Social Navigation Support in a Course Recommendation System

Rosta Farzan[1] and Peter Brusilovsky[1,2]

[1] Intelligent Systems Program and [2] School of Information Sciences
University of Pittsburgh, Pittsburgh PA 15260, USA
rosta@cs.pitt.edu, peterb@pitt.edu

Abstract. The volume of course-related information available to students is rapidly increasing. This abundance of information has created the need to help students find, organize, and use resources that match their individual goals, interests, and current knowledge. Our system, CourseAgent, presented in this paper, is an adaptive community-based hypermedia system, which provides social navigation course recommendations based on students' assessment of course relevance to their career goals. CourseAgent obtains students' explicit feedback as part of their natural interactivity with the system. This work presents our approach to eliciting explicit student feedback and then evaluates this approach.

1 Introduction

Information technology (IT) has rapidly changed many aspects of receiving a college education. The volume of course-related information available to students is rapidly increasing. This abundance of information has created the need to help students find, organize, and use resources that match their individual goals, interests, and current knowledge. One of the concerns students have is to make decisions about which courses to take. The concern is more serious for graduate students who have more freedom to choose courses while they care more about taking courses that contribute to their progress towards career goals. To make these decisions, they use information from course catalogs and schedules, consult with their advisors, and seek guidance from their classmates, especially those with similar interests. To give better decision-making support to students who wish to make relevant course choices, we have developed a course recommendation system, CourseAgent, which integrates all available information about courses and provides personalized access to it.

CourseAgent is a community-based recommendation system that employs *social navigation* [5] to tackle the problem of information overload. Community-based systems integrate explicit and implicit feedback provided by the community of users regarding information items and distill the collective wisdom of the community to help individuals. *Explicit feedback* is registered when a user rates an item as interesting or relevant. *Implicit feedback* is extracted from user actions that indirectly provide some evidence about item quality or relevance - such as link selection, reading time, bookmarking, etc. A challenge for recommendation systems is to encourage users to provide explicit feedback. Explicit feedback is considered the most

V. Wade, H. Ashman, and B. Smyth (Eds.): AH 2006, LNCS 4018, pp. 91 – 100, 2006.

reliable source of information for personalization; however, users rarely provide it since they don't perceive this activity as essential to their work with the system [4].

CourseAgent provides community-based recommendations of courses using explicit feedback - students' assessment of course relevance to their various career goals. To elicit feedback from users, the system employs a specific "do-it-for-yourself" approach. The main theme of this approach is to obtain students' explicit feedback *implicitly*, as part of their natural interaction with the system. This research study presents our approach for eliciting feedback from the students, and then evaluates our approach. The rest of the paper is organized as follows: section 2 describes background information and related work, section 3 provides details about different parts of CourseAgent system, details on adaptation and social recommendation, and our approach for eliciting user feedback. Section 5 presents our evaluation methods and the results of this evaluation. We conclude the paper in section 6 and provide several ideas for the future direction of this work.

2 The Under-Contribution Problem in Adaptive Community-Based Systems

There is an increasing focus on creating community-based adaptive Web systems that provide navigation support or collect recommendations based on feedback from the users of the system. Amazon.com recommends items to buy based on activities of other users. MovieLens [9] recommends movies to watch based on the feedback provided by similar users. The I-Spy search engine uses the information provided by their community to re-rank search results [6]. The functionality and precision of these community-based systems is strongly dependent upon the amount of feedback provided by users of the systems. In many cases, the insufficient quantity of contributions from users has damaged the value of these systems. Encouraging users to contribute has become one of the most important challenges to this field.

Since the discovery of the "users do not like to rate" phenomenon, different systems have tried different approaches to collecting user feedback, in order to fuel the recommendation mechanisms. Early works focused on substituting *explicit* feedback, such as relevance rating, with *implicit* feedback, such as time spent reading a page, time spent scrolling a page, or number of clicks [4]. While several studies have demonstrated the potential of implicit feedback in several contexts, it has not emerged as the ultimate solution. In many cases, implicit feedback lacks the required accuracy, damaging the system's precision.

The idea of a more recent "economy" approach is to encourage users' explicit contribution by building a reward mechanism into the system. In their early work, Bretzke and Vassileva [1] tried several reward mechanisms for encouraging contributions to their system resource-sharing system COMTELLA. The system rewards more cooperative users by such incentives as more bandwidth for download, or higher visibility in the community. More recent version of COMTELLA used the rewarding mechanism to regulate the quality of participation [3]. Harper et al. [7] designed an economic model to analyze users' contributions to a movie recommendation web site. The model compares the effort required for providing ratings with the direct and indirect benefits of the contribution. The model provides ideas on how to motivate users' ratings, such as improving the interface to increase the fun and non-predictable personal benefits of rating, and improving the interface for browsing collections of one's own ratings. Ling et al [8] employed social psychology theories to address the problem of under-contribution in online movie recommender community. The results of their study show that uniqueness of contribution can play an encouraging role for the users. Moreover, they found that users are more likely to contribute when the goal is very specific and challenging.

Our work explores an alternative approach to eliciting user feedback that we call "do-it-for-yourself." The main theme of this approach is to encourage users' participation by turning their feedback into an activity that is important and meaningful to them. In other words, we make the achievement of a personal goal dependent upon their contribution to the community. This approach stands somewhat between the two approaches analyzed above. On one hand, we encourage users to provide reliable explicit feedback. On the other hand, this feedback might be considered implicit by a recommendation system since it was provided not for the system (as in the "economy" approach), but rather to achieve the users' own goals.

3 CourseAgent

CourseAgent is an adaptive community-based hypermedia system that provides personalized access to information about courses. CourseAgent was developed for students and instructors in the School of Information Sciences at the University of Pittsburgh and incorporates information about courses offered at the School. However, it can easily be adopted for different programs by merely integrating the program-specific course data into the system.

3.1 Social Recommendation in CourseAgent

CourseAgent is a social navigation support system. It provides recommendation in the form of in-context adaptive annotations instead of generating an out-of-context sorted list of recommended courses. Course information is annotated with adaptive visual cues that help students to select their most appropriate courses. Fig. 1 demonstrates the use of in-context adaptive community-based annotations on the Schedule screen of CourseAgent. The Schedule screen provides different information about courses offered in a specific semester. As does any course schedule, it provides various information about each offered course, such as course number, course title, date and time, location, and information about the instructor. If the student finds a specific course relevant and interesting, she can use the provided link to register for this course or to plan to pursue this in the future (right column). To help the student register and plan decisions, the system attempts to enhance each link with two kinds of community-based annotation displayed as icons to the left of the links. One icon expresses the expected course workload (one shovel for low, two for average and three for a high workload). The other icon expresses the expected relevance of the course to the career goals of the given student (from one thumb up for a relevant course to three for a highly relevant course). The estimated workload and relevance of a specific course is calculated using community feedback about past offerings of this course, as taught by the same instructor. In addition, another kind of icon in the relevance column indicates that the student's advisor considers this course to be relevant for the given student.

Schedule of spring 2006

Taken Courses, Planned Courses, Currently Taken Courses, : Recommend by Advisor, : Degree of Relevance to Career Goals

CRN	Course No	Title	Day	Time	Instructor	Workload	Relevance	Action
2692	TELCOM 2940	PRACTICUM	apt		[illegible]			Plan It
16084	INFSCI 2120	INFORMATION AND CODING THEORY	tue	6:00-8:50 P	[illegible]			Plan It
16077	INFSCI 2130	DECISION ANALYSIS AND DECISION SUPPORT SYSTEMS	wed	6:00-8:50	[illegible]			Plan It
16088	LIS 2194	ETHICS IN THE INFORMATION SOCIETY	mon	3:00-5:50 P	[illegible]			Plan It
16099	INFSCI 2350	HUMAN FACTORS IN SYSTEMS	thu	6:00-8:50 P	[illegible]			Register It
16056	INFSCI 2470	INTERACTIVE SYSTEM DESIGN	wed	6:00-8:50 P	Peter Brusilovsky			Evaluate It

Fig. 1. Checking the schedule in CourseAgent

-Cognitive Science Area

Course No	Course Title	Workload	Relevance	Action
INFSCI 2300	HUMAN INFORMATION PROCESSING			View Feedback
INFSCI 2330	FOUNDATIONS OF COGNITIVE SCIENCE			Plan It
INFSCI 2350	HUMAN FACTORS IN SYSTEMS			

-Cognitive Systems Area

Course No	Course Title	Workload	Relevance	Action
INFSCI 2410	INTRODUCTION TO PARALLEL DISTRIBUTED PROCESSING			Plan It

Fig. 2. The Course Catalog screen in CourseAgent

Similar social navigation support is provided in the Course Catalog screen of the system. In this screen, courses are grouped by *areas of study* defined by the program as shown in Fig. 2. For example, an Information Science degree includes areas such as Cognitive Science, Cognitive Systems, and Mathematical and Formal Foundation. Each course in the catalog is annotated with social recommendation information representing the relevance and workload of the course. However, since different instructors might teach the same course, the average relevance and workload of each course is based upon the average score over all instructors who taught the course.

3.2 Providing Feedback

CourseAgent provides social navigation support by collecting three kinds of information from the community of students: a) the student's self-selected career goals, b) the students' explicit evaluation of course workload, and c) the student's personal rating for career goal relevance for the courses that they have already taken. We have defined an extendable list of 22 career goals that cover different ranges of careers related to the information science field. Students are able to add career goals that they wish to pursue to their profile. In addition, the system provides an interface to evaluate courses already taken. Students are asked to evaluate the relevance of each taken course to each of their career goals on a scale of 1 to 5 and to evaluate the workload of the course on a scale of 1 to 3. Fig. 3 presents the evaluation interface .

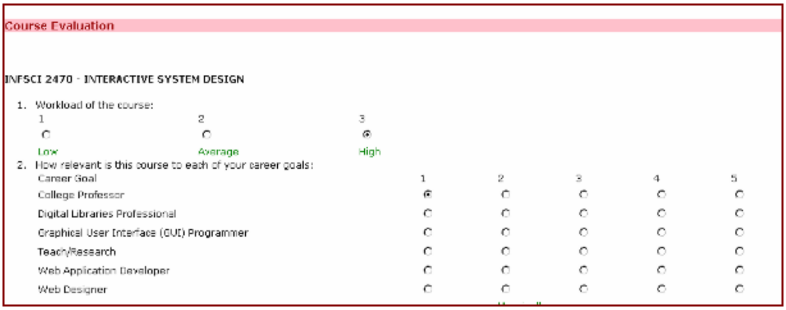

Fig. 3. Evaluation interface of the CourseAgent system

The collected information is used to deliver adaptive annotations presented in the previous section. The overall workload level of the course is computed by simply averaging all ratings provided by the students. The relevance of a course to a student is computed based on the relevance of the course to each of the student's career interests. To compute total relevance, we cannot easily average the relevance to all career goals of the student: A worthy course might be irrelevant to most of the students' career goals while being critical to only one goal. In this case, a simple average will give this a poor relevance rating, while the student might actually be especially interested in taking the course since it is essentially relevant to one of his career goals. To overcome this, we designed a simple algorithm to compute course relevance. The relevance of a course to each career interest of the student ranges from 1 to 5 - where 1 is not relevant and 5 is relevant in an essential way. Courses with a relevance level of 3 and above to at least one of the student's career goals contribute to the overall relevance of the course to the student. The relevance of the course to the student is visualized with a thumb-up icon (1 icon means reasonable relevance and 3 means the highest relevance). Table 1 presents part of our algorithm for computing course relevance. For example, if a course is essentially relevant (relevance level of 5) in 2 of the student's career goals, the course will be considered highly relevant to the student. The complete set of rules consists of 16 cases. The current version of the algorithm is derived from our preliminary assumptions and needs to be evaluated with real users. The evaluation of this algorithm is part of our future work.

Table 1. The Algorithm for computation of course relevance

# of career goals with Relevance 5	# of career goals with Relevance 4	# of career goals with Relevance 3	Total Relevance
>=2	*	*	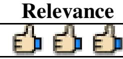
1	>1	*	
….			
0	1	0	
0	0	2	

3.3 Motivation for Providing Feedback

Similar to any other community-based adaptive system, the success of CourseAgent is highly dependent upon the feedback provided by the community. Course recommendation is a good example of a domain where community-based recommendation is useful while item-based recommendation [10] is not, since students are typically interested in taking courses that are *different* from those already taken, in order to learn the wide variety of knowledge that will be relevant for a career in this field. Moreover, unlike some community-based recommenders, such as MovieLens [9], recommendations that are provided to a specific student do not take into account her own ratings, but only the ratings of students who took potentially interesting courses earlier. As a result, ratings provided by the students in CourseAgent are beneficial solely to the community but not to the author of the ratings. This typical contradictory situation requires us to find some way for the system to encourage students to provide explicit feedback. As explained in the introduction, our goal has been to use a "do-it-for-yourself" approach. Therefore, our challenge has been to design an activity that is both attractive and meaningful for the students and can use course ratings provided by the student for the benefit of the author of the ratings. In our context, career planning looks like an attractive candidate for this kind of activity. To integrate career planning with student course evaluation, we developed the Career Scope interface, which is presented in this section.

In Career Scope, students can view the progress they have made towards each career goal. Courses they have taken and evaluated are used to compute their progress towards the career goal. The more relevant the course to the career goal, the more progress they will make towards the goal. Also, the difficulty level of the course will affect this rating. A low-load course would not necessarily cause the same progress as a high-load course. To visualize progress, we have assumed that a specific career goal can be "covered" by taking four relevant courses with medium level difficulty. More difficult courses with higher relevance contribute more to "covering" a career goal while courses with less relevance contribute less. To give more weight for courses taken earlier, we chose to use a logarithmic contribution function instead of a linear one. The current contribution function and all the parameters are considered to be pilot settings that will need to be validated with real users.

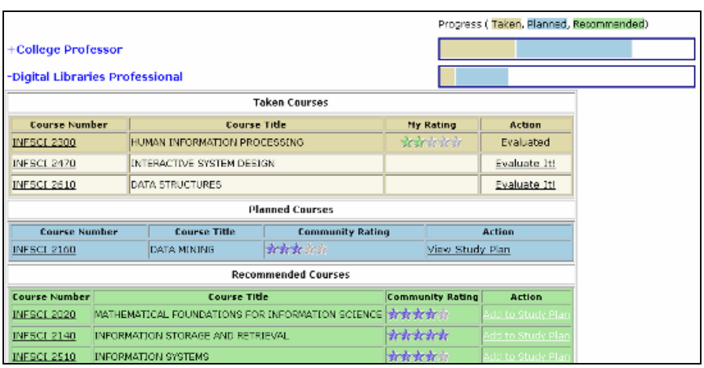

Fig. 4. Career Scope in CourseAgent

Fig. 4 shows a screenshot of the Career Scope section. For each specified career goal, the Career Scope section displays a progress bar that displays the contribution of relevant taken and planned courses towards achieving this goal. A taken course contributes to a career goal if the student rated it as being relevant to this goal. The amount of contribution depends upon the relevance and workload rating. The total contribution of the students' planned courses is computed from the average relevance and average difficulty level provided by the community of students. To distinguish actual progress from future progress, the contribution of planned courses is shown in the progress bar with a different color.

As shown in Fig. 4, the system lists three possible groups of courses for each career goal: taken, planned, and recommended. The students are able to see their own evaluation of taken courses. Taken but not evaluated courses are presented in the Taken Courses table with a lighter background. This prompts the students to evaluate the course (using the link to the right) in order to be count as a part of progress toward the career goal. Students can also re-evaluate the courses they have previously evaluated by clicking on the original rating. They are also able to view the community's evaluation of their planned courses, as rated by relevancy to each specific career goal. The list of recommended courses (based on the community's evaluation) is provided for each specific career goal and students are able to plan any of the recommended courses.

The design of Career Scope is based upon the assumption that the main goal of students is to take courses that will help them to find an interesting career in the future. By rating the relevance of courses, students are better able to take advantage of the system and observe their progress towards each of their career goals. This employs the methodology of "do-it-for-yourself" that is the main focus of our current work. By visualizing the contribution of planned courses to students' progress, we tried to encourage students to specify courses they plan to take. Specifying planned courses can then serve as implicit feedback for generating recommendations for the community. Social navigation support provided by the current version of the system does not take into account implicit feedback. As future work, we are planning to add implicit feedback into social navigation support.

4 Evaluation

We have completed the first study of the CourseAgent system at the School of Information Sciences in the University of Pittsburgh. The main goal of the study was to assess whether "do-it-for-yourself" approach increases student contribution to the system. To evaluate this hypothesis, we prepared two different versions of the system. The controlled version does not include the Career Scope screen that was designed to provide motivation to rate and plan courses. The rest of the system is exactly the same for both versions. The system was advertised to graduate students of the School of Information Sciences for two weeks before the registration deadline. When a student requested to use the system, they were randomly assigned to one of the two groups. For evaluation purposes, we logged all user interactions with the system.

We hypothesized that students in the control group would provide fewer evaluations and career interests, plan fewer courses to take in the future, and provide fewer taken courses. To evaluate our hypothesis, we looked at the average number of times that each group saved an evaluation, added a course, planned a course, and added a career interest. Table 2 presents the result. As shown in the table, the control group has planned fewer courses, added fewer career interests, and provided less evaluation. This means that the control group has provided less implicit and explicit feedback. However, the difference is not significant.

Table 2. Contribution of users in from the control and experimental group

	# of students	Average # of added courses	Average # of planned courses	Average # of added career interests	Average # of saved evaluations
Control	11	5	2	0.91	4.55
Experimental	9	5.89	5	2.22	6.22

For a deeper analysis, we looked at the usage of Career Scope by the experimental group. We observed that about half of the students in the experimental group did not actually use Career Scope. This might be due to interface problems such as the name of the section or the position of the section in the system. Also users might be lacking a good description of this part of the system. (We will investigate this issue as part of our future work.) As a result, for better analysis of the effect of Career Scope, we divided the users into 3 groups: control group, experimental group I who did not use Career Scope, and an experimental group II who used Career Scope. Table 3 presents the same result as Table 2 for these 3 groups. As shown in the table, the contribution of users from experimental group who did not actually use Career Scope is close to the contribution of users from the control group. The data shows that students who actually used Career Scope contributed significantly more to the system by providing more evaluations, planning more courses, and adding more career interests and taken courses. In all cases, the difference is statically significant (t-test, α=0.05).

Table 3. contribution of users with respect to usage of Career Scope

	# of students	Ave. # of added courses	Ave. # of planned courses	Ave. # of added career interests	Ave. # of saved evaluation
Control Group	11	5	2	0.91	4.55
Experimental group I	4	2.25	1.5	1.25	3.75
Control + Experimental I	15	4.27	1.87	1	4.33
Experimental group II	5	8.8	7	3	8.2

We were also interested in observing the activity patterns among these three groups. We looked at the fraction of providing feedback (explicit & implicit) compared to other actions, to measure the extent that the rating had been encouraging. The following graph presents the percentage of different types of activity among the three groups. The results suggest that the experimental group II, who received more encouragement for providing feedback, spent a higher fraction of their time on activities that would provide feedback to the system. This result is another indication that the encouragement caused by presenting career progress was beneficial to creating more feedback.

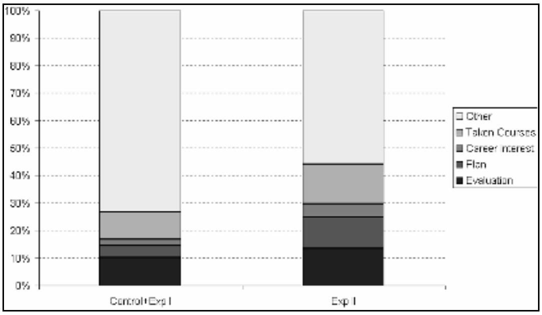

Fig. 5. Activity analysis of 3 groups

5 Discussion and Future Work

CourseAgent is a social navigation system that strives to automate "word of mouth" to help students making decision about courses to take [11]. Similar to any other community-based adaptive system, the success of CourseAgent is highly dependent on the feedback provided by the community. In CourseAgent we have tried to address the problem of under-contribution by employing a "do-it-for-yourself" approach and emphasizing the direct benefit of providing feedback. In CourseAgent, students are able to provide feedback in implicit and explicit ways. They can directly evaluate courses with respect to the relevance to each career goal as well as the difficulty level of the course. They are also providing implicit feedback when they plan or register for a course. Registering or planning a course represents an implicit interest in the course, which may be due to the relevance of the course to the students' career goals. The basic and obvious benefit of the system to the students is as a course management system that keeps information about courses they have taken and facilitates communication with their advisors. Providing social navigation support and community-based recommendation provides more benefit and encouragement to use the system. However, to encourage students to evaluate the courses they have taken, we have designed the Career Scope section of the system. Our results suggest that the "do-it-for-yourself" approach succeeds in providing more course recommendations. Observing progress toward each career goal is an important motivation for the students to use the system while also providing more explicit and implicit feedback to the system.

Currently, we are trying to advertise the use of this system among a larger number of students in the School of Information Sciences at the University of Pittsburgh, to validate our hypotheses with a larger number population. We have also designed a user study to conduct interviews and surveys. We plan to collect subjective feedback from students about the community-based support provided by CourseAgent. Using subjective data from real users, we will adjust our adaptation algorithm and different parameters used in the algorithms. We plan to modify the constant parameters in the algorithms (e.g. number of courses to cover a career goal) to variable parameters that are adjustable to students' goal and interests and specification of the area of the study. In the next version of the system we plan to improve the adaptation algorithm by taking into account the implicit feedback such as course planning. We hope that extending the development of this system and its evaluation will provide us with more ideas, in order to improve our approach for eliciting user feedback, an essential tool for building community-based adaptive hypermedia systems.

Acknowledgement

This research is supported by the National Science Foundation under Grant No. 0447083 and the National Science Foundation graduate fellowship. We would also like to thank Li-Chen Mao for developing the earlier version of the system.

References

1. Bretzke H. and J. Vassileva (2003). Motivating cooperation on peer to peer networks. Proceeding of 9th International Conference on User Modeling .
2. Brusilovsky, P. (2001). Adaptive hypermedia. User Modeling and User Adapted Interaction 11(1/2): 87-110.
3. Cheng, R. and Vassileva, J. (2005) Adaptive Reward Mechanism for Sustainable Online Learning Community. In Proceedings of 12th International Conference on Artificial Intelligence in Education, AIED'2005.
4. Claypool, M., Le, P., Waseda, M., and Brown D. (2001). Implicit interest indicators. In Proceedings of ACM Intelligent User Interfaces (IUI'01).
5. Dieberger, A., Dourish, P., Höök, K., Resnick, P., and Wexelblat, A (2000). Social navigation: Techniques for building more usable systems. Interactions, 7(6), 36-45.
6. Smyth, B., E. Balfe, Freyne, J., Briggs, P., Coyle, M., and Boydell, O. (2004). Exploiting Query Repetition and Regularity in an Adaptive Community-Based Web Search Engine. User Modeling & User-Adapted Interaction 14(5): 383-423.
7. Harper F. M., Li X., Chen Y., and Konstan J. (2005). An economic model of user rating in an online recommender system. In Ardissono L., Brna P., and Mitrovic A. (Eds.), Proceedins of 10th International Conference on User Modeling (UM 2005), Edinburgh, Scotland, UK.
8. Ling, K., Beenen, G., Ludford, P., Wang, X., Chang, K., Li, X., Cosley, D., Frankowski, D., Terveen, L., Rashid, A. M., Resnick, P., and Kraut, R. (2005). Using social psychology to motivate contributions to online communities. Journal of Computer-Mediated Communication, 10(4), article 10.
9. Miller, B., Albert, I., Lam, S.K., Konstan, J., and Riedl, J. (2003). MovieLens Unplugged: Experiences with a Recommender System on Four Mobile Devices, Proceedings of the 17th Annual Human-Computer Interaction Conference.
10. Sarwar, B., Karypis, G., Konstan, J., and Riedl J. (2001) Item-based Collaborative Filtering Recommendation Algorithms, In proceedings of 10th International World Wide Web conference.
11. Shardanand, U. and Maes, P. (1995). Social information filtering: Algorithms for automating 'word of mouth'. In Proceedings of Computer Human Interaction.