Recommending Research Colloquia: A Study of Several Sources for User Profiling

Shaghayegh Sahebi
Intelligent Systems Program
University of Pittsburgh
Pittsburgh, 15260, USA
shs106@pitt.edu

Chirayu Wongchokprasitti
School of Information Sciences
University of Pittsburgh
Pittsburgh, 15260, USA
chw20@pitt.edu

Peter Brusilovsky
School of Information Sciences
University of Pittsburgh
Pittsburgh, 15260, USA
peterb@pitt.edu

ABSTRACT

The study reported in this paper is an attempt to improve content-based recommendation in CoMeT, a social system for sharing information about research colloquia in Carnegie Mellon and University of Pittsburgh campuses. To improve the quality of recommendation in CoMeT, we explored three additional sources for building user profiles: tags used by users to annotate CoMeT's talks, partial content of CiteULike papers bookmarked by users, and tags used to annotate CiteULike papers. We also compare different tag integration models to study the impact of information fusion on recommendations outcome. The results demonstrate that information encapsulated in CiteULike bookmarks generally helps to improve several aspects of recommendation. The addition of tags by fusing them into keyword profiles helps to improve precision and novelty of recommendation, but may harm systems ability to recommend generally interesting talks. The effects of tags and bookmarks appeared to be stackable.

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval - *Information Filtering*

General Terms

Experimentation, Performance

Keywords

Recommendation System, User Profile Fusion, Talks, Papers, Tags, Keywords

1. INTRODUCTION

In academic environment, a short (typically one hour long) research colloquium is one of the common ways to disseminate new ideas and obtain valuable community feedback. A research university holds dozens to hundreds of colloquia every semester. Talks feature a range of speakers from world known researchers to PhD students. Some colloquia are arranged well in advance while others are organized when opportunity comes in a matter of days. A way of sharing information about the colloquia is by posting paper flyers, announcing the talks on a dedicated department page, and sending e-mails to mailing lists and colleagues. While this approach may work well in a small university with well-positioned centers of expertise, it is not efficient in the context of large universities and, especially, -located universities where talks on a similar subject could be organized by many different departments. This is what we are dealing in Pittsburgh where two large research universities, Carnegie Mellon and University of Pittsburgh are located within 15 minutes of walk from each other. To improve the awareness of local researchers about relevant talks organized on both campuses, our research team developed CoMeT, a collaborative system for sharing information about research colloquia. The system was launched in Fall 2009 as a typical collaborating tagging system, which allows its users to announce, find, bookmark, and tag talks (Fig. 1). Over the first year of its use, the system gained some reasonable popularity. Over 750 talks were posted to the system and most popular talks have been accessed over 150 times.

Among our original goal was extending CoMeT with a recommendation feature. Given a relatively short-lived nature of colloquia, we focused on a content-based recommender, which can build profile of interest of individual users and recommend new talks immediately after their posting. The first version of our recommender used a traditional keyword vector profiling approach and considered talk abstract and title for profile construction. This version debuted in early 2010 and was explored by some frequent users of the CoMeT. However, the impact of the recommender component was harmed by "under-contribution", a common phenomenon of social systems [6]. While talks posted to the system were getting reasonable number of hits, the number of bookmarks was too small. It caused a lower-than-expected quality of recommendations. Worse, since many users never bookmarked a talk, the system was not able to generate any recommendation for them.

To resolve these two problems, we considered additional sources for building user profiles and specifically explored two ideas. One idea was just mildly innovative: user tags. As has been already argued, user tags provide good information about user interests [4; 7; 13]. Another idea was much less explored: using information about user interests from some other system. In our case, we

HetRec '10, September 26, 2010, Barcelona, Spain.

considered research paper sharing systems such as CiteULike[1] or BibSonomy [8]. While research colloquia and academic papers are different artifacts, they looked close enough to explore cross-modeling of user interests. Among paper sharing systems, we picked CiteULike. Informal interviews with CoMeT users demonstrated that a number of them had well-established CiteULike accounts with dozens of bookmarked papers. A collection of papers bookmarked by users can provide two more sources for profiling – titles and abstracts of papers and user tags.

The study reported in this paper is an attempt to improve CoMeT's content-based recommendation using the three sources of user data mentioned above: tags used by users to annotate CoMeT's talks, tags used to annotate CiteULike's papers and partial content of CiteULike papers (such as abstracts and titles). We compare different integration models to study the impact of information fusion on recommendations outcome. In section 2, we will review some related works. An introduction to CoMeT system is provided in section 3. After that, we will present different integration models and the way user profiles are built in section 4. We will see the experimental results in section 5 and conclusion in section 6.

2. RELATED WORK

Cold start [12] issue is a classic problem for recommender systems. With the growth of the number of various social and personalized systems, which maintain user profiles, the idea of using profiles from one system to improve the quality of personalization in another system is becoming more and more popular [9; 14; 15]. Indeed, even though a lot of users themselves are new in a specific system, they might have data or even user profiles in other similar systems. Using this data is one way to help solve the cold start problem. Moreover, an elaborate fusion of user profiles or data from several systems can also increase the precision of recommendation. In the area of recommender system this multi-system approach and integrating user models built by other personalization systems is called user model mediation. In [9] use of the Unified User Context Model (UUCM) from different personalization systems is proposed. Each personalization system can extract its required data from UUCM and update it separately. User model mediation [1] has been studying in recent years. In [1] integration of user models is done via a user model mediator, which imports user models collected by a group of personalization systems. Multiple ways for integrating user models in adaptive learning systems, collected by multiple educational systems is provided in [14]. Cross-item mediation is one of popular ways applied to existing recommendation techniques, such as content-based filtering [17], collaborative filtering [11], and hybrid approaches [3]. This mediation technique assumes that the items are similar to the past items, including cross-items from the other remote systems users liked, should be recommended to them.

Use of different data sources in recommendation systems is not limited to integrating user models of various personalization systems. There has been some prior research on integrating multiple recommendation methods, integrating collaborative filtering and content-based methods and using ontologies in recommendation systems, as hybrid recommendation systems [2; 3]. The use of tags as a source of additional information in recommender system is also growing in popularity. Specifically, in the area of content-based recommendation, a number of papers suggested the user of tags for building user profiles and reported an improved recommendation quality [4; 7; 10; 13; 16]. A multivariate Poisson model for naive Bayes text classification is used in [4] to infer content-based and tag-based user profiles. Tags are used in [7] to build tag-based user profiles for providing recommendations for a music search portal and discussed a search-based recommendation method. Folksonomy tags are used in [10] to classify Web pages for building a Web page recommendation system.

3. COMET: SYSTEM FRAMEWORK

The CoMeT is the collaborative colloquia sharing system built for Pittsburgh-based researchers. It supports both passive and active dissemination. Every user can post a talk by filling a simple form. Mandatory fields include title, speaker, date, time, and location. The talk description field is not mandatory, although almost all posted talks included talk abstract and information about the speaker in the description. Posted talks can be browsed by calendar (Fig. 1), standing series, organizing departments, and other ways. The system can also disseminate talks using iCal, Google Calendar and RSS feeds.

When users find the interesting talks they can bookmark them and add tags and comments (Fig. 2). The users also can share talks with their friends by email. The user cumulative activity related to a talk (viewing, tagging, sending by e-mail) is visualized in all lists where the talk is shown (Fig. 1). This social link annotation feature is known as social navigation [5]. It provides a simple way to use community wisdom for guiding users to good talks. Recommendations in the current version of CoMeT are also displayed in the form of link annotation. I.e., instead of showing all recommendations as a ranked list, the system adds red tag "Recommended" to recommended talks in all contexts where a link to this talk is shown (Fig. 1).

4. RECOMMENDATION APPROACHES

We explored two ways to improve recommendations in CoMeT: by using information about bookmarked papers from CiteULike in addition to the standard use of information about bookmarked talks from CoMeT and by using tags for better representing information about talks (and user interest) in addition to standard use of text-only information from talk descriptions. In addition, we combined both approaches – i.e., used both kinds of information (descriptions and tags) from both systems.

While fusing information from two systems is relatively straightforward (from the recommender engine's point of view, a bookmarked talk and a bookmarked paper is simply a bag of words plus a bag of tags), fusing tags and text in both item representations and profiles is not obvious and can be done in several ways. We start the presentation of our recommendation approaches with introducing several representation models, which explored various ways to fuse keywords and tags. After that we explain the user profiling and recommendation approaches based on these representations.

[1] http:///www.citeulike.org

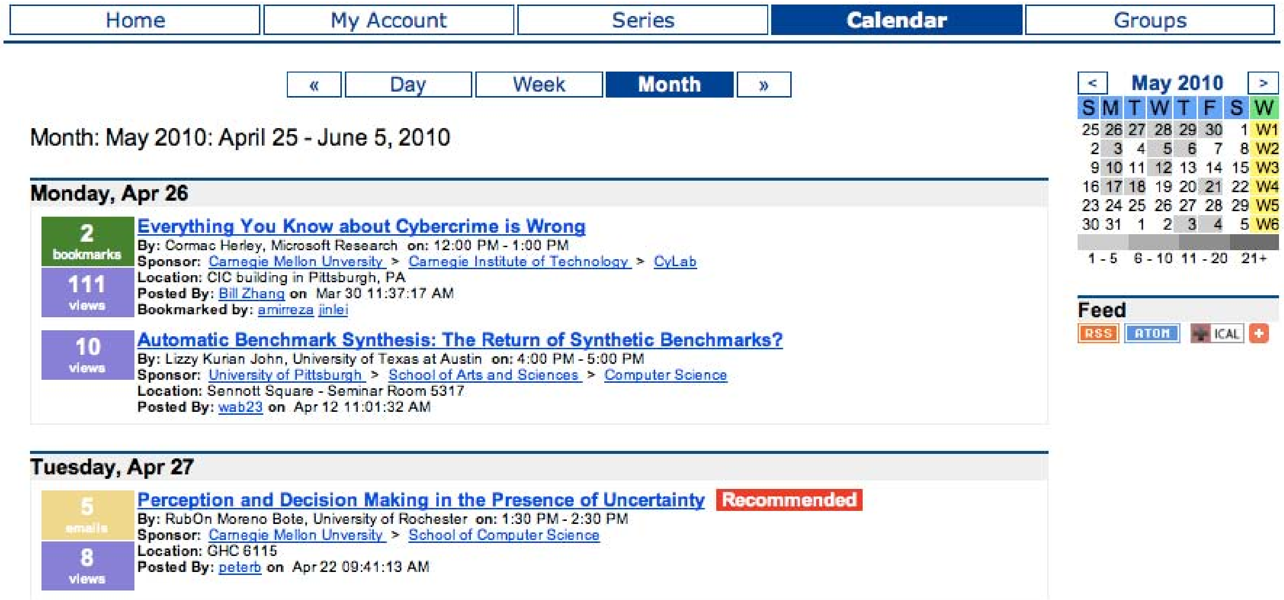

Figure 1. A Calendar view of CoMeT interface

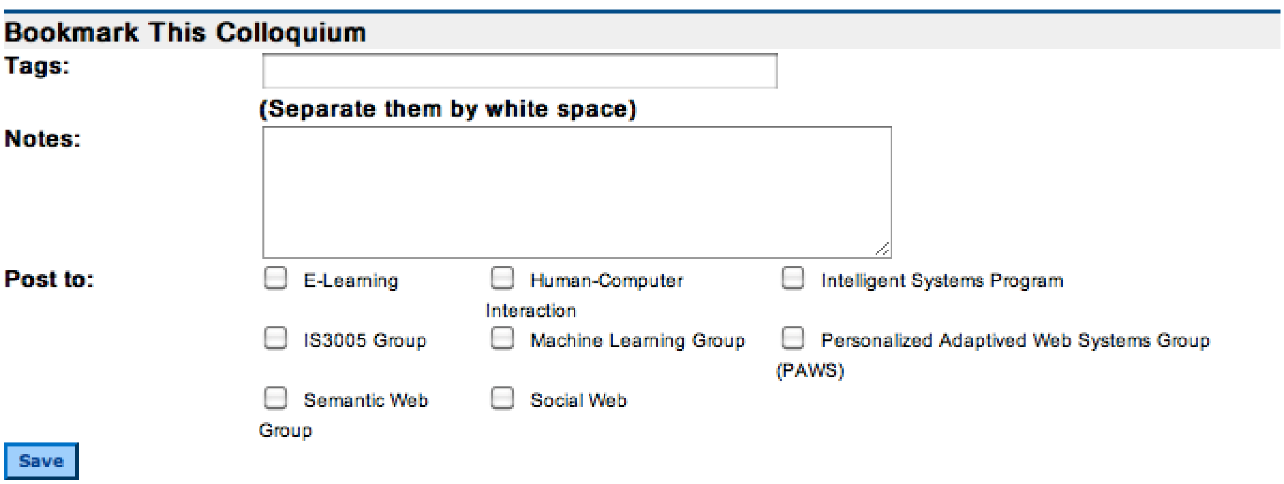

Figure 2. Bookmark the Talk

4.1 Document Representation Models

There are many ways to combine various sources of information for building user profiles. We have used tags and keywords utilized in abstracts and titles of CoMeT talks and CiteULike papers. To construct user profiles, we used the following models:

Keywords Only (KO): To represent documents in this model, only keywords extracted from documents' titles and abstracts are used. Each document is considered as a bag of words and represented as a vector in keywords vector space weighted by TF.IDF weightening scheme ($d_c = (w_{1,c}, w_{2,c}, \ldots, w_{l,c})$ which $w_{i,c}$ shows the weight of i^{th} keyword in document c and l is the total number of keywords).

All CiteULike papers can be represented as a $k \times l$ matrix D_c (k is the number of CiteULike papers and l is the number of keywords used in those papers) and each CoMeT talk is represented in an $e \times m$ matrix D_t (with e as total number of talks and m as total number of keywords). To integrate these two sources of data in this model, we obtain a $(k + e) \times (l + m - o)$ matrix D where o is the number of common keywords between two CoMeT and CiteULike systems.

Keywords+n*Tags (KnT): In this model, tags are considered as regular keywords and we treat each document as a bag of words containing document's abstract, title and tags. Each tag appears n times in this bag of words. Each document is represented as a vector in keywords and tags vector space weighted by TF.IDF weightening scheme ($d_c = (w_{1,c}, w_{2,c}, \dots, w_{l+j-p,c})$ which $w_{i,c}$ shows the weight of i^{th} keyword in document c, l is the total number of keywords, j is the total number of tags, and p is the number of common terms between tags and keywords).

In this case, we can obtain merged documents matrix just like the previous model.

Keywords Concatenated by Tags (KCT): In this model, we consider tags as a separated source of information and treat them separately. We obtain a bag of keywords and a bag of tags distinctly for each document. Using TF.IDF weightening scheme, a tags vector and a keywords vector is built for each document. Next, each document is represented by concatenating keywords and tags vectors as one vector in keywords and tags vector space ($d_c = (w_{1,c}, w_{2,c}, \dots, w_{l,c}, t_{1,c}, t_{2,c}, \dots, t_{j,c})$ which $w_{i,c}$ shows the weight of i^{th} keyword in document c, l is the total number of keywords, $t_{i,c}$ shows the weight of i^{th} tag in document c, and j is the total number of tags).

In this case we will have an $e \times (l + j)$ CoMeT talk's matrix (D_t) which e is the number of CoMeT talks, l is the total number of keywords in CoMeT, and j is the total number of tags in it. We will also have a $k \times (m + i)$ CiteULike paper's matrix (D_c) which k is the number of CiteULike talks, m is the total number of keywords in CiteULike, and i is the total number of tags in it. After merging these two matrixes, we will have an $(e + k) \times (m + I + l + j - o - p)$ matrix D showing all documents in keywords and tags vector space in which o and p are respectively the number of common keywords and tags between two systems.

To study the impact of various sources of information on recommendation systems, we utilized each of the aforesaid models once by only CoMeT system's data, and another time including CiteULike's data sources.

4.2 Recommending Talks to Users

To recommend related talks to users, we have used K-nearest neighbor method. In this method, top K closest documents (talks) to the user profile are recommended to each user. To do this, we have to represent user profiles in documents' vector space.

User profiles are built based on users' bookmarked and rated talks and papers. We represent each user's bookmarked and rated talks and papers, weighted by user ratings, in a vector in talks and papers vector space. To obtain keyword-based user profiles, we used the document representation models presented in section 4.1. Keyword-based user profile (UP) is obtained by multiplying the vector of user documents in documents vector space (U) by the matrix of document keywords represented in keywords and tags vector space (D):

$$UP = U.D' \qquad (1)$$

Resulted user profile (UP) is a vector consisted of user's related keywords (and tags), weighted based on the importance of each keyword (or tags). In this case, which user profiles are represented in the same vector space as CoMeT system's talks, we can find K closest talks, which user has not seen yet, and recommend them to user. To measure the distance between talks and user profiles, we use cosine distance measure.

5. EXPERIMENTAL RESULTS

To evaluate the recommendation approaches presented above, we run a small-scale user study. For the study we selected 8 real users of CoMeT system who also had accounts in CiteULike. All users were PhD students at the School of Information Sciences, University of Pittsburgh. For each user, we ran the explored recommendation approaches. The results were merged, randomly ordered and presented to the user for evaluation. To evaluate each item in the merged recommended list, a user had to answer three questions measuring relevance to research interest, overall interest, and novelty:

- Is this talk related to your interest? (yes/no question)
- How interesting this talk to you? (in 5-point scale)
- If the talk is related to your interests, how novel is this talk to you? (in 5-step scale)

Altogether, we compared five models: KO, KnT (with n = 1, 2, 5), and KCT. Each model was used twice – once to recommend talks using only CoMeT data and one using both, CoMeT and CiteULike.

For each model we calculated three measures (to explore relevance, interest, and novelty correspondingly) for each position in the top-10 ranked list produced by the model. To evaluate relevance (which was measured by yes/no answers) we used traditional precision (Table 1). To evaluate interest (measured by 5-point scale) we used nDCG (Table 2). To evaluate novelty, we averaged the novelty ratings of recommendations for first to tenth recommendation in the ranked list produced by each approach (Table 3). Non-relevant recommendations were considered as having zero novelty. To reduce the volume of reported data, in all tables, we only show results for the best n value for KnT model, which was n=1.

As we can see in Table 1 and Figure 3, adding tag using fusion approach (KnT), results in better cumulative precision for top 10 recommendations. The exceptions are the very top positions in the recommendation list where KO model works better in both cases. This interesting effect is caused by different behavior of these two models. For KO model the precision is high at the top positions, but then drops rapidly. The precision of KnT model is more stable and overruns KO at position 5.

Adding CoMeT data in both KnT and KO models increases the precision drastically. This both tags (with KnT) and CoMeT can help with the precision. Moreover, these two effects seem to be stackable: KnT model which includes both CoMeT and CiteULike (CUL) data has the best cumulative precision.

At the same time, adding tags using KCT model degrades system's precision. This might be because of high dimensionality of our vector space model when we concatenate keywords vector with tags vector. In this case, the distance between documents and user profile increases and decreases the variance between similarities of user profile to different talks. We can see that different sources of information may result in less precise results if we don't integrate them in a right way.

Table 1. Precision results for different models with different number of recommendations

Precision		1	2	3	4	5	6	7	8	9	10
Only CoMeT Data	**KO**	0.83	0.67	0.72	0.63	0.6	0.56	0.57	0.5	0.51	0.51
	KnT *n* **= 1**	0.5	0.5	0.58	0.59	0.57	0.58	0.57	0.58	0.6	0.57
	KCT	0.5	0.33	0.39	0.46	0.47	0.53	0.52	0.5	0.5	0.53
CoMeT + CiteULike Data	**KO**	0.83	0.83	0.67	0.75	0.73	0.69	0.64	0.63	0.56	0.57
	KnT *n* **= 1**	0.63	0.69	0.71	0.72	0.73	0.73	0.71	0.7	0.68	0.67
	KCT	0.38	0.44	0.42	0.47	0.48	0.52	0.5	0.49	0.53	0.55

Table 2. nDCG Results for different models and different number of recommendations

nDCG		1	2	3	4	5	6	7	8	9	10
Only CoMeT Data	**KO**	0.9	0.88	0.89	0.93	0.92	0.94	0.95	0.95	0.95	0.96
	KnT *n* **= 1**	0.9	0.85	0.82	0.83	0.87	0.88	0.89	0.9	0.91	0.93
	KCT	0.84	0.88	0.89	0.9	0.9	0.91	0.92	0.92	0.94	0.95
CoMeT + CiteULike Data	**KO**	0.84	0.91	0.9	0.92	0.93	0.94	0.95	0.96	0.96	0.96
	KnT *n* **= 1**	x	0.9	0.89	0.88	0.9	0.92	0.92	0.94	0.94	0.95
	KCT	0.77	0.85	0.84	0.81	0.83	0.84	0.86	0.88	0.91	0.92

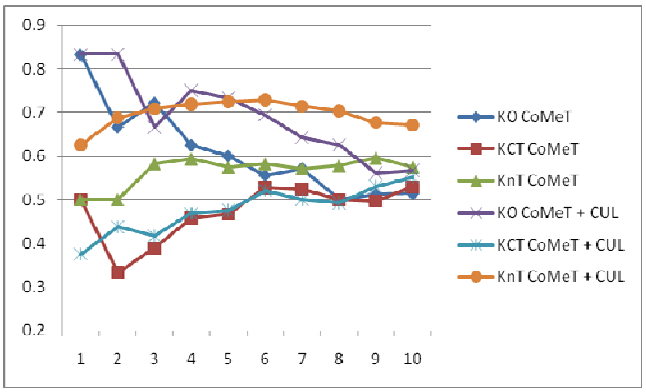

Figure 3. Precision Results for different models with different number of recommendations

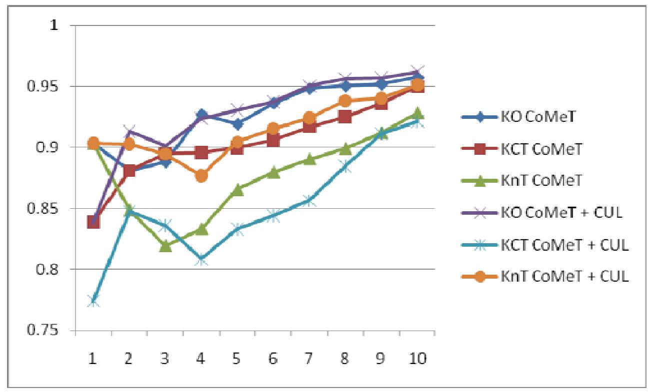

Figure 4. Interest results (nDCG) for different models with different number of recommendations

The results for "interest" in Table 2 and Figure 4 are a bit different. Generally, we can see the positive effect of using CiteULike data. In both, KCT and KnT models, using CiteULike data in addition to CoMeT data boosts user cumulative interest in returned talks, although this difference is very small for KO model. At the same time, best results are produced by tagless KO model both with and without CiteULike data. These results are very close to each other and better than other models in general.

To explain these results, we should stress that we deliberately separated "relevance", which was understood as a fit to user research work and "interest", which was understood as an overall attraction of an item. Through the CoMeT experience, we observed that many users seem to be interested in some talks on general topics (like art and politics), which had little in common with their research interests. A separation of relevance and interest allowed cases, where a talk is rated as interesting, yet non-relevant. The analysis of user rating data confirmed that we were correct – there were a number of talks like that for almost all users. We think that the decrease of system ability to recommend *interesting* talks with the addition of tags can be caused by the increased focus of relevance encapsulated in tags. It naturally decreased system's ability to recommend interesting, but not relevant talks. This is a natural outcome of user tagging behavior, which was focused mostly on their research interests. As a result, a simpler KO model is able to better grasp user overall interests.

The evaluation of novelty gives rather opposite results (Table 3 and Figure 5). Here, adding tags using KnT fusion model provides the largest positive impact. As we can see, both KnT models (using CoMeT data or both CoMeT and CUL data) produce more novel recommendations to users. The impact of CUL data is not consistent. KnT model using both CoMeT and CUL data recommends more novel talks to users. This shows that adding different sources of information, especially tags in this case, can improve the novelty of recommendations. This makes sense if we notice that tags are provided by users and include a broader range of vocabulary for describing a talk and paper. Each user uses tags to describe a document from his/her own point of view which might be different from the terms included in the document's abstract or title. On the other hand, we can see that in KO model, adding CiteULike data decreases the average of novel talks recommended to users. This is due to the distinctive natures of CoMeT and CiteULike systems. Users usually use CiteULike for adding, reviewing and rating related papers to their research field. This needs a user to spend a noticeable amount of time on a paper and users prefer to review related papers to their field of research. On the other hand, CoMeT contains information about talks happening within a specific time given on a particular date. It is more plausible for a user to bookmark a more interesting, more novel, even less relevant talk knowing that he/she might miss this amount of information given in a limited time. As a result, CoMeT user profiles include wider area of user interests with respect to CiteULike user profiles, which usually contain more relevant documents.

Table 3. Novelty Results for different models with different number of recommendations

Precision		1	2	3	4	5	6	7	8	9	10
Only CoMeT Data	**KO**	1.75	1.69	1.67	1.72	1.7	1.65	1.66	1.55	1.49	1.44
	KnT $n = 1$	1.88	1.75	1.67	1.88	1.88	1.88	2	2.03	1.99	1.93
	KCT	2	1.5	1.54	1.56	1.55	1.6	1.63	1.58	1.5	1.5
CoMeT + CiteULike Data	**KO**	1.88	1.44	1.33	1.5	1.5	1.52	1.61	1.47	1.44	1.36
	KnT $n = 1$	1.75	2.19	1.79	2.06	2.2	2.08	2.02	2.19	2.06	1.96
	KCT	1.38	1.31	1.38	1.47	1.58	1.6	1.52	1.47	1.61	1.64

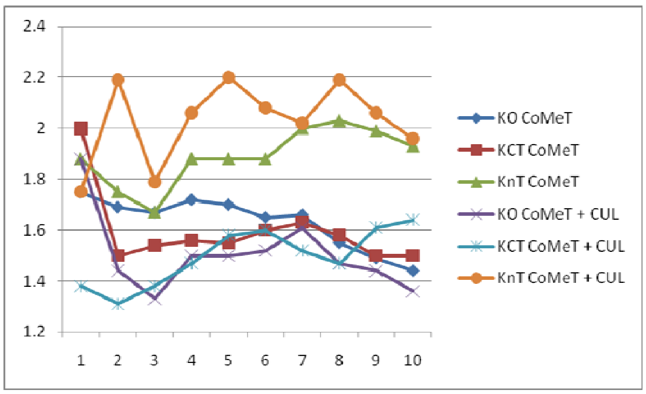

Figure 5. Novelty Results for different models with different number of recommendations

Comparing the results, we can see that there are several trends. First the addition of tags using KnT fusion model helps to improve both novelty and relevance of results. This effect is more pronounced for novelty. In contrast, when user general interests are considered, tagless KO model produced slightly better results. As for KCT concatenation model, it seems to have real problems: it decreases system performance for all kind of measures producing worst performances in most of the cases.

The effect of adding CiteULike is more consistent. It typically produces better results for all measures, although its effect for interest measure is negligible. Interesting enough is that the effects of adding tags and adding data appeared to be stackable. i.e., an approach, which uses both tags and CiteULike "stacks" the separate effects of the component approaches resulting in best approaches for relevance and novelty measures.

6. CONCLUSION

In this work, we used various sources of information to build user profiles for recommending talks in CoMeT system. To do this, we utilized both CoMeT and CiteULike documents (talks and papers) abstracts, titles, and tags. We built content-based and content and tag-based user profiles from both systems and recommended top ten talks to CoMeT users.

Based on discussions in experimental results section, we can see that including another reliable user profile would increase precision of recommendations but the way to augment the additional profile to existing user profile matters. Results may be different for separate injection features. Specially, when we take tags into account, we should be more concerned about the augmentation method; otherwise in some cases it might degrade the recommendation system's performance.

In addition, we can see that relevancy of recommended documents measured by precision increases using CiteULike data for all models, while results of interestingness of recommended talks vary by including CiteULike data. Adding tags increases the novelty of recommendations both using CoMeT and CiteULike data while it increases their relatedness in larger number of recommendations. As a conclusion, injection of keywords from another source of data, for obtaining relevant content-based recommendations, is more reliable than including tags while for getting more interesting or novel recommendations, including tags from various sources of information.

For future work, we plan to explore deeper the issue of relatedness vs. interestingness of talks to users. Usually recommendation systems recommend relevant items to users but users may prefer interesting items rather than only relevant ones. Additionally, we have to include user behavior, like sequence of user visits, as another source of information to our recommendation models.

7. REFERENCES

[1] BERKOVSKY, S., KUFLIK, T., RICCI, F. Mediation of user models for enhanced personalization in recommender systems, User Modeling and User-Adapted Interaction (2008) 18, 3, 245-286.

[2] BURKE, R. Hybrid recommender systems: survey and experiments. User Modeling and User-Adapted Interaction. (2002) 12, 4, 331-370.

[3] CHUNG, R., SUNDARAM, D., SRINIVASAN, A. Integrated personal recommender systems in Proceedings of the 9th International Conference on Electronic commerce (ICEC 2007),, . (USA 65-74.

[4] DE GEMMIS, M., LOPS, P., SEMERARO, G., BASILE, P. Integrating tags in a semantic content-based recommender in Proceedings of the 2008 ACM conference on Recommender systems, RecSys '08 (Lausanne, Switzerland), ACM 163-170.

[5] FARZAN, R., BRUSILOVSKY, P. AnnotatEd: A social navigation and annotation service for web-based educational resources, New Review in Hypermedia and Multimedia (2008) 14, 1, 3-32.

[6] FARZAN, R., DIMICCO, J., MILLEN, D., DUGAN, C., GEYER, W., BROWNHOLTZ, E. Results from deploying a participation incentive mechanism within the enterprise in Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems, CHI '08 (Florence, Italy), ACM 563-572.

[7] FIRAN, C.S., NEJDL, W., PAIU, R. The benefit of using tag-based profiles in Proceedings of Fifth Latin American Web Congress (LA-Web 2007) (Santiago de Chile, November 2007), IEEE 32–41.

[8] HOTHO, A., JÄSCHKE, R., SCHMITZ, C., STUMME, G. Information Retrieval in Folksonomies: Search and Ranking, In: The Semantic Web: Research and Applications, 2006, 411-426.

[9] MEHTA, B., NIEDERÉE, C., STEWART, A. Towards cross-system personalization in Proceedings of 3rd International Conference on Universal Access in Human-Computer Interaction, 22-27 July 2005.

[10] NIWA, S., DOI, T., HONIDEN, S. Web Page Recommender System based on Folksonomy Mining for ITNG'06 Submissions in Proceedings of Third International Conference on Information Technology: New Generations, ITNG 2006. 388-393.

[11] RAFTER, R., SMYTH, B. Passive profiling from server logs in an online recruitment environment. In Proceedings of ,, . in Proceedings of the IJCAI Workshop on Intelligent Techniques for Web Personalization (ITWP 2001) (USA) 35-41.

[12] SCHEIN, I.A., POPESCUL, A., UNGAR, H.L., PENNOCK, M.D. 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 '02) (Finland) 253-260.

[13] SIERSDORFER, S., SIZOV, S. Social recommender systems for web 2.0 folksonomies in Proceedings of HT '09: Proceedings of the 20th ACM conference on Hypertext and hypermedia (Torino, Italy), ACM 261-270.

[14] SOSNOVSKY, S., BRUSILOVSKY, P., YUDELSON, M., MITROVIC, A., MATHEWS, M., KUMAR, A. Semantic Integration of Adaptive Educational Systems, In: Kuflik, T., Berkovsky, S., Carmagnola, F., Heckmann, D., and Kruger, A., Eds. Advances in Ubiquitous User Modelling, 2009, 134-158.

[15] SOSNOVSKY, S., DOLOG, P., HENZE, N., BRUSILOVSKY, P., NEJDL, W. Translation of overlay models of student knowledge for relative domains based on domain ontology mapping in Proceedings of 13th International Conference on Artificial Intelligent in Education, AI-ED 2007 (Marina Del Rey, CA, July 9-13, 2007), IOS 289-296.

[16] TSO-SUTTER, K., MARINHO, L., SCHMIDT-THIEME, L. Tag-aware recommender systems by fusion of collaborative filtering algorithms in Proceedings of the 2008 ACM symposium on Applied computing, SAC '08 (Fortaleza, Ceara, Brazil, March 16 - 20, 2008), ACM 1995-1999.

[17] WHITE, R., RUTHVEIN, I., JOSE, J.M. Finding relevant documents using top ranking sentences: An evaluation of two alternative schemes in Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '02) (Finland) 57-64.