User Model in a Box: Cross-System User Model Transfer for Resolving Cold Start Problems

Chirayu Wongchokprasitti[1](✉), Jaakko Peltonen[2,3], Tuukka Ruotsalo[2], Payel Bandyopadhyay[4], Giulio Jacucci[4], and Peter Brusilovsky[1]

[1] School of Information Sciences, University of Pittsburgh, Pittsburgh, USA
{chw20,peterb}@pitt.edu
[2] Helsinki Institute for Information Technology HIIT, Aalto University, Espoo, Finland
{jaakko.peltonen,tuukka.ruotsalo}@aalto.fi
[3] School of Information Sciences, University of Tampere, Tampere, Finland
[4] Helsinki Institute for Information Technology HIIT, Department of Computer Science, University of Helsinki, Helsinki, Finland
{pbandyop,giulio.jacucci}@cs.helsinki.fi

Abstract. Recommender systems face difficulty in cold-start scenarios where a new user has provided only few ratings. Improving cold-start performance is of great interest. At the same time, the growing number of adaptive systems makes it ever more likely that a new user in one system has already been a user in another system in related domains. To what extent can a user model built by one adaptive system help address a cold start problem in another system? We compare methods of cross-system user model transfer across two large real-life systems: we transfer user models built for information seeking of scientific articles in the SciNet exploratory search system, operating over tens of millions of articles, to perform cold-start recommendation of scientific talks in the CoMeT talk management system, operating over hundreds of talks. Our user study focuses on transfer of novel explicit *open user models* curated by the user during information seeking. Results show strong improvement in cold-start talk recommendation by transferring open user models, and also reveal why explicit open models work better in cross-domain context than traditional hidden implicit models.

Keywords: Cross-system user modeling · Recommender systems

1 Introduction

Recommender systems often face the *cold-start problem*, where recommendations are required for users for whom not enough preference information is available, because the users have rated few or no items in the system [9]. In this article, we investigate the idea of *user model transfer* to enable warm start in cross-system recommendation scenario. That is, we consider user models that can be established in a *source system* and then used in another *target system*. While this idea

F. Ricci et al. (Eds.): UMAP 2015, LNCS 9146, pp. 289–301, 2015.
DOI: 10.1007/978-3-319-20267-9_24

is not new, past research on user model transfer produced mixed results. Our paper expands earlier research by exploring transferability of *open user models*. We investigate a scenario where user or the source system have the ability to explore and curate their model by visual interaction. We believe that open user modeling approach produces better quality user models that could be especially valuable for cross-system transfer. To assess this hypothesis, we investigate how the cross-system transfer of open user models improves performance of recommender methods both in the extreme cold-start setting when no preferences are available from the user, and over time when some preferences from the user become available. We compare different transfer strategies in an academic information setting where the source system is a search system for scientific papers and the target system is a system for sharing academic talks. The open user models in the source system are built from visual interaction with keywords and other available information includes views and bookmarks of query results; models in the target system are built from ratings of talks.

We believe that our work is the first one to explore the transferability of open user models. Its main contributions are: 1) we show cross-system transfer of open user models greatly improves cold-start recommendation performance, 2) we investigate different ways of transferring open user models from an information seeking system to a talk recommendation system, as well as transfer of more traditional implicit and explicit document information, and show the open user models bring the greatest benefit, for which we provide an explanation by analysis of cross-system similarities of the different information types.

2 Related Work

In recent years, cross-system and cross-domain recommendation research is grown in popularity due to the growing number of personalized system that collect information about the users. It this context, it becomes natural to export information collected about a user in one system and transfer it to another system (maybe in a different domain) to improve the quality of recommendations. It has been argued that this transfer might be most valuable in cold-start context when a recommender system has insufficient information about new users in a target system [8]. Despite the overall interest in this field, there are still very few studies exploring real cross-system user information transfer due to the lack of sizable datasets that have pairs of users in two systems. As a result, a major share of research on cross-domain recommendation focused on transfer learning approaches that do not assume a common user set for different domains [4,6].

According to the most recent review [3], the work on cross-domain recommendation could be split into two large groups - those using collaborative filtering and content-based approaches. The classic examples or collaborative filtering transfer are [1,2]. These papers offered an extensive discussion of cross-domain recommendation problems and suggested interesting generic approaches but were not able to explore these approaches in a true cross-domain context using instead artificial datasets produced by separation of single-domain user movie ratings into subdomains. More recent work explored collaborative transfer approaches in more realistic settings with hundreds of users having ratings in both domains [8]. Content-based cross-domain recommendation appeared to be a harder challenge. No immediate success was reported for simple keyword-based profile transfer approaches. As a result, the majority of research and success in this category focused on using shared semantic-level features such as social tags [11] or Wikipedia [5]. Unfortunately, it leaves open the case where user preference data in the source domain includes no tags and can't be associated with an extensive ontology such as Wikipedia.

In this context, our work attempts to re-examine the prospects of the keyword-level user model transfer across related, but different domains. To fight the known problems of keyword profile transfer, we explored a different kind of profile. While more traditional implicit keyword-level user model is served as a baseline in our studies, main emphasis is on transfer of an explicit model of interest that is open to the users in the source systems and explicitly curated by them.

3 Setting: Two Academic Information Systems

We study cross-system transfer of user models between two recent academic information systems, CoMeT and SciNet.

CoMeT is a system for sharing information about research talks at Carnegie Mellon University and University of Pittsburgh. The system is available online at http://halley.exp.sis.pitt.edu/comet/; it is a collaborative tagging system, which allows any individual to announce, find, bookmark, and tag talks. To help users locate interesting talks, CoMeT includes a content-based recommender system which builds an interest profile of individual users and recommend new talks to users immediately after such talks are posted.

SciNet [7] is an exploratory search system. SciNet indexes over 50 million scientific documents from Thomson Reuters, ACM, IEEE, and Springer. Going beyond text-based queries, SciNet helps users direct exploratory search by allowing them to interact with an *open user model* discussed below. The approach (called "interactive intent modeling" in [7]) significantly improved users? information seeking task performance and quality of retrieved information [7], thus the open user models are promising for cross-system transfer.

Open User Models in SciNet. Unlike traditional information seeking systems, SciNet opens its user model by allowing users to directly see a visual representation of the model and interact with it. User intent is modeled as a vector of interest values over a set of available keywords. The vector can be seen as a bag-of-keywords representation of an "ideal" document containing keywords in proportion to their interest values. Figure 1 shows the interface. The open user model is visualized as a radial layout where the estimated search intent and alternative intents are represented by scientific keywords, organized so that keywords relevant to the user are close to the center and similar intents have similar angles. Users start by typing a query and receive a list of documents which they can bookmark, and then direct the search by interacting with the user model.

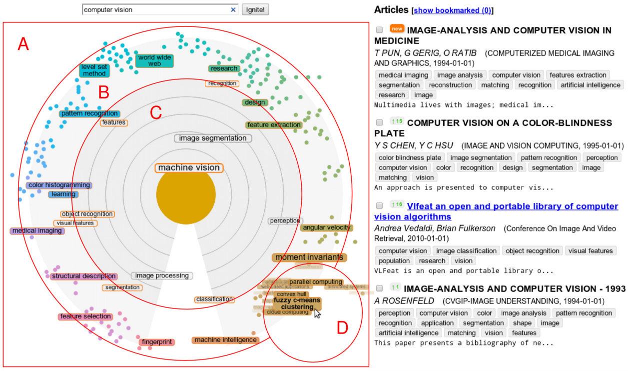

Fig. 1. The SciNet system. The radial display (A) represents the open user model by showing keywords for interaction; inner keywords (C) represent estimated intent and outer keywords (B) represent alternative intents. The user can inspect keywords with a fisheye lens (D) and bookmark documents.

Users can inspect model keywords shown on the radar and *curate* the model by dragging keywords to change their importance. After each iteration, the user model is inferred from the whole set of user actions, documents are searched based on the updated model, and the radial visualization is updated.

Our interest is to use (1) the whole content of the open user model and (2) its curated subset (the keywords the user moved in the process of curation). As a baseline, we also explore transfer of more traditional information: (3) the set of relevant documents selected by the user in the process of search (which could be considered as hidden, implicit user model) and (4) a broader set of all documents retrieved in response to user queries that is a weaker reflection of user interests.

4 Model Transfer for Cross-System Recommendation

The goal of our cross-system setup is to recommend users relevant talks based on features extracted from the description of the talk and their model of interests . Each CoMeT talk is represented by a unigram model, that is, as a vector of word counts in the description of the talk, normalized by the maximum word count, and converted into term frequency–inverse document frequency (TF-IDF) representation. We create a warm start for CoMeT talk recommendation by transferring user models from SciNet. We use two approaches to transfer the explicit *open user model* from SciNet, and two approaches to transfer implicit user models from the SciNet search trace, detailed below.

A. Ranking based on the CoMeT user model alone (denoted 'Baseline'). The baseline method ignores SciNet, and ranking is based only on the CoMeT user model. In the pure cold-start case where no bookmarks are available, the baseline is unable to give recommendations.

B. Transfer the explicit open user model from manipulated keywords (denoted 'ma.keywords'). The SciNet open user model represents the user's interest over keywords. While the model predicts importance over all keywords, a subset of most promising ones are shown for interaction, and a further subset out of those are manipulated by the user. Here we take this last subset: the set of all scientific keywords dragged by the user on the SciNet interface during the search session is treated as a pseudodocument containing the keywords. Each keyword is associated with a weight corresponding to the user's interest (radius where the user dragged the keyword). The pseudodocument is converted into a vector of unigrams by taking each unigram within each keyword (e.g. "support" and "vector" within "support vector"), associating it with the corresponding weight of the keyword (or sum of weights if the unigram occurs in several keywords), and discarding unigrams that do not appear in the CoMeT corpus. This extracts from the SciNet open user model the unigram information common with the CoMeT information space. Since the open user model is represented as a single pseudodocument instead of a corpus, we do not convert the corresponding unigram vector into TF-IDF representation, instead we only normalize it by its maximum value; the corpus of CoMeT talks is then converted into TF-IDF over CoMeT talks only. The resulting unigram vector of the pseudodocument is added into the user's cold-start set of bookmarked CoMeT talks.

C. Transfer the explicit open user model from shown keywords (denoted 'sh.keywords'). The SciNet open user model displays the subset of most important keywords to the user at each search iteration (area A in Figure 1 right). We take the subset in each iteration as a pseudodocument, where each keyword is associated with a weight corresponding to the user interest predicted by SciNet. We convert each such pseudodocument into a vector of unigrams in the same way as in **B**; the user thus gets one vector of unigrams for each search iteration, which together represent the evolution of the user model over the search session. As in **B** this is not a corpus, hence we only normalize each vector of unigrams by its maximum value and CoMeT talks apply TF-IDF weighting over CoMeT talks only. The unigram vectors of the pseudodocument are added into the user's cold-start set of bookmarked CoMeT talks.

D. Transfer an implicit user model from bookmarked documents (denoted 'bm.papers'). Scientific documents bookmarked by the user during the SciNet search session provide implicit information about the user's interests. We convert all bookmarked documents the same unigram representation as CoMeT talks as in **B**. Since the bookmarked documents form a corpus, we convert unigram vectors of CoMeT talks and SciNet documents into a TF-IDF representation computed over both corpuses. For each user we add the resulting unigram vectors of the SciNet bookmarks into the cold-start set of that user's bookmarked CoMeT talks.

E. Transfer an implicit user model from shown documents (denoted 'sh.papers'). We do the same as in **D** but using all documents seen by the user during the SciNet search session (not only bookmarked documents). Note that these documents had been retrieved for the momentary user models during the session; **C** represents the momentary models, whereas **E** represents the momentary best matches to the models in the corpus. Documents as in **E** are available in all interactive search systems whereas keywords in **C** require an open model.

5 User Study for Data Collection

To perform experiments on cross-system model transfer, we had to collect data that capture interests of the same users in the two explored systems. This data was collected through a user study In Section 6 we will perform experiments on cross-system transfer from SciNet to CoMeT, for cold-start prediction of CoMeT talk attendance using information from the SciNet search trace. Before we perform the recommendation experiments, we conduct a task-based user study in a laboratory setting, to collect data for a set of users over both systems:

1. We collected the search trace and open models of the users when they conducted an exploratory search in SciNet for scientific literature corresponding to their research interests. Participants were asked to imagine that they are preparing for a course or seminar on their research interest.
2. We collected user preference in attending academic talks indexed in the CoMeT system. Participants were asked to bookmark interestin talks and rate to what extent they would like to attend it.

Task Descriptions. For SciNet, we chose a search task complex enough that users must interact with the system to gain the information needed to accomplish the task, and broad enough to reveal research interests of users. The task is: "Write down three areas of your scientific research interests. Imagine that you are preparing for a course or a seminar for each research interest. Search scientific documents that you find useful for preparing for the courses or seminars." To determine which documents were relevant to the task, we asked the users to bookmark at least five documents for each research interest. For CoMeT, we used the following rating task: "Please rate all the scientific talks whether you would like to attend the talks or not. If you don't want to attend the talk then just click "no" button and go to next talk. If you want to attend the talk then you click "yes" button and fill the ratings." We provided the following guidelines for ratings:
"5" : This talk matches my research interest and I would definitely attend it.
"4" : This talk matches my research interest and I would likely attend it.
"3" : This talk somewhat matches my research interest and I might attend it.
"2" : This talk somewhat matches my interest, but its unlikely that I attend it.
"1" : This talk somewhat matches the research interest but I wouldn't attend it.

Participants. We recruited 20 researchers (14 male and 6 female) from University of Helsinki to participate in the study. All participants were research staff (10 PhD researchers and 10 research assistants) in computer science or related fields. The participation was limited to researchers because the nature of SciNet and CoMeT required participants having experience in scientific document search and having interest in attending research related talks or seminars.

Prior to the experiment, we conducted a background survey of the participants to ensure that they have conducted literature search before and have also attended research related talks or seminars.

Procedure. The study used a within-subject design in which all the participants performed both the tasks using both the systems alternatively. To minimize the impact of experience in one system on another, we counter balanced the order of system use. Ten participants used the SciNet system first and then used the Comet system while the remaining ten participants used the CoMeT system first and then used the SciNet system. The protocol for each system had two stages: for SciNet the stages were demonstrating the system (7 minutes) and then performing the search task by the participant (30 minutes); for CoMeT the stages were demonstrating the system (7 minutes) and then performing the rating task by the participant (75 minutes).

Data Logging. When the participants were performing the search task in SciNet, we logged all their interactions with the system Data logged from each interaction included details of the scientific documents displayed, keywords for the estimated intent predicted by the system and for alternative intents (areas C and B in the right-hand subfigure of Figure 1), curated keywords, search query of the users, abstracts of scientific documents viewed, scientific documents bookmarked by the users and the corresponding timestamps of all the user interactions. When the participants were performing the rating task in CoMeT, we logged the ratings of the scientific talks or seminars and the ratings of the novelty of the scientific talks or seminars.

6 Cross-System Recommendation Experiment

The user data gathered in Section 5 was used to create experiments on cross-system recommendation of CoMeT talks by transporting user models from Scinet to CoMeT. We compare the four modes of transporting user models described in Section 4. We evaluate both the global impact in a setting with several rated talks available from each user, and performance in the harder cold-start setting.

Data Processing and Demographics. There were 500 CoMeT talks selected from January 10 to February 5, 2013, containing 8,406 unique unigram terms. SciNet indexes over 50 million scientific articles, out of which users see a subset based on their interest and search behavior: there were 9,457 unique SciNet articles returned to our participants with 30,848 unique terms. SciNet also records keywords shown to the user and manipulated by the user; in total the participants were shown 3,474 unique keywords and manipulated 178 unique keywords. Scientific documents and keywords from Scinet and research talks (talk descriptions) from CoMeT were cleaned from html tags, stop words were removed, and words were stemmed by the Krovetz algorithm. In each cross-system transfer approach, the documents, talks, and keywords were converted into unigram vectors with according to term frequency–inverse document frequency (TF-IDF) schemes, as described in Section 4.

6.1 Experiment 1: Global Impact of Cross-System Models

We first consider the traditional (non-cold-start) learning setting where much training data is available within CoMeT, that is, recommendation for users that have used both CoMeT and SciNet for some time. We call this evaluation of global impact of the transfer. We compare the CoMeT-only baseline and four transfer approaches: explicit open user model from manipulated or shown keywords, and implicit user model from bookmarked or shown documents. For each approach we use three methods to make recommendations based on the transported user model: Centroid, k-Nearest Neighbors, and positive-sample-only k-Nearest Neighbors. In all three approaches, We assign equal weight on any vector from either CoMeT or SciNet source.

- **Centroid**: We take the centroid (mean) of the unigram vectors of the user's bookmarked CoMeT talks and any vectors transferred from SciNet. The unigram vectors are extracted from the content of CoMeT talks and SciNet papers (title and abstract) users have bookmarked. Test talks are ranked by cosine similarity of their unigram vector to the centroid.
- **k-Nearest Neighbors (k-NN)**: We treat unigram vectors of bookmarked CoMeT talks and vectors transferred from SciNet as positive samples, and vectors of non-bookmarked CoMeT talks as negative samples. For each test talk we find its k neighbors (nearest positive or negative samples) by cosine similarity of unigram vectors. The k neighbors can be all positive samples or negative ones or combination of them. Test talks are then ranked by $s_{pos} - s_{neg}$, where s_{pos} is the sum of cosine similarities from the test talk to the positive neighbors and s_{neg} is the sum of cosine similarities to the negative neighbors.
- **Positive-Sample-Only k-Nearest Neighbors (denoted k-NN.PO)**: This method is similar to the k-NN approach but we find the k neighbors of a test talk from the positive samples only, and rank the test talk by the sum of cosine similarities from the test talk to the k positive neighbors [10].

We use a ten-fold cross-validation setup: bookmarked and non-bookmarked CoMeT talks of each user were randomly divided into ten equal-size bins, and evaluation was run ten times with each bin in turn held as the test set and the other nine bins as the learning set used to construct user models. SciNet user model data was transferred to each learning set by the transfer approaches to augment the resulting user model. Results (rankings of test talks) are evaluated by mean average precision (MAP): mean of precision values at locations of positive test talks in the ranking, averaged over users and cross-validation folds.

Results of Experiment 1. Table 1 shows the results of global transfer models from Scinet data to CoMeT system. Performance of the CoMeT-only baseline varies by recommendation approach: Centroid and k-NN perform comparably and k-NN.PO is slightly better. For the transfer approaches centroid performs comparably to other approaches; k-NN or k-NN.PO can yield slightly higher MAP but depend on value of k. Overall, in this non-cold-start setting transfer

Table 1. The Mean Average Precision (MAP) of Global Impact from Transferring Models by Transporting Data Methods and by Recommending Algorithms

Mean Average Precision		**Centroid**	**k-NN**				**k-NN.PO**			
			5nn	**10nn**	**20nn**	**30nn**	**5nn.po**	**10nn.po**	**20nn.po**	**30nn.po**
baseline		0.47	0.45	0.47	0.48	0.46	0.48	0.50	0.50	0.50
Implicit	**ex.papers**	0.42	0.44	0.45	0.45	0.44	0.43	0.44	0.44	0.44
User Model	**im.papers**	0.36	0.36	0.36	0.35	0.35	0.36	0.37	0.36	0.36
Explicit Open	**ex.keywords**	0.48	0.46	0.48	0.48	0.47	0.49	0.51	0.51	0.50
User Model	**im.keywords**	0.47	0.46	0.48	0.49	0.49	0.48	0.49	0.49	0.48

of SciNet information did not yield much improvement over the CoMeT-only baseline: transferring an implicit model from shown SciNet documents underperformed the CoMeT-only baseline and the other transfer approaches were not significantly different from that baseline. In this non-cold-start situation user profiles in CoMeT had enough data to work well on their own; thus we expect greatest benefit of transfer in the cold-start setting.

6.2 Experiment 2: Cold-Start Recommendation

Here we study the main focus of the paper: cold-start settings. We again transfer four kinds of user models from Scinet (explicit open user models from shown and manipulated keywords, and implicit models from shown and bookmarked documents).

Experimental Setup: Ten-round-ten-fold Cross-Validation. We consider a range of cold-start setups where the user has bookmarked 0-20 CoMeT talks. For each number of bookmarked talks we use a ten-fold cross-validation setup: we divide data into ten bins; in each fold we hold out one cross-validation bin as test data, and from the remaining nine bins we randomly sample a pool of bookmarked talks and a pool of non-bookmarked talks; within each fold we perform that random sampling 10 times (10 "rounds") and report average results over rounds. For each number of bookmarked talks, the number of sampled non-bookmarked talks was chosen to keep the same ratio of bookmarked to non-bookmarked talks as overall in the data of that user. We evaluate results by the same MAP criterion as in the first experiment.

Results of Cold-Start Recommendation. We first tested cold-start recommendation for the CoMeT-only baseline using the three recommendation methods (centroid, k-NN, k-NN.PO as in Table 1). centroid and k-NN.PO models performed equally (results omitted for brevity; essentially no visible difference regardless of the profile size) and outperformed the k-NN models. Given that the centroid model is simpler and faster than k-NN.PO, it was used for all transfer methods in the cold-start recommendation. Figure 2 shows the results of transfer in the cold-start setting. Explicit transfer of the SciNet open user models (ma.keywords and sh.keywords) helped the transition of the new users into the CoMeT system. The MAP results of keyword models are significantly better than the baseline until the user has two bookmarked talks in the Scinet implicit

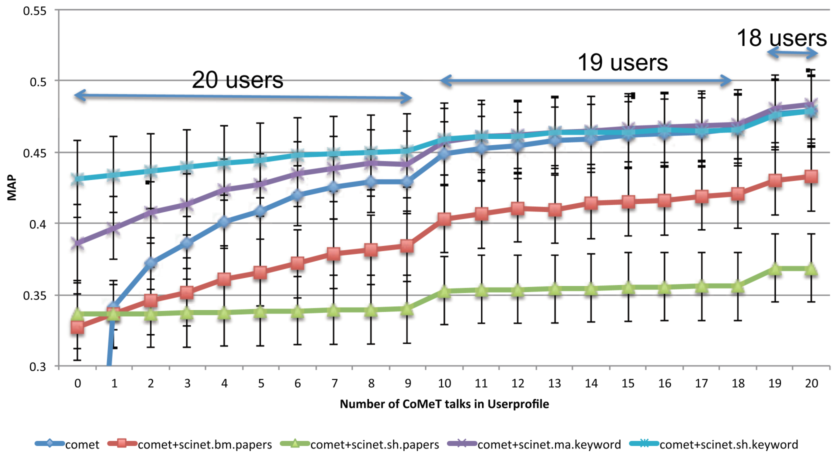

Fig. 2. Cold-start-effect MAP Results of Centroid Models. The MAP is shown for each method and each number of bookmarked CoMeT talks. Error bars show 95% confidence interval of the mean over users and cross-validation folds. "20 users"–"18 users" denote how many users had enough data within cross-validation folds to have the desired number of bookmarked talks.

keyword transfer centroid model, and until five bookmarked talks in the Scinet explicit keyword transfer centroid model. While the transfer of explicit curated models worked well, the transfer of implicit models built from retrieved and bookmarked papers harmed the MAP performance. Even the model built from explicitly selected talks performed poorly comparing to the baseline. We analyze reasons for this performance difference in the next section.

7 Analysis of Cross-System Transfer Performance

As expected, the transfer of explicitly curated user models yielded better results than alternative transfer approaches and the baseline. The reason for the good performance of open curated models can be analyzed by comparing the spaces formed by vectors of the bookmarked and non-bookmarked CoMeT talks with vector spaces of different types of transferred information to the information contained .

The results of this analysis are shown in Figure 3 which plots the distribution across users of average similarities between five representations of user interests – (bookmarked talks in CoMeT and four kinds of models transferred from SciNet) on one side and bookmarked/non-bookmarked CoMeT talks (CP/CN) on the other side. The foure kinds of transferred models are SciNet implicit/explicit keywords (SIK/SEK) and implicit/explicit paper (SIP/SEP). As expected, bookmarked talks in CoMeT are pairwise closer to each other (CP_CP boxplot) than to non-bookmarked ones (CP_CN boxplot). More interesting is that both manipulated SciNet keywords (SEK_CP boxplot) and shown SciNet keywords (SIK_CP

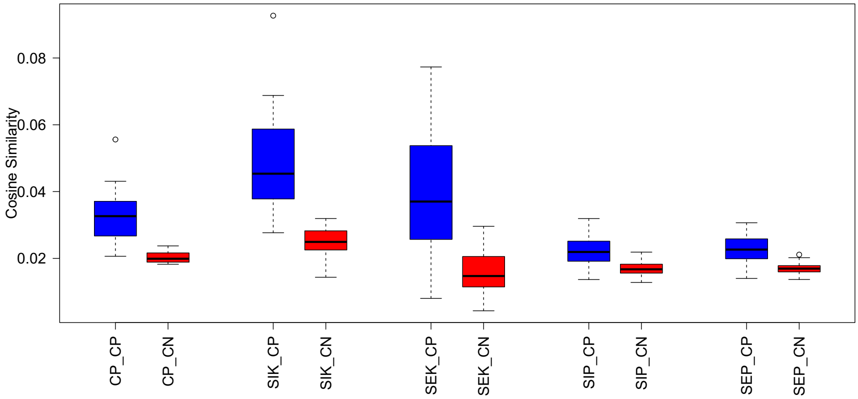

Fig. 3. Distributions of similarities between CoMeT talks five representations of user interests (bookmarked talks and 4 kinds of transfer models). The box-and-whiskers plot shows the median (band inside the box), first and third quartiles (edges of the box), 9th and 91th percentiles (whiskers), and values outside the percentiles as circles. Details for each of plot are in the section 7.

boxplots) are even closer to the space of bookmarked talks, separating them quite well from non-bookmarked ones (SEK_CN and SIK_CN boxplots, respectively). As a result, incorporating open curated keywords into the Centroid model helped alleviate the cold-start problem and improved recommendation performance. In contrast, implicit models built from all shown SciNet papers and bookmarked SciNet papers (SIP_CP and SEP_CP, respectively) are quite far from the space of bookmarked talks and offer poor separation of this space from the non-bookmarked talks (SEP_CN and SIP_CN, respectively). Consequently, implicit models add more noise than value and damage recommendation.

8 Conclusions

This paper explores a novel approach to cross-system personalization based on transferring an explicit, open, and editable user model maintained by one system to another system in a similar, yet different domain. More specifically, we explored whether an open user model maintained by SciNet exploratory literature search system could help recommend relevant research talks to users of CoMeT talk sharing system who have established SciNet interest models. The impact of open model transfer was compared with a baseline case that uses no information from SciNet and with more traditional implicit model transfer based on information about retrieved and marked papers. The comparison examined an overall impact of the transfer as well as its value in a cold-start situation when the user has none or too few talks bookmarked in CoMeT. The results showed that cross-system model transfer is a challenging task. We were not able to register a significant impact of model transfer in a general situation. However, the study demonstrated a significant positive impact of the open model transfer in the *cold-start* case. It also demonstrated that the use of open, explicitly curated user models is critical for the success of user model transfer: the transfer of SciNet data in the form of implicit model damaged the performance of talk recommendation. The analysis of differences between explicit and implicit model hinted that the use of explicitly curated models reduces the noise in modeling user interests. An interesting research question that we will leave for the future work is to what extent machine learning approaches could simulate model curation focusing on better representation of user interests. We plan to explore several feature selection approaches and compare them with the impact of manual curation.

References

1. Berkovsky, S., Kuflik, T., Ricci, F.: Mediation of user models for enhanced personalization in recommender systems. User Modeling and User-Adapted Interaction **18**(3), 245–286 (2008)
2. Cremonesi, P., Tripodi, A., Turrin, R.: Cross-domain recommender systems. In: 2011 IEEE 11th International Conference on Data Mining Workshops, pp. 496–503. IEEE, December 2011
3. Fernadez-Tobis, I., Cantador, I., Kaminskas, M., Ricci, F.: Cross-domain recommender systems: a survey of the state of the art. In: the 2nd Spanish Conference on Information Retrieval
4. Li, B., Yang, Q., Xue, X.: Can movies and books collaborate?: cross-domain collaborative filtering for sparsity reduction. In: Proceedings of the 21st International Jont Conference on Artifical Intelligence, IJCAI 2009, pp. 2052–2057. Morgan Kaufmann Publishers Inc, San Francisco (2009)
5. Loizou, A.: How to recommend music to film buffs: enabling the provision of recommendations from multiple domains. Thesis (2009)
6. Pan, W., Liu, N.N., Xiang, E.W., Yang, Q.: Transfer learning to predict missing ratings via heterogeneous user feedbacks. In: Proc. of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Three, IJCAI 2011, pp. 2318–2323. AAAI Press (2011)
7. Ruotsalo, T., Peltonen, J., Eugster, M., Głowacka, D., Konyushkova, K., Athukorala, K., Kosunen, I., Reijonen, A., Myllymäki, P., Jacucci, G., Kaski, S.: Directing exploratory search with interactive intent modeling. In: Proceedings of the 22nd ACM International Conference on Conference on Information & Knowledge Management, CIKM 2013, pp. 1759–1764. ACM, New York (2013)
8. Sahebi, S., Brusilovsky, P.: Cross-domain collaborative recommendation in a cold-start context: the impact of user profile size on the quality of recommendation. In: Carberry, S., Weibelzahl, S., Micarelli, A., Semeraro, G. (eds.) UMAP 2013. LNCS, vol. 7899, pp. 289–295. Springer, Heidelberg (2013)

9. Schein, A.I., Popescul, A., Ungar, L.H., Pennock, D.M.: Methods and metrics for cold-start recommendations. In: Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2002, pp. 253–260. ACM, New York(2002)
10. Schwab, I., Pohl, W., Koychev, I.: Learning to recommend from positive evidence. In: Proceedings of the 5th International Conference on Intelligent User Interfaces, IUI 2000, pp. 241–247. ACM, New York (2000)
11. Shi, Y., Larson, M., Hanjalic, A.: Tags as bridges between domains: improving recommendation with tag-induced cross-domain collaborative filtering. In: Konstan, J.A., Conejo, R., Marzo, J.L., Oliver, N. (eds.) UMAP 2011. LNCS, vol. 6787, pp. 305–316. Springer, Heidelberg (2011)