User-Controlled Hybrid Recommendation for Academic Papers

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

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

ABSTRACT

This demo paper presents Paper Tuner, a user-controlled interface for recommending papers and presentations in a research conference context. The availability of multiple sources of information about user interests make hybrid recommendation approach attractive in a conference context, but traditional static parallel hybridization makes is hard to generate a single ranking that can address different needs. We introduce a novel slider-based user interface that allows users to control the importance of different relevance source and even reverse the impact of specific sources. The log analysis of system usage during in a real conference context revealed an extensive use of sliders. Moreover, nearly half of the users applied the reverse functionality while using the sliders.

ACM Classification Keywords

H.1.2 User/Machine Systems: Human factors; H.3.3 Information Search and Retrieval: Information filtering; H.5.2 User Interfaces: User-centered design

Author Keywords

User-Controlled Recommendation; Hybrid Recommendation

INTRODUCTION

Modern recommender systems routinely fuse together multiple sources of information (browsing trails, ratings, social links, etc.) to increase the accuracy of generated recommendation. Parallel hybrid recommendation approaches perform this fusion by assigning optimal static weights to the contributing sources. However, this static approach might not perform well when users refer to a recommender to address different information needs or in different contexts. To address this problem, the field of intelligent user interfaces introduce user-controller personalization where different subsets of sources or different weights could be engaged for different needs.

In our recent work, we explored the use of sliders to support user-controlled hybrid recommendation in a research conference context where it was applied to suggest most relevant attendees meet [2]. While slider-controlled fusion performed well, we discovered an important weakness of this approach: it was hard to use when the recommendation context required *reversing* the weight of a specific recommendation source. For example, while a recommendation source based on co-authorship links ranks attendees by its social similarity with the target user, the recommendation case might require to find attendees who are interested in similar topics while being most likely unknown to the target user (i.e., having the *weakest* social similarity). In this demo paper we present *Paper Tuner* system, which explores two extensions of our earlier work: (1) using slider-based controlled fusion to recommend research papers rather than attendees and (2) providing the option to reverse relevance of each contributing source.

IUI'19, March 17–20, 2019, Los Angeles , CA, USA

ISBN X-XXXXX-XX-X/XX/XX.

DOI: http://dx.doi.org/10.1145/0000000.0000000

THE PAPER TUNER

The Paper Tuner [1] is a user-controllable hybrid recommender system for papers and presentations at an academic conference. The system combines several features that has been found beneficial by the past work including slider control of source importance and *stackable bars* for visualizing relevance.

Design

The Paper Tuner consist of three main parts (Figure 1). **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) and left (decrease). Setting a value of each criteria to 0 will disable the contribution of that source to the final results. **Section B** located in 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 1: **Section C**). Users can click on the link on *Paper Title* and *Author(s)* columns to get more information.

Recommendation Sources

The current implementation of Paper Tuner uses three personalized and two community-level recommendation sources

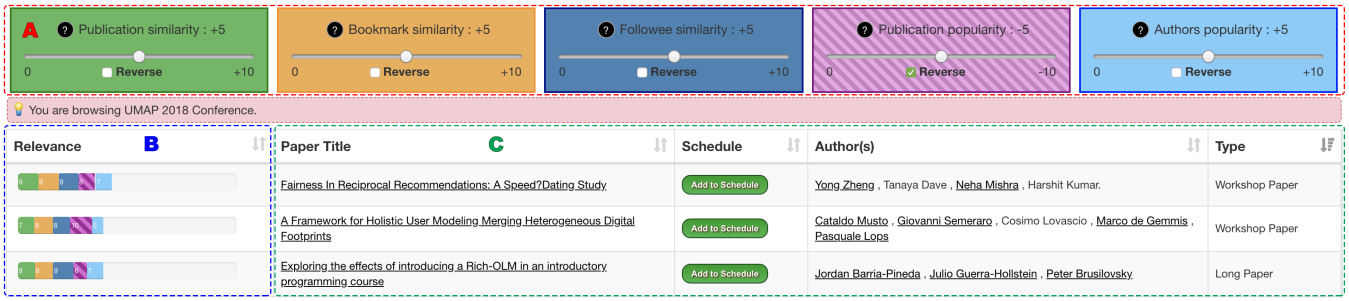

Figure 1. Interface Design: *Paper Tuner* main page interface. when the reverse function is activated, both "slider" and "bar" design will change.

to generate hybrid recommendation. Each source uses different type of information to estimate relevance of each recommended paper to the target user.

Publication Similarity estimates item relevance as the degree of text similarity between the user's past publications and the recommended paper. We create a bag of words by concatenating Title, Abstract and Keywords of each publication and 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 degree of text similarity between papers bookmarked by the target user papers in the system and the recommended item. Similarly to *Publication Similarity*, we create a weighted vector of keywords for all papers in user's bookmark list and compare it to the vector of the recommended item.

Followee Similarity is based on information about following other users. As many modern social systems, *Conference Navigator* users are able to follow fellow users when they are interested in their research area. We estimate item relevance a the of cosine similarity between publications of the user's followees and the recommended item.

Publication popularity is determined by the total number of bookmarks received by an item in the *Conference Navigator* system. We normalized this number and use it to rank papers by popularity.

Author popularity of each recommended item as the average popularity of its co-authors. In turn, popularity of each author calculated by the average number of bookmarks received by the author's publications in the system.

Reversible Relevance

To choose an an approach for reversing source similarity, we compared two different designs for controlling (Figure 2) and reviewing (Figure 3) reversed relevance. The first slider design (Figure2, left) offered a checkbox to reverse the slider scale of from 0:10 to 0:-10. The second design (Figure2, right) allowed a smooth transition from negative to positive relevance using the scale from -5 to 5. The stacked bars in the first design (figure3, left) attempt to display the contribution of negatively weighed relevance sources as positive *dissimilarity*. Dissimilarity is defined by the distance between the current value of relevance and the maximum available value for that criteria. For example, if a recommended item has 65% similarity to the user as measure by *Publication Similarity* source, then the reverse value (dissimilarity) for this source is 35%. In the stacked bar, dissimilarity is shown using mesh pattern. This approach ensures "natural" behavior of reversed relevance sources: as the reversed slider moves to the right towards -10, dissimilarity increases. The second version (figure3 - right) presents the total relevance as a balance of negative relevance stacked on left (red) side and positive similarity stacked on the right (green) side, according to positive and negative value of the sliders.

To select a design that is most clear to the target users, we conducted a user study in which participants were asked to chose from the proposed designs.66.7% of participants favoured the first design for both sliders (Figure2, left) and stacked bars (Figure3, left). We used these preferred design choices in the final version of the Paper Tuner.

Figure 2. Reversible slider designs

Figure 3. Stackable relevance bar designs

The Field Study

We evaluated the *Paper Tuner* in a field study by making it available to authors and attendees during the EC-TEL 2018 conference. The analysis of log data revealed that all system users who explored the Paper Tuner component used the sliders to control the weights of relevance sources and 47% of users (15, M=1.25, SD = 1.14) applied the reverse functionality while using the sliders.

REFERENCES

1. `https://goo.gl/MXyqyM`
2. 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, Article 0, 12 pages. DOI:`http://dx.doi.org/10.1145/3172944.3172959`