Supporting 'word-of-mouth' Social Networks via Collaborative Information Filtering

Mimi M. Recker

Utah State University

Andrew Walker

Lehigh University

September, 2002

Abstract

Altered Vista is an instructional system that supports a form of ‘contextual’ collaborative learning. Its design incorporates an information filtering technique, called collaborative information filtering, which, through computational and statistical means, leverages the work of individuals to benefit a group of users. Altered Vista is designed to provide, upon request, personalized recommendations of Web sites. It can also provide recommendations of like-minded people, thus setting the stage for future collaboration and communication. An empirical study involving in-service and pre-service teachers was conducted using Altered Vista and presents results from an empirical study. The study examined the feasibility and utility of automating the well-known social feature of propagating word-of-mouth opinions within educational settings. It also examined the impact of Altered Vista’s ability to recommend a social network of potentially unknown people.

Supporting ‘word-of-mouth’ Social Networks via Collaborative Information Filtering

Introduction

Much has been written about information technology support for intentional, extended, and intensive collaborative learning (e.g., Dillenbourg, 1999; Roschelle & Teasley, 1995). This form of collaboration falls into the category that Bruce (2001) has dubbed *conceptual* collaboration. Bruce (2001) calls another form of collaboration *contextual*, wherein individual participation is occasional and less intensive. In this case, the pursuit of personal goals by individuals creates incidental by-products, which also contribute to the common good.

This article describes a system, called Altered Vista, which was designed to support such contextual collaboration. Specifically, its design incorporates a recent information filtering technique, called collaborative information filtering which captures an individual’s preferences so as to benefit a group of users (Resnick & Varian, 1997). In its individual, intentional form, Altered Vista solicits ratings and opinions from users about the design, usefulness, and quality of web sites on particular topics. These data thus become a repository of community knowledge.

In its contextual form, the system can mine the data using collaborative filtering techniques in order to provide personalized recommendations of Web sites to an individual user. Because of this capability, Altered Vista is an example of what is called a ‘recommender system’ (Resnick & Varian, 1997). In addition, because of the underlying collaborative filtering algorithm, Altered Vista can also provide recommendations of like-minded people. Thus, Altered Vista sets the stage for future collaboration and communication.

The next sections describe collaborative information filtering, and discuss its implementation within the Altered Vista system. Results from an empirical study involving mostly in-service and pre-service teachers enrolled in classes at two U.S. universities are then reported. In particular, via analyses of user surveys, user comments in an online bulletin board, and system usage, the article reports the feasibility and utility of providing personalized recommendations. Specifically, the ability of the system to support and automate the well-known social feature of propagating word-of-mouth opinions from trusted people is examined.

Second, the article examines the feasibility and utility of Altered Vista's ability to recommend like-minded users. In particular, users' reported interest in people recommendation, related privacy issues, and the broader question of the role of a computer system in suggesting social networks where none previously existed are discussed.

Collaborative Information Filtering Systems

Within the human-computer interaction (HCI) literature, an approach to categorizing, collecting and filtering information has emerged, called collaborative information filtering. It is based on propagating word-of-mouth opinions and recommendations from trusted sources about the qualities of particular items (Malone, Grant, Turbak, Brobst, & Cohen, 1987; Maltz & Ehrlich, 1995; Shardanand & Maes, 1995). For example, you've arrived in a brand new city, and hunger pangs have erupted. How do you make that all-important decision: Where to dine? You might consult restaurant guides, newspapers, or the phone book. More likely, you would ask friends with similar tastes in cuisine to recommend their favorite spots. In the end, you want trusted sources to provide you with information about the quality of restaurants in order to help you make the best selection.

This solution to the ‘restaurant problem’ is the basic insight underlying research in collaborative information filtering. In general, collaborative filtering systems approach the problem of information filtering by estimating the desirability of items under consideration. These estimates are generally made, unlike content-indexing search engines, without any knowledge about the *content* of the items. Instead, desirability can be inferred *explicitly* by directly soliciting data from users. Typically this takes the form of a likert scale ranking or “vote” (Herlocker, Konstan, Borchers, & Riedl, 1999), but it may also involve anything from a binary “like/dislike” to detailed annotations (Hill, Stead, Rosenstein, & Furnas, 1995). Explicit data may also take the form of general user demographic information that is relevant to the domain. For example, in an education application, a user profile might specify the teacher’s subject areas and grade levels. The benefit of explicit data is its accuracy. Its difficulty lies in the effort and overhead required by the users in providing such data.

Desirability estimates can also be inferred *implicitly* by leveraging information collected for other purposes, usually as a by-product of user actions (Herlocker et al., 1999). For example, the system might infer that desirable items are used more frequently or more recently (Recker & Pitkow, 1996). An example in the domain of Usenet News articles is the time a user spends reading an article, which turns out to be a reliable way to infer user preference irrespective of article length (Konstan, Miller, Maltz, Herlocker, Gordon, & Riedl, 1997; Morita & Shinoda, 1994). In this way, the collaborative filtering approach attempts to generate inferences based on the “social” aspects of information, rather than simply its content (Brown & Duguid, 2001). Implicit data are collected more easily and at lower cost to the user, but inferences about item desirability are generally much less accurate than explicitly supplied ratings (Resnick & Varian, 1997).

Systems built on the collaborative filtering approach are also frequently called recommender systems because of their ability to provide recommendations of items to users (Resnick & Varian, 1997). Specifically, by statistical mining of users' preference data, recommender systems can automatically provide personalized recommendations to a particular user. Such systems have been implemented in a variety of domains, including recommending books, movies, research reports, and Usenet news articles (Resnick & Varian, 1997). More recently, recommender systems have become a staple element of e-commerce, as Internet vendors attempt to provide personalized recommendation of products to their customers. For example, the Internet vendor Amazon.com uses a recommender system to recommend products to its users. However, applications within education are much less common (for a review of the literature, see Walker, Recker, Lawless, & Wiley, 2002).

It is important to note that collaborative filtering systems are most useful in situations and domains with the following characteristics:

- The system can collect numerous data and metrics (e.g., ratings) about items, from many different users. Similarly, coverage (the proportion of items with data for estimating desirability) must be high. In general, the accuracy of the predictions made by recommender engines increases as the data pool for estimating item desirability (either explicitly or implicitly) also increases (Breese, Heckerman, & Kadie, 1998).

- The set of resources is better described by more subjective labels, such as opinions and taste, than objective labels, such as topic or keyword frequencies (Herlocker et al., 1999). Examples of such resources are jokes (Gupta, Digiovanni, Narita, & Goldberg, 1999) or political commentaries.

- Traditional information retrieval methods are less effective. This might be true in domains where content-indexing of resources as performed by traditional search engines) is impractical or difficult (e.g., multimedia items).

There are several different approaches to collaborative information filtering. While the specific techniques vary, all of them utilize the following steps:

1. **Data gathering.** Collaborative filtering depends critically on gathering information about the items under consideration and the people who use them. The more information known about people, and their preferences for various items, the more accurate the system's predictions will be. Through interacting with the system, a user builds a profile of his/her preferences by supplying opinions about the quality of different items. As previously noted, these opinions may be explicitly collected and/or implicitly inferred. Typically, these data constitute a detailed level of information about users, which raises difficult privacy concerns. Some users may be reluctant to provide such detailed personal information. In addition, it can be difficult to motivate users to contribute necessary preference data (Avery & Zeckhauser, 1997). In the end, users must perceive a reward for their efforts, either through receiving high-quality recommendations, or appropriate incentives (Avery & Zeckhauser, 1997; Swearingen & Sinha, 2001).

2. **Prediction and recommendation**. In the case of prediction, systems respond to a user's request to predict how much they would like a specific item. The systems may also recommend a set of items to the user (Karypis, 2000). This usually consists of a list of the items with the highest predicted value. Alternatively, the collaborative filtering algorithm may perform both tasks. At the heart of deriving these predictions and recommendations is the algorithm driving the filter. The Altered Vista system, described below, relies on a class of algorithms called neighborhood-based (Herlocker et al., 1999). This approach is primarily concerned with determining similarities in preferences between users, and splits the prediction/recommendation task into two distinct parts.

a. Neighborhood identification. The collaborative filtering system identifies for each user, other users with similar profiles. This is called the active user's *neighborhood*. User similarity is often computed by correlating users on the basis of their ratings data; users with high correlation are placed in the same neighborhood. If the system recommends people as well as resources, then the set of recommended people will come from this neighborhood.

b. Prediction/Recommendation. Once the neighborhood has been formed, predictions can be made on a set of items which the user supplies by using some form of a weighted average of all the preference data provided by neighborhood members. Alternatively, predictions can be made for all items unseen by the active user, and items with high predicted ratings are presented as recommendations.

3.**Algorithm Evaluation.** As an ancillary step, the algorithm's speed, coverage (how many predictions an algorithm is able to make with the available data), and accuracy are evaluated. It can be beneficial to pass these evaluations on to users to help them assess the quality of the predictions. In addition, there appears to be benefits in providing explanations of predictions or recommendations to users. If users do not know how a recommendation or prediction is made then they will not know what level of confidence to place in the suggestion (Herlocker, Konstan, & Riedl, 2000). However, exactly how these should be described and displayed to end-users is still an area of active research.

Altered Vista: System Description

In its current implementation, Altered Vista is specifically aimed at teachers and students who use Web resources in education. Using Altered Vista, users submit reviews about the design, quality, and usefulness of Web resources for online education. These ratings become part of the recommendation database. Users can then access and search the recommendations of other users. The user can also request personalized recommendations from the system. In this way, a user is able to leverage the opinions of others in order to locate relevant, quality information, while avoiding less useful sites. An additional benefit of this approach is that it allows a user to locate other users (e.g., students or instructors) who share similar interests for further communication and collaboration.

Design Considerations

When developing a collaborative filtering system that gathers explicit user opinions, several design dimensions must be considered. These are 1) the ontology of the review or rating scheme, 2) how user data are collected, 3) how user data are aggregated, 4) how user data are used, and 5) the level of user anonymity (Resnick & Varian, 1997). Each of these dimensions is discussed below in terms of the design of the Altered Vista system.

Review Scheme. A fundamental issue in the design of a collaborative filtering system is defining the data users will supply to infer their preferences for resources in the domain. Together, these data comprise what is typically called a review or rating scheme.

The review scheme in Altered Vista is specific to the domain under consideration. Table 1 shows the current scheme for one domain (online education) implemented within Altered Vista. This review scheme consists solely of explicitly collected preference data. It was derived and refined after several iterations of testing the scheme with a variety of professional educators and researchers in educational technology. In particular, each group of professionals was asked to comment on the current version of the scheme, and its utility in terms of the kinds of preference data collected. After each test, the review scheme was revised prior to its presentation to the next group.

Collection of Review Data*:* Altered Vista relies on explicit, active collection of preference information from users as defined in the review scheme. To enter their review data, users interact with a series of interface elements, including Likert scales, text entry boxes, and multiple selection lists.

Aggregation*.* Once a rating is complete, the user submits the review form and all values are stored in a database. This database of aggregated reviews becomes a mechanism that supports search and automated recommendation of resources.

Usage: Searching. Because the rating scheme is searchable, it provides an alternate to content-indexing for discovering resources of interest. A user can display all reviews, search by keyword, or search for reviews by a specific contributor.

Table 1

Review scheme for the domain of "online education"

Name	Description	Format
Web Site Title	The title of the site	Text box
Internet Address	The URL of the site	Text box
Keyword(s)		multiple selection list
Added by	User email	automatically generated
Overall Rating		5 point likert scale
Navigation Ease	How easy is it to get around the site?	5 point likert scale
Accuracy of Information	Is the information on the web site correct?	5 point likert scale
Educational Relevance	How useful the site is for educators or their students.	5 point likert scale
Description	Any information not represented in the other review criteria, as well as justification for any extremes.	text box
Grade Level	What is the target audience for this site?	multiple selection list
Would you use this web site while teaching?	Can you picture yourself using this web site as part of your own instruction?	5 point likert scale

Usage: Recommendation. Upon user request, the aggregated database of user reviews can be analyzed to provide automated, personalized recommendations. As previously noted, the recommendation algorithm relies upon a specific implementation of the neighborhood-based approach to collaborative filtering. Such algorithms contain a number of parameters that the designer must set during implementation. For example, the designer must decide what counts as a threshold correlation between user ratings when defining that two users “agree.” As described below, parameter values were set using guidance from Herlocker et al. (1999).

When a user requests recommendations, Altered Vista first determines the *neighborhood* for the current, active user. In a pair-wise fashion, the overall rating for resources provided by the active user is correlated with all other users. To be considered, users must have mutually reviewed at least two resources and have a correlation of at least 0.5. This set comprises the active user’s neighborhood. The thresholds for the number of overlapping reviews and correlation level were pre-determined to result in an approximation of the ideal neighborhood size, which, as defined by Herlocker et al. (1999), is about ten.

Resources rated highly by users within the neighborhood but unseen by the active user form the basis for automated recommendations. The system calculates a predicted rating for the unseen resource for the active user. This predicted rating is a weighted average of the ratings of users in the neighborhood, based on their correlation level with the active user. The current system only recommends resources with a predicted rating greater than or equal to 4.0 (on a 5-point scale). Using 4.0 or greater as the definition of high rating was again based on research by Herlocker et al. (1999), who discussed user consumption decisions in terms of “signal” and “noise” on a five-point scale. They defined “noise” or poor resources as those rated less than 4.0, while “signal” was defined as those rated greater than 4.0.

Members of the active user’s neighborhood can also be recommended. In this way, the active user can locate other users that share similar interests for further communication and collaboration.

Identity of Contributors. To maximize the value of contributed information, Altered Vista was designed to make user identity salient. Hence, users must log-in prior to using the system, and the email address of the author of particular ratings is both a searchable item and available for inspection within search results.

System Interactions

To access Altered Vista, users log into the system, and select the currently implemented area, online education, in which they will contribute reviews for particular Web resources. Figure 1 shows an example screen shot for entering a review. As can be seen, on one side of the screen, the user views the target Web site, while on the other side of the screen, the review of the site is entered using the pre-defined review scheme described above.

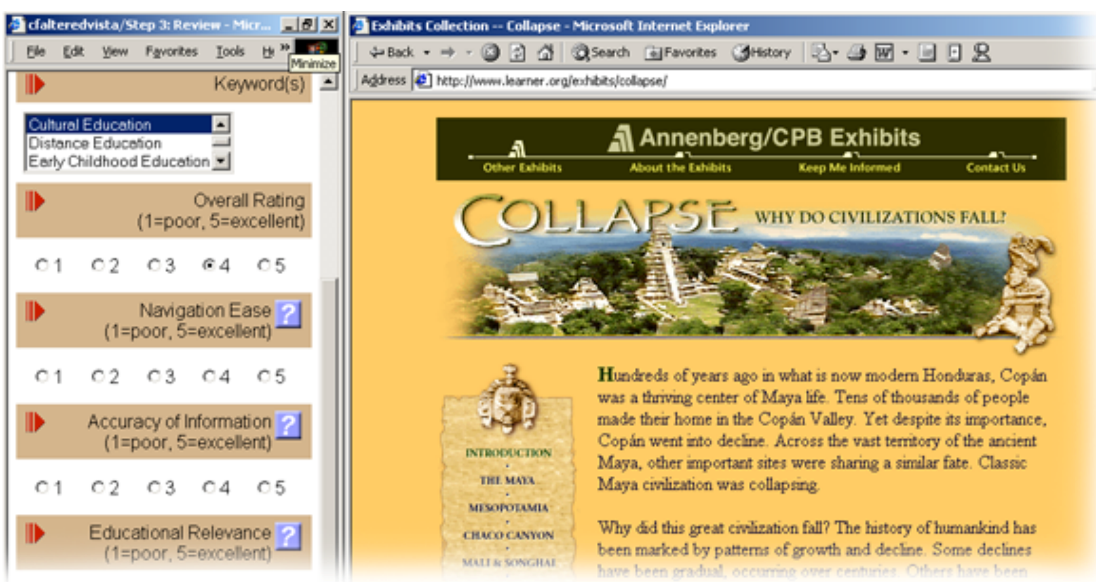

Figure 1. Adding a review to Altered Vista.

These reviews are then stored in the Altered Vista database. Users can then search the reviews submitted by other users. Alternatively, as previously described, they can request personalized recommendations of unseen Web resources. Figure 2 shows a screen shot of a composite review, based upon several user ratings of one resource.

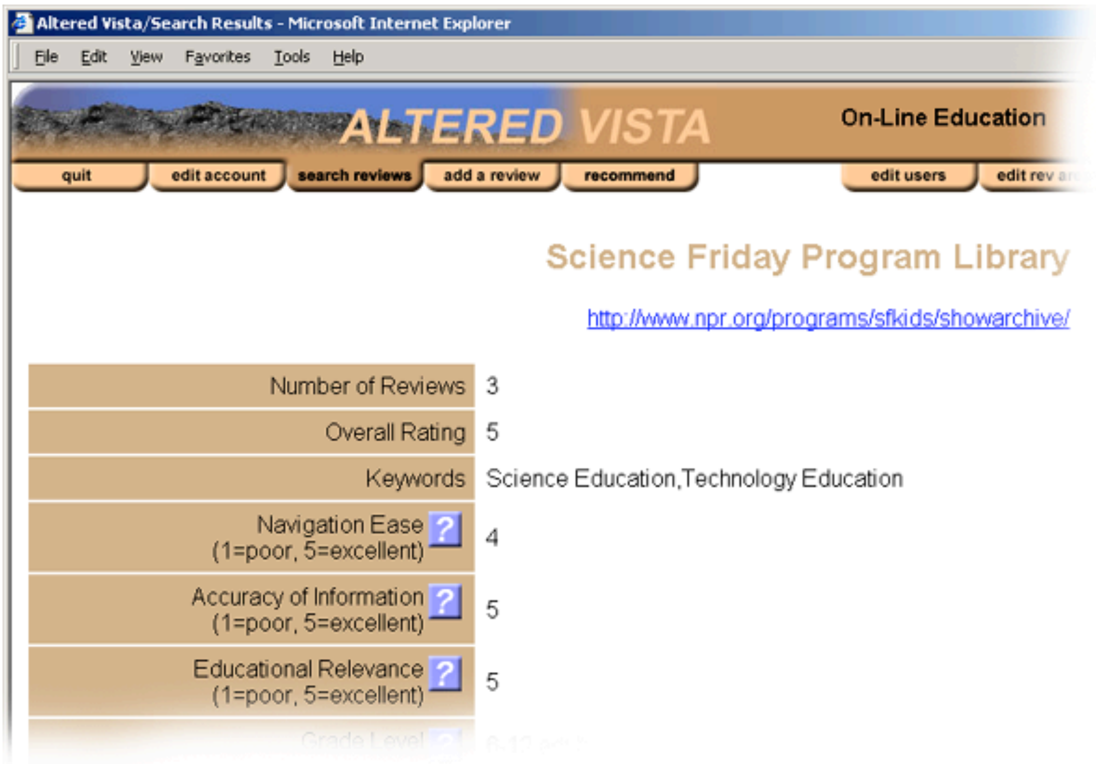

Figure 2. A screen shot showing a composite review from Altered Vista.

System Specifications

Altered Vista is implemented on a Linux machine, running the Apache Web server. Reviews are stored in database, and communication between it and the server is accomplished using PHP. Users may access the system using any browser supporting Javascript (or VB Script) and Cascading Style Sheets.

Empirical Evaluation

A 3-month trial involving 63 participants was conducted using Altered Vista. Walker et al. (2002) reported results concerning system usability and algorithm performance. The present analysis focuses on two questions:

1. Can Altered Vista automate and support the well-known social feature of propagating word-of-mouth opinions? Specifically, did participants find automated, personalized recommendations of resources useful?
2. To what extent does reviewing and receiving recommendations of Web resources within a community of users support and promote collaborative and community-building activities? Specifically, did participants find personalized recommendations of people useful?

The next section describes the study's methods and participants. Then, analyses of system usage, user questionnaires, and participants' comments in an online bulletin board are presented.

Participants and Methods

Sixty-three students (41% male and 59% female) from two universities in the United States participated in the study as part of course credit. As shown in Table 2, most participants comprised a mix of current classroom teachers taking additional professional development classes, and students preparing to become teachers.

In the context of the educational technology courses in which they were enrolled at their respective institutions, students were asked to use Altered Vista to review web resources related to the domain of 'online education'. Initially, students were asked to review five sites from a pre-selected list of Web resources. An expert in online learning, who had taught numerous classes on the design and evaluation of Web-based educational sites, selected these sites. The list of sites was drawn from the expert's teaching experience and was intended to represent a broad, cross section of the type of resources that teachers would typically encounter. As such, they were intended to run the gamut of quality in terms of content, design, and overall utility.

Table 2

Participant background

Participant descriptor	Frequency
In-service teachers	22 (35%)
Pre-service teachers	19 (30%)
Religious education	10 (16%)
Other	9 (14%)
University instructor	3 (5%)

Participants were also asked to review five sites of their own choice and related to the domain. This means that participants found Web sites that they wanted to review, then added them to the database along with their review information. Finally, they were asked to review five sites reviewed by other users in the Altered Vista database. Thus, at a minimum, they were asked to contribute fifteen reviews during the course of the trial evaluation period. The goal was to ensure a critical mass of overlapping reviews in order to provide data to the recommender algorithm.

Prior to using Altered Vista, all participants completed an online questionnaire, which asked basic demographic information. At the end of the trial, students completed an exit survey that asked participants to rate the usability, usefulness, and accuracy of Altered Vista. Fifty-two (82%) of the participants completed this exit survey. The surveys consisted of 11 5-point likert scale (1=strongly disagree; 5=strongly agree) and two short answer questions. Open-ended comments were also collected from an online course bulletin board used by participants.

Usage Results

Table 3

Usage results

Total number of participants	63
Total number of resources reviewed	242
Total number of reviews submitted	934
Mean number of reviews per resource (SD)	3.9 (7.2)
Mean number of reviews submitted per user (SD)	14.8 (2.3)

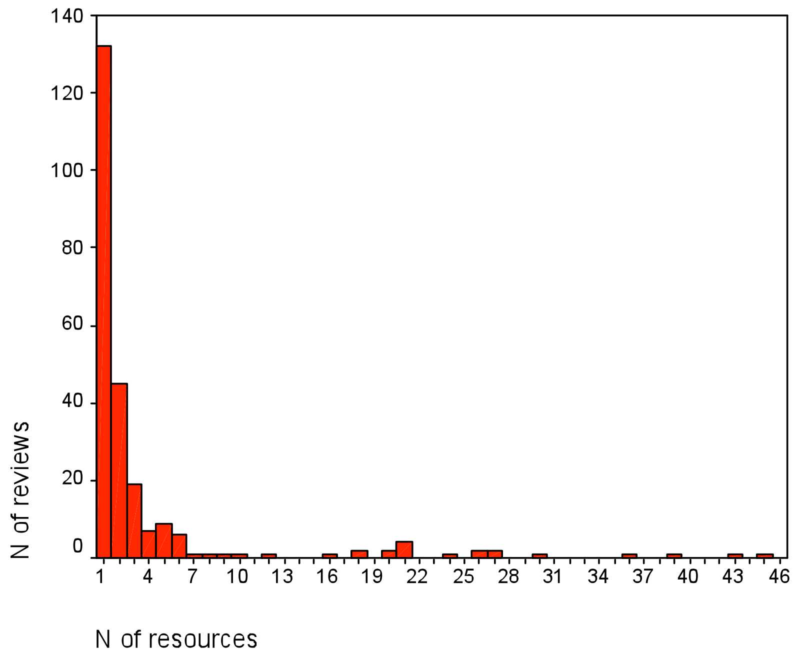

Figure 3. Frequency graph of number of reviews per resource.

As shown in Table 3, almost 1000 reviews were submitted for over 240 unique Web resources. Resources received a mean number of 3.9 reviews, but their frequency distribution is skewed. Figure 3 shows that a handful of resources received a large number of reviews, while most resources had a small number of reviews.

As previously described, the recommender algorithm employed by Altered Vista relies upon a neighborhood-based method (Herlocker et al., 1999). As shown in Table 4, users received a mean number of approximately 46 recommended resources and 16 recommended people.

Table 4

Performance of the recommender engine

Mean number of recommended resources per user (SD)	46.5 (28.0)
Minimum-Maximum number of recommended resources	0-76
Mean number of recommended people per user (SD)	16.5 (6.8)
Minimum-Maximum number of recommended people	0-31

Table 5

Summary results from exit survey

Results from exit survey (Likert scale: 1= strongly disagree; 5=strongly agree)	**% rating 4 or 5**	**Mean**	**SD**
A. AV is a useful tool for finding quality resources.	87	4.2	0.9
B. AV provided me with useful recommendations of resources.	65	3.8	1.0
C. AV helped me find resources that I would otherwise not have found.	73	3.8	1.0
D. I would use AV even if it weren't a course requirement.	45	3.3	1.1
E. AV is a useful tool for finding people with shared opinions.	74	3.9	0.8
F. AV provided me with useful recommendations of people with similar opinions.	54	3.5	1.0
G. AV allows me to find and communicate with other professionals in my field to whom I would not normally have access.	54	3.5	1.0
H. AV allowed me to see opinions about the quality of resources from people with different expertise.	36	3.5	1.0

Automating 'word-of-mouth'

As indicated by survey results, most respondents reported that Altered Vista was a useful tool for finding new and quality resources (A and C in Table 5). In addition, many respondents (65%) appeared to like the personalized resource recommendation (B in Table 5).

As noted by one participant:

> It takes a lot of time to evaluate websites and if there were a place where teachers could go to see evaluations already completed (and have a list of other sites that they may be interested in) it would save a lot of time in the long run. In this way teachers will find encouragement and resources that will help them integrate technology into their curriculum without having to reinvent the wheel. If teachers had a place to share their impressions about sites they had looked at, and all teachers had access to the data, just think of the work and time that could be saved. Especially if the data was searchable by grade level and subject.

Another participant saw the value of the recommender function, despite the fact that the set of reviewed sites laid outside of her field of interest:

> Regarding the recommender function in Altered Vista, I think this is an excellent function. Altered Vista is not as "real" to me as other sites since I teach at the university level and many of the sites available for critiquing are geared to younger audiences. However, if the choice of web sites were more numerous and more developed, the tool would be excellent. I would

> appreciate having a program such as this to sort through the endless possibilities on the Internet.

However, the recommender algorithm, is not always optimal, as discovered by this participant:

> I went to the Altered Vista recommend link and found eleven people that had similar preferences in the area of Education. Based on the reviews of the people that I tend to agree with, Altered Vista recommended four pages of sites. Browsing through the sites leads me to believe that I must have erroneously rated them because most of the sites on the recommended pages don't appear to be ones that I would use.

Unlike word-of-mouth opinions, which propagate naturally through social networks, Altered Vista usage is not deeply embedded within an offline group or informal social experiences. As a consequence, slightly less that a half of the respondents indicated that they would use the system if it weren't a course requirement (D in Table 5). As best described by one participant:

> Reviewing web sights (sic) is not something I would do without some kind of motivation.

This result highlights difficult issues relating to people's motivation for sustained use of the system. Clearly, most participants saw value in receiving recommendations. However, most also noted that they required incentives (in this case, course credit) to provide their reviews. This issue of incentives has been investigated within the collaborative filtering literature (e.g., Avery & Zeckhauser, 1997; Swearingen & Sinha, 2001). Overall, these researchers conclude that if benefits of using the system (high quality recommendations) do not clearly outweigh costs of participating (providing reviews), then incentives are required for users to participate. Results from the present study suggest that participants believed that the cost of participating outweighed the value of receiving personalized recommendations.

Automating Social Networks

This section examines the extent that the 'people' recommender function can support the formation of social networks by analyzing survey results and participants' comments. The designers of Altered Vista wrestled between supporting user privacy and promoting social interaction. Currently, Altered Vista only lists the email addresses of participants, allowing for a somewhat high level of privacy.

In the exit survey, three quarters of the respondents indicated that the system allowed them to find people with similar opinions (E in Table 5). However, just over half of the respondents liked the "people" recommender function (F), and saw the value of using Altered Vista to find and communicate with other users (G & H in Table 5).

One respondent clearly saw value in the people recommender:

> This is a COOL feature!! Like [person x] mentioned, What a time saver this could be if all the teachers could have access to this kind of a system. What was fascinating to see how the top few people on my list responded almost exactly as I had done. Knowing this kind of a trend, I could then search through the sites they rated high in order to find some thing of interest to me, with very few exceptions.

It seems clear that participants grappled with the notion of having people recommended to them. For example, in the exit survey, participants were asked to list people that were recommended to them. Of those that responded, half provided the names (not the emails) of people. Thus, even though Altered Vista only identifies users via their email addresses, some users were able to recognize known peers. One user was clearly able to recognize his fellow religious educators (a small subset of the total group):

> … I for one enjoy the recommendations. I noticed a few Seminary Teachers having the same interest as me in our group.

Other participants did not like the pseudo-anonymity of people. Instead, they wanted to know more about the person behind recommendations:

> … Usually recommendations are more valuable if the credentials of the recommender are known. Is there a way (besides guessing from what they say) to display expertise level of the recommender?

> … I am not impressed with the fact that so-and-so and I have the same predicted evaluation of a web site.

> … Siskel (okay, the new guy!) and Ebert at least have a following and have reputation and a rating system to uphold. Although something like this may work in the future but the database must be huge.

One participant hit upon an aspect of Altered Vista without a strong parallel in the social world. This participant was interested in knowing more about dissimilar people (something Altered Vista could easily report), though the immediate usefulness of this feature is unclear:

> At first I thought the recommender function of Altered Vista was great, and in the final analysis I still do. There was only one thing that kind of bothered me. It allowed me to see the people who had similar preferences to me in the area of online education. However, it made me wonder what the preferences were of the people that were not similar to me.

Finally, a participant directly addressed one of the study's key concerns – that of privacy:

> I found that the recommender listed 18 email addresses of people with my common ratings. It was interesting to see what others had researched, but I don't know if I would agree to having this information widely available on the web - would this be an additional open invitation for the invasion of my privacy - if there is such a thing on the web?

Conclusion

This article described a system that applied collaborative information filtering techniques in an instructional setting. Through its focus on contextual collaboration, Altered Vista attempts to leverage the work of many people to capture and propagate the opinions of its user community. It also contains features to support future communication and collaboration – and the establishment of social networks – by recommending like-minded users.

Results from the empirical study suggest that participants found Altered Vista a useful tool for finding quality resources and like-minded people. Its role in fostering community-building and collaboration is much less clear. Certainly, simply listing 'like-minded' users for a person does not guarantee that collaborations will occur. Instead, it seems that such collaborations must occur in the context of a larger goal or activity. Moreover, because of their relative novelty, users have had little previous exposure to systems that automatically recommend potential collaborators. As such, the results of this study are inconclusive in terms of the ability of recommender systems to support community building within the wider Internet.

This research also raises a number of important issues concerning the use of collaborative filtering in education, which are worthy of further study. First, results suggest that while most of the study participants saw great value in receiving personalized recommendations, they also needed incentives to provide reviews. In the present study, course credit was their incentive. As such, users wanted to 'free ride' on the work of others (Avery & Zeckhauser, 1997), but were reluctant to contribute their own efforts. However, if users do not provide reviews, it is difficult to seed and grow a review database, which impacts the system's reliability when recommending resources (Konstan et al., 1997). This is especially true if the user is an early contributor of review information and wishes to receive recommendations (Avery & Zeckhauser, 1997).

As previously noted, future studies must pay closer attention to the way Altered Vista is integrated into participants' routine use of the Web. For example, Amazon.com solicits user reviews as customers browse and shop for products. These reviews form the basis for recommendations. Similarly, future studies must design activities in which participants rate web sites as part of larger learning activities. In this way, users would be motivated to contribute reviews in meaningful and sustained ways.

Second, use of the system raises concerns surrounding user privacy in online environments. Specifically, it remains unclear if anonymity of participation (or even pseudo-anonymity via a proxy) impacts user acceptance and trust of the system and the recommendations it provides.

Finally, while great potential lies in the application of collaborative filtering in education, care must be taken in selecting appropriate domains. In particular, bounds must be established on the range of resources to be filtered. The unconstrained World Wide Web, because of its unlimited, heterogeneous, and ever-changing nature, is less ideal. Instead, collaborative filtering is more suitably applied in a bounded environment (for example, the domain of books and movies). Indeed, current research is applying the approach within Internet-based digital libraries of educational resources (Recker & Wiley, 2001). Ultimately, this may prove to be a more suitable domain, because items in a digital library are more stable, easily itemized, and indexed.

References

Avery, C., & Zeckhauser, R. (1997). Recommender systems for evaluating computer messages. *Communications of the ACM, 40*(3*),* 88-89.

Breese, J., Heckerman, D., & Kadie, C. (1998). Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, 43-52.

Brown, J.S., & Duguid, P. (2000). *The social life of information*. Cambridge, MA: Harvard Business School Press.

Bruce, B. C. (2001, April). *Collaboratories for elearning: How virtual and physical spaces index modes of working and learning*. Paper presented at the American Educational Research Association Annual Meeting, Seattle, WA.

Dillenbourg, P. (Ed.) (1999). *Collaborative learning: Cognitive and computational approaches.* Oxford, UK: Elsevier.

Gupta, D., Digiovanni, M., Narita, H. & Goldberg, K. (1999). Jester 2.0: Evaluation of an new linear time collaborative filtering algorithm. *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, 291-292

Herlocker, J., Konstan, J., Borchers, A., & Riedl, J. (1999). An algorithmic framework for performing collaborative filtering. In *Proceedings of SIGIR'99*, 230-237.

Herlocker, J., Konstan, J., & Riedl, J. (2000). Explaining collaborative filtering recommendations. *Proceedings of the ACM 2000 Conference on Computer Supported Cooperative Work*, 241-250

Hill, W., Stead, L., Rosenstein, M. & Furnas, G. (1995) Recommending and evaluating choices in a virtual community of use. *Proceedings on Human Factors in Computing Systems*, 194-201

Konstan, J., Miller, B., Maltz, D., Herlocker, J., Gordon, L., & Riedl, J. (1997). GroupLens. *Communications of the ACM*, *40*(3*),* 77-87.

Malone, T., Grant, K., Turbak, F., Brobst, S., & Cohen, M. (1987). Intelligent information sharing systems. *Communications of the ACM*, *30*(5*),* 390-402.

Maltz, D. & Ehrlich, K. (1995). Pointing the way: Active collaborative filtering. *ACM Conference on Human Factors in Computing Systems*, 202-209.

Morita, M. & Shinoda, Y. (1994). Information filtering based on user behavior analysis and best match text retrieval. *Proceedings of the Seventeenth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 230-237.

Recker, M. & Pitkow, J. (1996). Predicting Document Access in Large, Multimedia Repositories. *ACM Transactions on Computer-Human Interaction (ToCHI)*, *3*(4*)*, 352-375.

Recker, M. & Wiley, D. (2001). A non-authoritative educational metadata ontology for filtering and recommending learning objects. *Interactive Learning Environments. 1*, 1-17.

Resnick, P. & Varian, H. (1997). Recommender systems, Special Issue. *Communications of the ACM*, *40*(3*),* 56-58.

Roschelle, J. & Teasley, S. (1995). The construction of shared knowledge in collaborative problem solving. In C. O'Malley (Ed.), *Computer supported collaborative learning* (pp. 69-197). Berlin: Springer Verlag.

Shardanand, U. & Maes, P. (1995). Social information filtering: Algorithms for automating word-of-mouth. *ACM Conference on Human Factors in Computing Systems*, 210-215.

Swearingen, K. & Sinha, R. (2001). Beyond Algorithms, An HCI perspective on Recommender Systems. In *Proceedings of the 2001 Workshop on Recommender Systems.*

Walker, A., Recker, M., Lawless, K. & Wiley, D. (2002). *Collaborative information filtering: a review and an educational application*. Manuscript submitted for publication.

Author Note

This work reported was partially supported by grants from Utah State University and the National Science Foundation (NSF DUE-0085855). We thank the students who participated in our study. We also thank Richard Cutler, Jen Walker, and an anonymous reviewer for helpful advice.