the users to have an immediate effect on the recommendations [12], i.e., the users can further filter or re-sort the recommendation based on their preference or information need. It usually requires an interactive visualization framework that combines recommendation with visualization techniques to support user interaction or intervention into the recommendation process [11]. The idea of the user-controllable interface of different recommendation approaches was originally presented in [15]. Bostandjiev et al. [4] suggested a slider-based interface that the user can adjust the weights of the items and the social connections. Verbert et al. [19] encouraged users to choose the most appropriate sources of relevance for each case and provided a cluster-map interface to support user-driven exploration and control of tags, agents and users. Parra and Brusilovsky [14] attempted to increase both controllability and transparency of hybrid recommendation by using a combination of sliders for controlling the fusion and a Venn diagram to visualize results. Ekstrand et al. [9] discussed a recommender-switching feature to let the users choose recommender algorithms. Tsai and Brusilovsky [16] offered user controllable interfaces, a two-dimensional scatter-plot and multiple relevance sliders, to a social recommender system for conference attendees. Bailey et al. [3] further provide a visualization for data analytic task using the conference data.

User controllability has also been recognized as a crucial component in supporting the exploratory search, i.e., allowing the users to narrow down the number of items and inspect the details during the information seeking process [6]. Ahn et al. [1] present a summary of search results in the form of entity clouds, which allows the users to explore the results in a controllable interface. Han et al. [10] offered users an option to re-sort people search results based on multiple user-related factors. di Sciascio et al. [7] proposed a *uRank* interface for understanding, refining and reorganizing documents. di Sciascio et al. [6] integrated controllable social search functionality into an exploratory search system. An effective interactive visualization representation can enable users to control the process of recommendation [11].

3 USER CONTROLLED HYBRID RECOMMENDATIONS FOR ACADEMIC CONFERENCES

In this paper, we discuss a visual interface design with user-driven control function and meaningful visual encoding. The design aims to help the users to inspect or tune the ranked recommendations in a hybrid recommender system with multiple relevance sources. Our proposed interface combines several features that have been found beneficial by the past work including slider control of source importance [14, 16] and stackable bars for visualizing combined relevance [7, 16]. The design can be applied to general relevance exploration tasks or recommendation contexts. However, in this paper, we particularly focus on two conferences-focused information needs, i.e., people and paper recommendations.

We implemented the design as two recommender user interfaces: *Relevance Tuner* (for people recommendation) and *Paper Tuner* (for paper recommendation). The two interfaces were served as components of conference support system Conference Navigator 3 (CN3), which equipped with the recommendation functions to the conference attendees. CN3 has been used to support more than 45 conferences at the time of writing this paper and has data on approximately 7,045 articles presented at these conferences; 13,055 authors; 7,407 attendees; 32,461 bookmarks; and 1,565 social connections. The earlier version of the *Relevance Tuner* has been explored in a controlled user study [16], however, in the past we have not explored this approach in different contexts and have not assessed it in a field study.

3.1 People Recommendation: Relevance Tuner

3.1.1 Relevance Sources. To rank the recommended attendees by their relevance to the target user, the system uses five separate recommender engines that rank other attendees along five dimensions. Text similarity of their academic *publications*, *topic* similarity of research interests using topic modeling, social similarity through the *co-authorship* network, similarity of current *interests* measured as intersection of their bookmarked talks, and the *distance* of their place of affiliation to the target user. Each of the relevance is defined below.

- **Publication Similarity** is determined by the degree of publication similarity between two attendees using cosine similarity. The function is defined as:

$$Sim_{Academic}(x, y) = (t_x \cdot t_y) / \|t_x\| \|t_y\| \quad (1)$$

where t is word vectors for user x and y . We used TF*IDF to create the vector with a word frequency upper bound of 0.5 and lower bound of 0.01 to eliminate both commonly and rarely used words.

- **Topic Similarity** is a metric that measures the distance between topic distributions [8]. This is another approach to measure the similarity between the publications of two researchers. The approach assumes that a mixture of topics is used to generate a string (document), where each topic is a distribution of topical words. A recommender engine, based on the topic-based approach, can represent the scholars' research interests through the learned *topics*. The topic

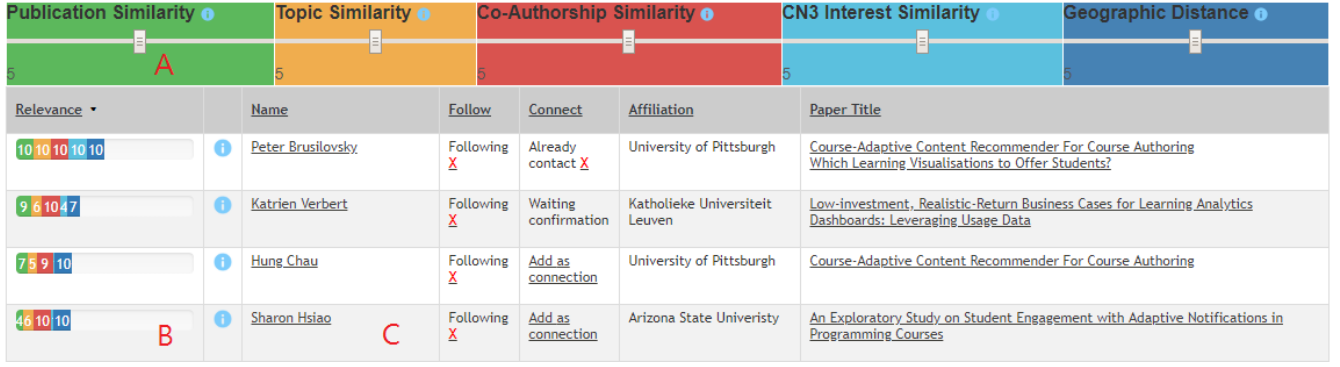


Figure 1: Relevance Tuner: the interface attached five controllable sliders to a people recommender interface, which allows the users to adjust the importance or preference of different relevance aspects.

similarity could be computed as the pairwise similarity of the topic distributions [17].

- **Co-Authorship Similarity** approximates the social similarity between the target and recommended users by combining co-authorship network distance and common neighbor similarity from publication data. We adopted the depth-first search (DFS) method to calculate the shortest path p and common neighborhood (CN) for the number n of coauthors overlapping in two degrees for users x and y .

$$Sim_{Social}(x, y) = p + n \quad (2)$$

- **CN3 Interest Similarity** is determined by the the number of co-bookmarked papers and co-connected authors within the experimental social system. The function is defined as

$$Sim_{Interest}(x, y) = (b_x \cap b_y) + (c_x \cap c_y) \quad (3)$$

where b_x, b_y represent the paper bookmarking of user x and y ; c_x, c_y represents the friend connection of user x and y .

- **Geographic Distance** is a measure of geographic distance between attendees. We retrieve longitude and latitude data based on attendees' affiliation information. We used the Haversine formula to compute the geographic distance between any pair of attendees.

$$Sim_{Distance}(x, y) = Haversine(Geo_x, Geo_y) \quad (4)$$

where Geo are pairs of latitude and longitude coordinates for user x and y .

3.1.2 *Visual Design.* The design of the Relevance Tuner is shown in Figure 1 and firstly introduced by [16]. The design can be summarized in three sections.

- **Section A** contains five controllable sliders with the different colors representing the features of the Personalized Relevance Model. The scale of the slider ranges from 0 to 10. The user can change the weighting on the fly to re-rank the ranked recommendation list. It provides controllability for the user to adjust the ranking to different recommendation needs and preferences.

- **Section B** shows the stackable relevance score bar of each recommended item in the ranked list. The color corresponds to the features in section A. It would adaptively adjust the bar score (length) from 0 to 20, based on the weighting percentage of the sliders. A stackable color bar interface is known for its ability to enhance controllability and transparency in a multi-aspect ranking [7]. In our system, the stackable color bars help the user to see how different relevant aspects of a recommended item are coordinated while adding transparency to the multi-aspect recommendation process.

- **Section C** shows the recommended scholar's meta-data, including name, social connection, affiliation, position, title, and country. The user can sort the ranked list by clicking the head of each column or can inspect the explanation tabs (same as Section C in Figure 1) by clicking the *explanation icon* [18], which is designed to enhance the algorithmic transparency by offering several visualizations regarding the recommendation relevance.

3.2 Paper Recommendation: Paper Tuner

3.2.1 *Relevance Sources.* The current implementation of Paper Tuner uses three personalized and two global social contexts to generate the hybrid recommendation. Each source uses a different type of information to estimate the relevance of each recommended paper to the target user.

- **Publication Similarity** estimates relevance by the degree of text similarity between the user's past publications and the recommended item. We first create a bag of words by concatenating Title, Abstract, and Keywords of each publication and then use TF*IDF to create the word frequency vector. This vector is compared with a similar vector created from user publications using traditional cosine similarity.
- **Bookmark Similarity** is determined by the degree of text similarity between "bookmarked" presentations (i.e., presentation that the user added to her personalized schedule in CN3) and the recommended item. Similarly to *Publication Similarity*, we create a weighted vector of keywords for all

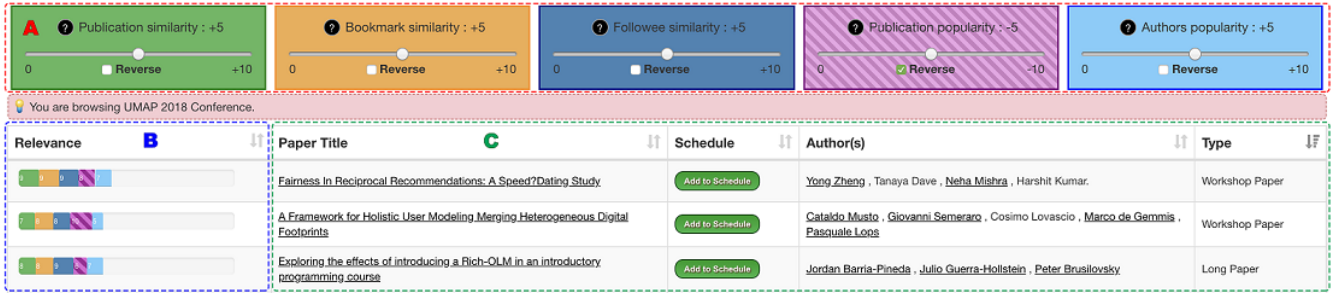


Figure 2: The interface of Paper Tuner with sliders, which allow to adjust the importance of different relevance aspects.

papers in the user’s scheduled papers list and compare it to the vector of the recommended item.

- **Followee Similarity** is based on the ability to follow another user provided by the conference support system as well as by many modern social networks. We create a weighed keyword vector from the entire collection of papers published by the user’s followees and estimate relevance as the cosine similarity between this vector and the vector of the recommended item.
- **Publication Popularity:** Unlike three previous relevance sources, *Publication Popularity* offers not personalized, but social relevance ranking. The *Publication popularity* is determined by the total number of bookmarks received by an item in CN3. We normalized this numerical value and use it to rank items by popularity.
- **Author Popularity:** Similar to *Publication popularity* this is a social relevance source based on the popularity for each author in the system. The popularity of each author is calculated by the average number of bookmarks received by the author’s publications in the system. Once we had this value for each author, we can define the *Author popularity* of each recommended item as the average popularity of its authors.

3.2.2 *Visual Design.* The Paper Tuner is an interface for user controllable recommendation of research papers. It consist of three main parts (Figure 2).

- **Section A** contains five sliders to control the importance of recommendation sources used to generate the ranked list of the results. Users can adjust the weight of each source from 0 to 10 by sliding to the right (increase) or left (decrease). Setting the value of each criterion to 0 will disable the contribution of that source to the final results.
- **Section B** located on the right side of the interface and displays a stacked relevance bar next to each result. The full length of the bar displays the combined relevance of a recommended item to the target user. Each colored segment displays how much a specific source contributed to the total relevance given the current position of the source slider. The segments of the stacked bars update each time the user

changes the sliders, i.e., the length of the “green” section will increase when the green slider is moved right.

- The ranked list of results also provides details for each recommended item (Figure 2: **Section C**). Users can click on the link on *Paper Title* and *Author(s)* columns to get more information such as the abstract of the publication, people planning to attending the presentation, etc.

4 FIELD STUDY

4.1 Context and Data Collection

To explore the value of the two interface designs, Relevance Tuner and Paper Tuner, we organized a field study in the EC-TEL 2018 conference held in Leeds (UK) from September 3 to 6, 2018. The two interfaces were released to all conference users as a part of their host system Conference Navigator 3 (CN3). To mitigate the cold start problem that occurs when users have no publications or co-authorship information related to the event for which the recommendations are produced, the system integrates the AMiner dataset. The live system is available at <http://halley.exp.sis.pitt.edu/cn3/>.

We sent out an invitation email seven days before the conference date to introduce the recommendation feature available in CN3 to all the 158 attendees of the conference. The user IDs were created for each conference attendee based on their registration data. The EC-TEL 2018 conference had accepted a total of 142 papers. We deployed and collected system log data from August 27 to September 14, 2018, which is one week before and after the official conference date. The conference attendees also received several reminder emails during the event date, and the CN3 system link was attached to the homepage of the conference website. The conference website is located at <http://www.ec-tel.eu/index.php?id=805>.

4.2 Log Analysis for People Recommendation

Table 1 presents the system usage for *Relevance Tuner*. A total of 44 users accessed pages with recommended authors or attendees. Around 30% of these users (14 users) interacted with the tuner function. The users in *Tuner Group* (those who click on the sliders at least once) tuned the recommendations 14.28 times on average. The slider *Publication Similarity* and related to it slider *Topic Similarity* received the highest user attention followed by *Co-Authorship Similarity*, however the slider of *CN3 Interest Similarity* was used less frequently. The data indicated that the conference attendees

Table 1: User Interaction Log: Relevance Tuner

Action	All Users		Tuner Group		Non-Tuner Group	
	M (SE)	User Count	M (SE)	User Count	M (SE)	User Count
View Author/Attendee Page	3.61 (3.37)	44/44	3.92 (2.46)	14/14	3.46 (3.74)	30/30
Name Link Clicks	0.95 (1.66)	19/44	1.21 (1.88)	8/14	0.83 (1.57)	11/30
Paper Link Clicks	0.50 (2.00)	8/44	1.00 (3.46)	2/14	0.26 (0.63)	6/30
Total Tuner Clicks	4.54 (11.16)	14/44	14.28 (16.20)	14/14	0 (0)	0/30
Tuner: Publication	1.61 (4.13)	8/44	5.07 (6.14)	8/14	0 (0)	0/30
Tuner: Topic	1.11 (3.03)	7/44	3.50 (4.63)	7/14	0 (0)	0/30
Tuner: Co-Authorship	1.02 (3.92)	5/44	3.21 (6.57)	5/14	0 (0)	0/30
Tuner: CN3 Interest	0.18 (0.62)	4/44	0.57 (1.01)	4/14	0 (0)	0/30
Tuner: Geographic	0.63 (2.70)	4/44	2.00 (4.60)	4/14	0 (0)	0/30
Total Explanation Clicks	1.54 (4.11)	8/44	1.00 (2.54)	3/14	1.80 (4.69)	5/30
Exp: Publication	0.34 (0.80)	8/44	0.28 (0.72)	2/14	0.36 (0.85)	6/30
Exp: Topic	0.43 (1.12)	7/44	0.42 (1.15)	2/14	0.43 (1.13)	5/30
Exp: Co-Authorship	0.50 (1.64)	7/44	0.21 (0.57)	2/14	0.63 (1.95)	5/30
Exp: CN3 Interest	0.11 (0.38)	4/44	0.07 (0.26)	1/14	0.13 (0.43)	3/30
Exp: Geographic	0.15 (0.56)	4/44	0 (0)	0/14	0.23 (0.67)	4/30

Table 2: User Interaction Log: Paper Tuner

Action	Total Adjustments	Median	Standard Deviation	User Count
Total Tuner Adjustment	112	9.33	5.35	12/12
Tuner : Publication Similarity	32	2.67	1.44	11/12
Tuner : Bookmark Similarity	22	1.83	1.85	8/12
Tuner : Followee Similarity	21	1.75	1.60	8/12
Tuner : Publication Popularity	19	1.58	1.78	7/12
Tuner : Author Popularity	18	1.50	1.09	10/12
Reverse Functionality	15	1.25	1.14	8/12
Paper Link Clicks	28	2.33	0.98	12/12
Author Link Click	19	1.58	1.16	9/12

emphasize the publication text in exploring the conference authors and attendees. Note, however, that for those who are new to the system, the *CN3 Interest* may be less useful due to lack of bookmarking data. It might explain the lower use of the slider.

4.3 Log Analysis for Paper Recommendation

Table 2 provides information about authors and attendees interactions with the paper tuner during the EC-TEL 2018 conference. The analysis revealed that all system users who explored the Paper Tuner component used the sliders to control the weights of relevance sources. As in the case of Relevance Tuner, the publication similarity was the most popular slider. For the case of the papers, however, the bookmark similarity was the second most popular one. Social relevance sources were adjusted less frequently than personalized sources. Altogether, it looks like the Paper Tuner provided valuable information to the users as they frequently requested additional details about recommended papers including 28 clicks to receive paper details and 19 clicks to receive author details.

5 SUMMARY AND CONCLUSION

In our past work, we explored sliders as a tool for the user-controlled hybrid recommendation of conference attendees in a controlled user study [16]. In this paper, we attempted to implement the idea of user-controlled hybrid recommendation in a broader set of conference context and explore it in a field study. The result of log analysis indicated that the users to a considerable extent adopted the user-controlled interface in exploring the hybrid people and paper recommendations of a conference. This result provides some early evidence about the effectiveness of the proposed user interface in a real-world, outside of controlled student context setting. Based on the preliminary finding, we see a great potential to deploy the controllable interface to other relevance exploration tasks or recommendation contexts.

We also are aware of some limitations in this experiment. First, the multiple relevances were combined with linear fashion. Second, we reported the observational findings due to the nature of the field. The experimental robustness requires further investigating. Third, in each interface, we had engaged only a small sample of users. All these limitations would be addressed in our future works.

REFERENCES

- [1] Jae-wook Ahn, Peter Brusilovsky, Jonathan Grady, Daqing He, and Radu Florian. 2010. Semantic annotation based exploratory search for information analysts. *Information processing & management* 46, 4 (2010), 383–402.
- [2] Jae-wook Ahn, Peter Brusilovsky, Daqing He, Jonathan Grady, and Qi Li. 2008. Personalized Web Exploration with Task Models. In *the 17th international conference on World Wide Web, WWW '08*. ACM, 1–10.
- [3] Samuel M Bailey, Justin A Wei, Chaoli Wang, Denis Parra, and Peter Brusilovsky. 2018. CNVis: A Web-Based Visual Analytics Tool for Exploring Conference Navigator Data. *Electronic Imaging* 2018, 1 (2018), 1–11.
- [4] Svetlin Bostandjiev, John O'Donovan, and Tobias Höllerer. 2012. TasteWeights: a visual interactive hybrid recommender system. In *Proceedings of the sixth ACM conference on Recommender systems*. ACM, 35–42.
- [5] Robin Burke. 2007. Hybrid Web recommender systems. In *The Adaptive Web: Methods and Strategies of Web Personalization*, Peter Brusilovsky, Alfred Kobsa, and Wolfgang Neidl (Eds.). Lecture Notes in Computer Science, Vol. 4321. Springer-Verlag, Berlin Heidelberg New York, 377–408.
- [6] Cecilia di Sciascio, Peter Brusilovsky, and Eduardo Veas. 2018. A study on user-controllable social exploratory search. In *23rd International Conference on Intelligent User Interfaces*. ACM, 353–364.
- [7] Cecilia di Sciascio, Vedran Sabol, and Eduardo E. Veas. 2016. Rank As You Go: User-Driven Exploration of Search Results. In *the 21st International Conference on Intelligent User Interfaces (IUI'16)*. ACM, 118–129. <https://doi.org/10.1145/2856767.2856797>
- [8] Jianguang Du, Jing Jiang, Dandan Song, and Lejian Liao. 2015. Topic modeling with document relative similarities. *IJCAI*.
- [9] 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*. ACM, 11–18.
- [10] Shuguang Han, Daqing He, Jiepu Jiang, and Zhen Yue. 2013. Supporting exploratory people search: a study of factor transparency and user control. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*. ACM, 449–458.
- [11] Chen He, Denis Parra, and Katrien Verbert. 2016. Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. *Expert Systems with Applications* 56 (2016), 9–27.
- [12] Dietmar Jannach, Sidra Naveed, and Michael Jugovac. 2016. User control in recommender systems: Overview and interaction challenges. In *International Conference on Electronic Commerce and Web Technologies*. Springer, 21–33.
- [13] John O'Donovan, Barry Smyth, Brynjar Gretarsson, Svetlin Bostandjiev, and Tobias Höllerer. 2008. PeerChooser: Visual Interactive Recommendation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 1085–1088. <https://doi.org/10.1145/1357054.1357222>
- [14] Denis Parra and Peter Brusilovsky. 2015. User-controllable personalization: A case study with SetFusion. *International Journal of Human-Computer Studies* 78 (2015), 43–67.
- [15] 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 (CIKM '02)*. ACM, New York, NY, USA, 43–51. <https://doi.org/10.1145/584792.584803>
- [16] Chun-Hua Tsai and Peter Brusilovsky. 2018. Beyond the Ranked List: User-Driven Exploration and Diversification of Social Recommendation. In *23rd International Conference on Intelligent User Interfaces (IUI '18)*. ACM, New York, NY, USA, 239–250. <https://doi.org/10.1145/3172944.3172959>
- [17] Chun-Hua Tsai and Peter Brusilovsky. 2019. Explaining Recommendations in an Interactive Hybrid Social Recommender. In *Proceedings of the 2019 Conference on Intelligent User Interface*. ACM, 1–12.
- [18] Chun-Hua Tsai and Peter Brusilovsky. 2019. Exploring Social Recommendations with Visual Diversity-Promoting Interfaces. *ACM Transactions on Interactive Intelligent Systems (TiiS)* (2019).
- [19] Katrien Verbert, Denis Parra, Peter Brusilovsky, and Erik Duval. 2013. Visualizing recommendations to support exploration, transparency and controllability. In *Proceedings of the 2013 international conference on Intelligent user interfaces*. ACM, 351–362.