

ASSIST: Adaptive Social Support for Information Space Traversal

Rosta Farzan
University of Pittsburgh
Intelligent Systems Program
Pittsburgh, PA, 15260, USA
rosta@cs.pitt.edu

Maurice Coyle
University College Dublin
School of Computer Science &
Informatics
Belfield, Dublin 4, Ireland
maurice.coyle@ucd.ie

Jill Freyne
University College Dublin
School of Computer Science &
Informatics
Belfield, Dublin 4, Ireland
jill.freyne@ucd.ie

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

Barry Smyth
University College Dublin
School of Computer Science & Informatics
Belfield, Dublin 4, Ireland
barry.smyth@ucd.ie

ABSTRACT

Finding relevant information in a hyperspace has been a much studied problem for many years. With the emergence of so called Web 2.0 technologies we have seen the use of social systems for retrieval tasks increasing dramatically. Each system collects and exploits its own pool of community wisdom for the benefit of its users. In this paper we suggest a form of retrieval which exploits the pools of wisdom of multiple social technologies, specifically social search and social navigation. The paper details the added user benefits of merging several sources of social wisdom. We present details of the ASSIST engine developed to integrate social support mechanisms for the users of information repositories. The goal of this paper is to present the main features of the integrated community-based personalization engine that we have developed in order to improve retrieval in the hyperspace of information resources. It also reports the results of an empirical study of this technology.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval – *Search process, Selection Process*; H.5.1 [Information Interfaces and Presentation]: Hypertext/Hypermedia – *Navigation*

General Terms

Design, Experimentation, Human Factors

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Hypertext'07, Sep 10–12, 2004, Manchester, UK.

Copyright 2004 ACM 1-58113-000-0/00/0004...\$5.00.

Keywords

Social navigation, social search, classroom study

1. INTRODUCTION

The emergence of several social Web technologies is commonly considered a major stage in the evolution of the WWW, as stressed by the popular nickname Web 2.0. The technologies that are typically listed as examples of Web 2.0 – blogs, Wikis and various tagged collections – stress the role of Web users as creators of information. These technologies harness the power of the user community to create and structure (with explicit or implicit links) new information artifacts. However, once recognized, the power of the community can be effectively used in a range of other contexts. This paper explores new ways of harnessing the community power in a context that is more traditional for the Hypertext community – finding relevant information in a hyperspace of information resources.

Consider the users of a large information repository, such as the large collection of academic papers provided in the ACM Digital Library (<http://portal.acm.org/>) or the collection of movie-related information provided by the IMDB (<http://www.imdb.com/>). These repositories support state-of-the art search and browsing facilities, yet they leave users on their own when it comes to finding the most appropriate article or movie. In contrast, while traditional repositories such as libraries and movie rental stores may provide relatively poor search and browsing facilities, they allowed their users to rely on community wisdom in this process. Users frequently look for interesting items in the stacks of books and movies recently returned by others. When browsing shelves with books and movies they pay attention to the “wear and tear” of the cover. The fact that a book or movie was retrieved or used by other community members in the past provides an additional source of information, which is especially reliable in small communities where user interests are relatively similar.

An interesting research issue is how can we combine the power of the search and browsing facilities available in modern information spaces with the power of the community that is available in more traditional repositories? While this issue has been explored in the past, most of the earlier work (i.e., research on social navigation and social search) focused on social enhancement of specific information access technologies. In contrast, this paper takes a user- and application- centered approach. Since our earlier studies indicated that both search-based and browsing-based access to information resources are important for the users of information repositories [4], the teams at University College Dublin and the University of Pittsburgh joined efforts to develop integrated social support mechanisms for the users of information repositories. Starting from the complementary work carried out by our research groups in the areas of personalized search and navigation, we developed ASSIST (Adaptive Social Support for Information Space Traversal), an approach to harness the power of the community to support and complement both search and browsing in modern information spaces. The goal of this paper is to present the main ideas of ASSIST, to introduce ASSIST-ACM – an integrated community-based personalization system that was developed for the users of the ACM Digital Library, and report the results of an empirical study of this system. A detailed presentation of community-based adaptation mechanisms is provided in an earlier published paper [7].

The remainder of this paper is structured as follows. In section 2 we review some related work on social information access and provide a summary of previous work on the integration of search and browsing. In Section 3, we describe the comprehensive system that we have developed to support community-based personalization in DLs. Section 4 describes the evaluation of the system in detail and highlights a number of interesting outcomes and results. Section 5 concludes with a discussion of directions for future research in this area.

2. RELATED WORK

2.1 Social Information Access

Over the last decade, a range of information systems have looked at how community-based personalization can help users to find their way in the rapidly expanding information space. Community-based personalization is mainly offered through social information access technologies. These technologies capitalize on the natural tendency of people to follow direct and indirect cues about the activities of others; for example, we often prefer restaurants that appear to attract lots of customers and our movie preferences are often informed by the opinions of others.

The pioneering work on social information access, from the early 90's, attempted to formalize this social phenomenon in two ways: collaborative filtering and history-enriched environments. Collaborative filtering attempts to propagate information items between users with similar interests. This technology enabled a social form of information filtering and recommendation. For example, the pioneering collaborative filtering system GroupLens [16] allowed cross-recommendation of NetNews articles. In contrast, history-enriched environments attempt to make the aggregated or individual actions of community users visible to others, mainly to facilitate social navigation through a given information space. For example, the 'Read Wear and Edit Wear' system in [9] visualized the interaction history of a document's

authors and readers, enabling third parties to quickly locate the most viewed or edited parts of a document. Knowledge Sea II [3] is a more elaborate example of a history-enriched environment. It was developed to help students of the same class discover the most useful pages in multiple open corpus textbooks. It supports information access through search, visualization, and browsing and guides the users with extended footprint-based and annotation-based social navigation support.

More recently, the set of available social information access technologies has been extended by social search and social bookmarking systems. Social search systems such as I-SPY [18], attempt to help new searchers by harnessing past successful searches by similar users, promoting results that proved successful during these sessions. The important challenge of remembering and retrieving items found in large information spaces is the basis of social bookmarking systems. Social bookmarking systems allow users to store individual collections of bookmarks. The collections are centrally stored and as a result the users can access their bookmarks from anywhere. Moreover, the users are able to share complete or partial sets of their personal bookmark collection with other users. Therefore, they are able to benefit from information discovered by others. Del.icio.us (<http://del.icio.us>) was one of the pioneers of social bookmarking. Dogear is a social bookmarking service designed for helping users with information discovery in a large enterprise. Millen et al present how a large enterprise can also benefit from a social bookmarking service. Moreover, they discuss how Dogear the social bookmarking system can be integrated with other corporate collaborative applications [15].

Traditional classification is too rigid and expensive to address users' need in social bookmarking systems. In contrast, tagging is an informal, bottom-up approach to classification, which allows users to organize resources without any pre-determined classifications. Users describe each resource with a set of *tags* that they find relevant to the resource. A *collective* user classification emerges when large numbers of people tag large amounts of content. Systems such as WebTagger [12] apply tagging to help new users locate useful information already discovered and classified by others. Information access using social tagging systems has recently been popularized by such systems as del.icio.us and Flickr (<http://flickr.com>).

Social networking tools are another example of the application of social technologies that help people connect to others with similar interests. Social networking tools such as LinkedIn and Orkut have successfully attracted many users. Finally, Golbeck and Wasser have tried integrating browsing with social networking by highlighting social information in movies using social networking tools [8].

2.2 Integrating Browsing and Search

At a fundamental level, the goals of browsing and searching are the same: to locate information and other resources online. Thus, it seems intuitive to attempt to integrate these two related but functionally different tasks so that the user experience for each may be augmented.

The work of [6] describes a system known as SearchGuide that seeks to integrate browsing and searching paradigms in the field of community-based Web search. Briefly, Web search is enhanced by using browsing information to supplement search

results. For example, if users frequently use a given search result as a starting point to browse to some other page, then this other page may be promoted in future result-lists. Similarly, search information can be used to augment the content of pages that users browse to from a result-list.

The work of [17] uses a formal perspective to blur the distinction between searching and browsing, resulting in the development of a system that allows a user to cross the boundaries between the two activities more easily and intuitively. In this work, we investigate some of the core concepts contained in [17], by recognizing and enhancing the potential synergy between browsing and searching in a digital library. We allow users to share their experiences implicitly and use data from one activity to directly affect user experience in the other. That is, usage data collected from users searching through a DL can be used to augment the browsing experience for other users and vice versa.

3. The System

To explore the problem of personalization in DL, we have developed a system that integrates community-based searching and browsing. Starting with social search technology employed in I-SPY [18], and the social navigation component, AnnotatEd, used in Knowledge Sea II [3], we designed a community-based information access system where these two technologies reinforce and complement each other.

The system was designed to provide community-based access to the articles in the Communications of ACM (CACM) magazine. CACM is well known for its broad scope and provides a good example of a collection of documents accessed by multiple communities. For the purpose of this paper we will call our integrated system *ASSIST-ACM*. This section explains how community-based access to CACM works from a user perspective, stressing the integrative aspects of the system.

The core data structure used to collect and store “community wisdom” is called a *hit matrix*, which can be thought of as a matrix-like data structure used to map relationships between user actions and useful content. Separate hit matrices are used for storing browsing and searching knowledge, where the information stored includes queries submitted, results selected (for the searching component), annotations provided for articles and hyperlinks selected from within pages (for the browsing component).

The ASSIST-ACM system monitors users as they interact with the system and builds up community wisdom, which it uses to augment pages with social cues. The search component uses the search history information in its hit matrix in future search sessions to re-rank the returned result-list so that it reflects the interests (selection histories) of the community. The browsing or *social navigation* component uses the browsing data in its hit matrix to identify parts of a page that have been interesting to previous users. Each hit matrix is used to enhance the overall user experience in both components, by enabling the user to see at a glance how their fellow community members have interacted with the links and content on the site. The integration of social navigation and social search paradigms can be seen in Figure 1, where the direction of an arrow indicates the flow of search or browsing knowledge. We can see that the search hit matrix is updated by the search component, but used to augment both

searching and browsing activities, and likewise the navigation hit matrix is updated by the browsing component but used to augment the user experience for both browsing and searching.. This synergy aids the user in better understanding which content is of interest to their specific needs and improves their ability to locate interesting information resources.

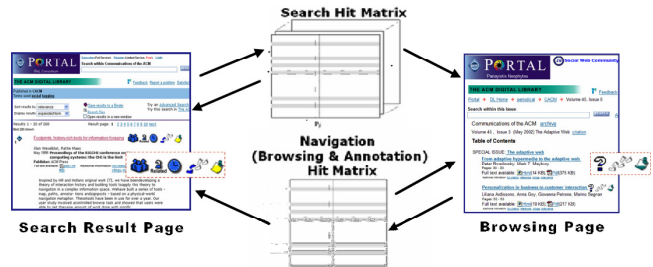


Figure 1: The interactions of the ASSIST-ACM browsing and search components with the system's hit matrices

From a user perspective, the system is an augmented version of the ACM Digital Library (which was used as the back-end of our system). As standard, the ACM Digital Library Web site (<http://portal.acm.org>) provides users with two information access strategies – searching and browsing. The following sections demonstrate in more detail how the augmented system helps students through each of these processes.

3.1 Search Component

Upon entering a query to the ACM Digital Library search box users are presented with a result page as seen in Figure 2. ASSIST-ACM alters the results returned by the standard ACM search facility in two ways. Firstly, I-SPY’s search technology re-ranks the result-list according to the preferences of the community. In short, articles that have been previously selected by classmates with similar queries are actively promoted by increasing their ranking in the result-list. Secondly, we utilize and extend a technique described in [5], and [6] whereby supplementary search and navigational access information is appended to any results in the result-list which the system has encountered previously. This information is added to the results in the form of a series of intuitive visual icons. In total, five icons are used, each denoting a different type of interaction information. Each icon appears with varying levels of fill to denote the level of interaction history recorded.

For example, the result that appears in position one of the result-list shown in Figure 2 is augmented with all 5 icons. The *popularity icon* (the first icon from left) shows the relevance of the result for the current query relative to the community in which it was entered. Its relevance value derives from the percentage of times this result has been selected by community members using the query. The *related query icon* (second icon from left) informs the user of any related queries for this result; that is, other queries that have also led to the selection of this result by the current community. A simple mouse-over reveals these queries; this provides the student with suggestions of other queries that may be relevant to their search task and also informs the user of other topics that the article may also be related to. The user can use one of these queries to initiate a new search, which enables them to refine or modify their query easily. The fill-level of the icon

represents the average relevance of the article in the result-list across all related queries.

The *recency icon* (third from left in Figure 2) provides users with information relating to the last time the result was encountered by users through searching, or browsing. This allows users to form a view of the freshness of the interaction trail. Results that have not been accessed in a number of months may not be as useful as results that have been accessed more recently. The *browse (footprint) icon* (fourth icon) explicitly connects the social search system to the social navigation system by indicating the browsing popularity of the article behind the link. The fill-level of this icon indicates the relative popularity of the article in question. For example, the article most visited by a community is 100% filled. The *annotation icon* (fifth from left in Figure 2) informs the user of the presence of annotations by community members. Annotated articles contain valuable additional user-provided information, which the authors of the annotations deemed relevant to members of the community. Once again a mouse-over of this icon shows annotation popularity.



Figure 2 - Search Result Page (Re-ranked and annotated based on community activity)

3.2 Browsing Component

An alternative way to access articles in the Social CACM is browsing. The user can browse the magazine hierarchically accessing a specific issue by year and month and then selecting a specific paper from a table of contents. In addition, cross-citation links between papers allow browsing the collection “horizontally”. Both ways of browsing are assisted by social icons, as can be seen in Figure 3.



Figure 3 - Social navigation support for browsing

Each link to a CACM article encountered during browsing is augmented with social information. It includes table-of-contents pages, as shown in Figure 3, as well as article information pages

with cross-citation links to other articles. The *annotation icon* (third icon from left) denotes the presence of annotations. The browsing popularity of the link is represented by the browse (footprint) icon (second from left). The query icon (the question-mark icon, first on left) complements the browsing interface with search information adding another connection between the search and browsing components. The presence of this icon indicates that the article has been selected for at least one query issued by the community members. A mouse-over the icon reveals a list of such queries. The user can use one of these queries to initiate a search for that query. Here we see an example of how this system facilitates the ability to seamlessly cross from browsing activity to search; this feature is discussed below.

Choosing an article, from either a search result-list or through a link on one of the browsed pages, brings the user to the article’s information page (see Figure 4). In addition to the standard article information provided by the CACM (authors, abstract, source etc.) the users are presented with additional social information pertaining to article’s search history. Integrated into the information page is a panel displaying the list of queries that have been used to locate the article. Each query is augmented with icons that represent the time of the last search and the query’s relevance to the article. This unique feature had several motivations. Firstly, listing past queries provides an insight to the article content, as mentioned in a previous paragraph. The second and more important motivation is to provide another connection between browsing and search knowledge and functions. By selecting a query in the list a search can be instantly initiated. Since this query was originally used to locate the viewed article, the results of the new search are also likely to be related to the viewed article. In effect selecting the query is like asking the system to find “similar” articles, or articles that are “more like this” without the user having to browse further, using the related query as an additional contextual cue.



Figure 4 - Article augmented with related community activities

The last context where our system augments the ACM Digital Library interface is the full content view of the article in HTML format. This view can be reached directly from a search result or table-of-contents pages as well as from the article information page. While reading the article, the users can contribute to yet another information access attribute, an annotation. This allows users to provide implicit feedback to the “community wisdom” gathered from the explicit search and navigation information. In our system, the usual article view page is augmented by a side panel which contains a full list of previous annotations and an interface allowing users to add and edit annotations. Annotations can be in the form of page highlights or notes. The notes (textual

annotations) can be categorized into positive or general annotations. The user can choose whether to make their annotations visible to just their community members or to the general public. Finally, the users also can choose whether to sign the note with their name, or keep it anonymous. The shared annotations are used to provide navigation support for the community as presented above.

4. The Study

To evaluate our ASSIST-ACM system, we designed a real-classroom study. The participants were a group of students enrolled in an "Introduction to Multimedia" course in the School of Computer Science and Informatics at University College Dublin. Four lectures of the course were dedicated to the topic of the social Web and the lectures included information pertaining to several social applications including social bookmarking, social networking, social navigation and social search. The functionality of the ASSIST-ACM system was detailed in lectures. In the evaluation students were asked to spend one hour using the system in order to complete a literature review on the topic of the social Web. The students were asked to locate papers which pertained to the topics introduced in lectures and to provide a very short explanation as to why the articles found were relevant to the topic of the social Web. The subjects were asked to fill out a short questionnaire at the end of the session providing information on their experiences with the system and their views of the various social features provided

4.1 Study Design

To assess the value of the community-based personalization functionality provided by the ASSIST engine to the ACM system, we deployed two versions of the system; 1) the *control* system which looks and feels like the original ACM Web site with the exception that users are provided with the facility to annotate any articles which they view and retrieve the annotated articles; and 2) the *experimental (test)* system, or the fully functional ASSIST-ACM system, which provided users with all the social support features detailed in Section 3. In both systems all user interactions are recorded and a pool of "community wisdom" was collected.

The participants were randomly assigned into the control and experimental group while controlling the number of participants assigned to each group. Each group was made up of 15 participants 7 of the control group's students used the system before the others and their data seeded the system for the experimental group. The remaining students in both groups used the system at the same time.

It is worth mentioning how the short-term nature of the study provided a limited ability to assess some system features such as annotations and annotation-based navigation support. In a simple information-finding task set over a short period of time there is little motivation to leave traces for oneself or to other users in the form of annotations. As expected, a low number of students added annotations to the digital library. However, as detailed below, the students reported that they did like the facility even if they didn't use it to its potential.

4.2 Hypotheses

The main goal of the study was to assess the values of social support in search and browsing and the integration of the social search and social browsing. We hypothesized that

1. Social support will help the students to increase their retrieval speed and allow them to find articles faster.
2. Social support will help the students to find more relevant articles.
3. Social support will attract student activity; i.e., results, and links with the social support augmentation will have a higher probability of being selected.
4. Search related icons will be useful in browsing while browsing related icons will prove to be useful during search.
5. Students in the experimental group will find the system to be more useful overall.

4.3 Results

4.3.1 Relevance of located papers

In total 54 unique articles from the CACM were returned by students of both groups. Each article was blindly examined and manually assigned a relevance score between 1 and 3, with 1 meaning "not relevant", 2 meaning "partially relevant", and 3 meaning "very relevant". Table 1 shows the average number of articles, average number of relevant articles, and average number of very relevant articles found by each group.

The results show that students using the ASSIST-ACM engine retrieved 13% more papers in the allotted time; this shows that the social features allowed students to increase their retrieval rate, thus supporting our first hypothesis. We also see an increase in the relevance of the results returned by the students in the experimental group: these students located 19.4 % more relevant articles than those in the control group. The difference is even pronounced (22.4%) when only the very relevant articles are examined. This result is in support of our second hypotheses, that social support helps students to find more relevant articles and thus perform more accurate retrieval.

Table 1 - Relevance of located papers by control and experimental group

	Control	Experimental
Average number of papers	4.33	4.93
Ave # of relevant papers (partially relevant & very relevant)	3.33	4.13
Ave # of very relevant papers	2.53	3.27

In order to discount any advantages that either group may have had over the other we examined the search and navigational behavior of both groups as they interacted with the system. Students in the control group submitted an average of 13.2 queries each, while students in the experimental group submitted an average of 9.13 queries each. Thus, the increase in the number of relevant articles reviewed cannot be explained by the experimental students performing more searches as they in fact performed 30% fewer searches than their classmates. Indeed using the number of queries per subject as a proxy for search effort, these statistics highlight how the test users are finding

significantly more relevant papers *per query* than the control group. For example, the former averaged 0.45 relevant papers per query compared to 0.25 papers per query for the latter, a relative benefit for the test group of approximately 80%.

Further to this we examined the queries that were used by each group in order to ascertain if the experimental group used significantly longer and therefore more specific queries, thus gaining them an advantage. Our analysis of search log data shows that students in both groups used queries of very similar length, 2.33 terms per query for the control group and 2.37 terms per query in the experimental group. This result confirms that the experimental group had no significant advantage in terms of more specific queries that would account for their increased performance [1], [2].

4.3.2 Search and Browsing Efforts

Since we are interested in showing that the social search and browsing system presented here provides a performance benefit to users in terms of their information retrieval abilities in a digital library, it is valuable to consider how hard users had to look to find interesting resources. We will consider how many search result pages a user had to visit, on average, to find a result of interest. Indeed, this aspect of any rank-based information retrieval system is very important for the overall effectiveness of the system. [11] found that 58% of users do not look beyond the first page of results, with 86% of users looking no further than the first 3 pages of results. A questionnaire-based study in [10] shows a very similar result, that 88% of users do not look past the 3rd result page. Thus, in the event that a user does not find an interesting result on the first page, it would appear vital to ensure that such a result appears at least within the next 2 pages.

Table 2 shows some figures for the average page number upon which users in the control and experimental groups selected a result. As we can see, both groups tended to select results on the first page, but what we are interested in here are the figures for result selections on pages 2 and higher. When the user selected a result that was not on the initial result page, it was on page 4.44 on average (standard deviation = 2.12) in the control group, but only page 2.29 (standard deviation = 0.54) in the experimental group. That is, users in the experimental group were required to look at only 1.29 additional result pages on average, compared to 3.44 pages for the control group. This is a significant saving in terms of user effort and speaks to the utility of promoting previously-selected search results within a community of users, even beyond the initial result page (to avoid swamping result lists, if there were more than 3 promotion candidates, the first 3 were promoted on the first page, the next 3 were promoted on the 2nd page and so on).

Table 2 - Result page hit number among control and experimental group

	Control	Experimental
Selections from 1st result page	92	99
Selections from 2nd and higher result page	18	23
Ave selected result page number	1.56	1.26

Ave. 2nd and higher result page number	4.44	2.29
--	------	------

As mentioned previously, many users will not bother looking past even the 1st page of results, with a rapidly diminishing percentage of determined users looking at result pages 2 and higher. Thus, it is important to ensure that the more determined users find useful results close to the initial result page, so that when selected they may be eventually promoted to the first page. This would increase the chances of all users finding an interesting result, not just the most determined users.

Similar results in regards of user efforts were obtained after examination of the navigational data. We learned that the control group students browsed an average of 12.53 links each in comparison to their classmates in the experimental group, who only browsed an average of 7 links each. Thus the students in the experimental group spent less effort not only in the search, but also in the browsing process. This allows us to exclude the possibility that the increased level of performance in the experimental group was achieved by larger investment of search or browsing efforts.

This result further supports our first hypothesis, that social support helps the users to find relevant results faster (and importantly, with less effort), particularly when those results do not appear on the first page of results.

4.3.3 Searching Success Rates and Browsing Access Rates

It is worth mentioning here the percentage of *successful* search sessions observed by both groups. In this context a successful search session is one where at least one result is deemed relevant and selected by the user and the *success rate* of a set of sessions is the percentage of the sessions that were successful. There was a 42% relative increase in success rates, from 32.97% for the control group to 46.7% for the experimental group, indicating that the social support provided by our system helps users to find apparently relevant articles more often.

From the results in this and the previous section, we can conclude that the social support offered by our ASSIST engine was the factor contributing to the improved performance of the students. We can see that the students in the control group had to do a lot more work in terms of entering search queries and browsing menus than the students in the experimental group and even at that they still failed to reach the success levels of the experimental group. Thus the provision of icons and the re-ranking of search results based on previous user behavior pointed the students in the direction of the most relevant material. This allowed them to locate relevant results faster and thus locate more results in the time allotted.

4.3.4 Effect of Social Support on Student Decision

We hypothesize that the addition of social cues will encourage the students to follow the links which have previously been found useful by the community. Specifically, we expect to observe more frequent selections when links are enhanced with the social cues described above. We are also interested in assessing the effect of mouse-over information. As mentioned before, mousing over each of the icons provides detailed information on the icons meaning. To measure the effect of social cues we define two probabilities

1. Impression-based selection probability: the probability that a link accompanied by social icons will be selected.
2. Event-based selection probability: the probability that a link for which the user mouses over one of its social icons will be selected.

Since the control group was not able to see any of the social cues, the data reported in this section is only related to the experimental group.

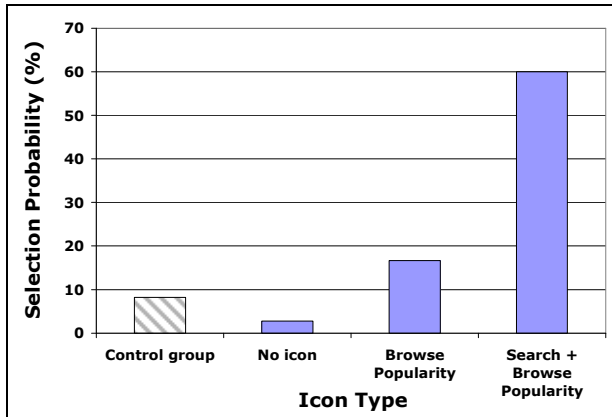


Figure 5: Impression based selection probability for socially-augmented links while browsing

Figure 5 presents the impression-based selection probabilities for links with social visual cues during browsing. The probability is computed by dividing the number of selected article links that are accompanied by a given icon type or combination of icon types by the total number of article links accompanied by that icon. We also calculated the same probability for links with no icons attached and a general selection probability of the control group while browsing to use as a baseline. The results show that social support has a significant effect on students' selection decisions. Firstly we see that results augmented with any form of social cues are more likely to be selected. Secondly, when browsing students are significantly more likely to select results that were augmented with both search and browsing information when compared to those that just contained browsing information. These findings directly support our 3rd and 4th hypotheses.

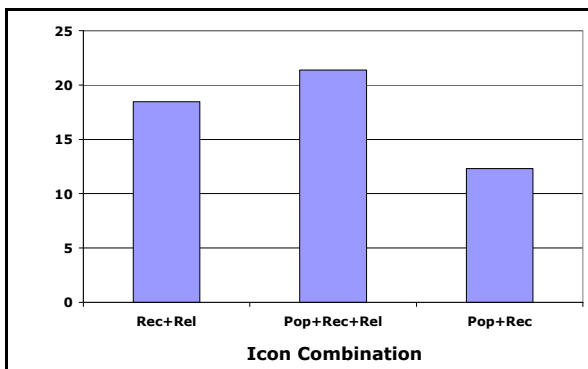


Figure 6: Impression-based selection probabilities for combinations of (search-related) explanation icons. In the above, POP, REC and REL refer to the popularity, recency and related queries icons, respectively.

Figure 6 shows the probability that a result with different combinations of the search explanatory icons attached will be selected. As we can see, the combination with the highest probability of a result selection contains all 3 icons. Since the recency icon is always shown, we cannot infer much as to its influence. However the fact that presence of a related query icon results in a much higher chance of a result selection, we can take this as evidence that the related query icon is the most influential in terms of user selection behavior.

Figure 7 presents the event-based selection probabilities for search results. We can see that when a user mouses over an icon, the chance that they will subsequently select the associated result increases. This is intuitive, the interesting point to note here is that the user mousing over the browse (footprint) icon is as good an indicator as any of the search-related icons that the user is likely to select the associated result, and so it seems that the integration of the browsing data to the search process can be quite useful in aiding the user to decide which result to select. This again supports our 4th hypothesis.

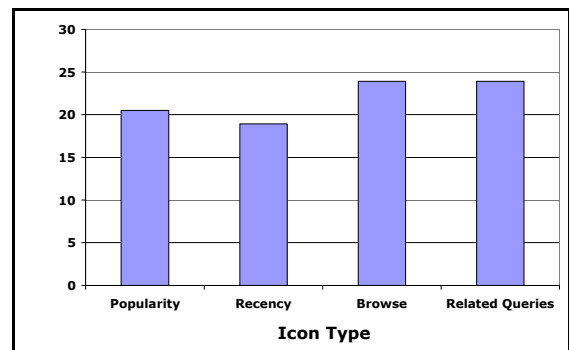


Figure 7: Event-based selection probabilities for search results

It is interesting to observe that the effect of social support is more significant throughout the browsing process in comparison with search in terms of the magnitude of the selection probabilities. We believe this is due to the fact that in browsing social navigation supports are the only decision support available to students while in the search the rank provides basic decision support. Also when searching the user goal is highly important and thus the goal specific wisdom i.e. the search wisdom will be more important to the user.




Figure 8: The selection probability for top-3 results with (Enhanced) and without (Standard) social explanation icons

Finally, we will attempt to reinforce our hypothesis that the addition of social cues to search results aids users in deciding which result to select. We will compare the selection probability

of results that have no social cues attached (*Standard* results) with the selection probability of those that do (*Enhanced* results). Since the enhanced results primarily occur in the top 3 result positions, they are likely to accrue a higher percentage of selections by virtue of their prominent position alone and so, we will limit our observations to the top 3 standard and enhanced results. Figure 8 graphs these selection probabilities. We can see that an enhanced result in the top 3 positions is more than twice as likely to be selected as a standard result. This is quite a strong indicator that our 3rd hypothesis, that social support will attract student activity, is correct.

4.3.5 Subjective Data Analysis

Students were asked to respond to a short questionnaire at the end of their session. We hypothesized that students in the experimental group would find the ASSIST-ACM system to be more useful in comparison to the control group students who used the regular ACM functional system. **Figure 9** compares the response of both groups to the question, “*Did you find the system useful to locate relevant information?*” In support of our final hypothesis, more than 70% of the ASSIST-ACM users (experimental group) found the system useful for locating relevant information, while only 40% of the regular ACM users (control group) found the system useful; alternatively, only 10% of ASSIST ACM respondents indicated that they did not find the system useful, compared with 40% of the regular ACM system users. This shows that with our social support system in place the ACM system became more useful in this retrieval task.




Figure 9: Ease of locating information

It has been mentioned that due to the student’s introduction to the topic and the system that they may be biased towards the ASSIST system. In the findings of the previous sections the usage of the features speak for themselves and in this section we feel that the students were giving feedback on either the regular ACM system or the ASSIST system and thus are reporting on their opinion on the system that they used and not comparing the two systems.

A feature shared in the experimental and a control version of the system is the annotation ability. The feature operated in the same way for both groups but the experimental group was able to view social support based on annotations as well. We asked users about the usefulness of the annotation feature. While 73% of the experimental group found the annotation feature to be useful only 26% of the control group found it useful. We expect that this difference in opinion is due to the differing levels of social support that the different groups observed. The control group

users could add annotations to pages but were unaware of whether annotations existed until they visited a page. The students in the experimental group were able to see the presence of annotations on a page by the presence of the annotation icon beside the link and thus annotations were much more visible to this group and so they found it more useful. This is in direct support of our 5th hypothesis.

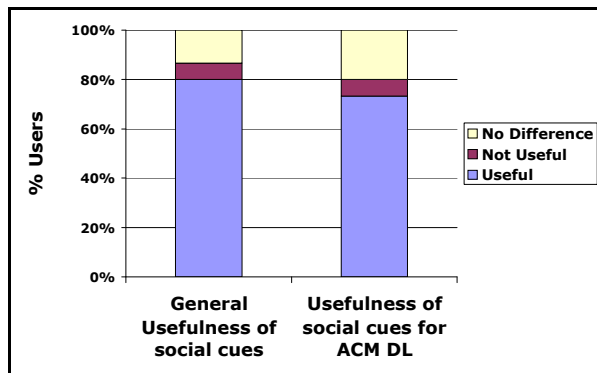


Figure 10: Usefulness of social cues

We also questioned the experimental group about the usefulness of social cues, in general and in the context of the ASSIST system. As shown in Figure 10 80% of the experimental group found the social cues useful in general and 73% find it useful in out ASSIST-ACM system. After speaking with the trial participants it was obvious that our social support was useful but that we needed to reconsider our use of social icons. Many students felt that while they understood the icons meaning there were often too many icons visible and that this confused the students a little.

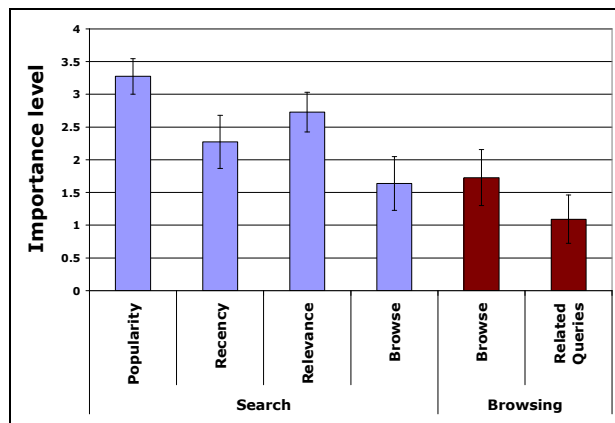


Figure 11: Users' opinion about importance of visual icons in searching and browsing

Figure 11 shows the results when students were asked which of the icons they found useful when searching and browsing. The possible answers ranged from always to never. To compare the usefulness of different cues we assigned numerical importance rating for each answer ranging from 4 for “always” to 0 for “never” and averaged the answers for each question. In a different question we asked the user whether they noticed the visual cues. 4 out of 15 users in the experimental group responded that they did not notice the visual cues. So, to evaluate the users’ opinion

about the importance of the icons, we excluded data from users who mentioned not noticing the icons.

We see in the searching scenario that users reported that they considered the popularity and relevance icons highly. As the trial was carried out over a few days the recency information was not of particular importance to the participants as all activity was in the same time frame, it is duly ranked lower than the other search information icons. As also mentioned users did not use the annotation functionality to its potential and thus the presence of the annotation icon was rare and thus its importance is lowered also. We see that the browse icon is also reported as being useful although not as useful as the search related icons. We feel that this reduction in importance is fueled by several facts. Firstly when searching the user has a specific goal in mind, the search icons that are displayed all relate to the users goal – similar queries recommended results etc and thus in a searching scenario the browse information is obviously less influential. Further to this the search value reflects the fact that the browsing activity during the trial was low and that the browse icons may not have appeared as often as it could have. However even taking these facts into consideration users still report that the icon is useful in a search scenario and thus this supports our 4th hypothesis.

Surprisingly the user opinion about the importance of social guidance icons in browsing context is quite low. To understand this result better, we looked at the usage pattern of browsing and searching for the students in the experimental group. As was mentioned above, ASSIST ACM supports two ways of browsing: (1) hierarchical browsing starting from table of content (2) horizontal browsing from paper to paper using cross-citation links. The first way is an alternative to the search process, while the second is rather a complement to search when the user starts browsing session from one of the results discovered during search. Our data shows that hierarchical browsing was used only in 3 cases, the rest of the sessions started from search. This could be explained by the nature of the task given to the users. As a result, browsing was mostly used as an extension of the search-driven access. It influenced different use of search and browsing. Out of total 162 accesses to articles, only 41 were done through browsing (25%). This result alone may explain the students' low opinion about the importance of browsing icons. Moreover, as we mentioned above, horizontal browsing using cross-citation links accounted for almost all cases of browsing. Unfortunately, the users had almost no chances to see social guidance icons during cross-citation browsing in this short-term study. While a typical paper offers a reasonable number of cross-citation links to other papers in ACM Digital Libraries, only links to CACM papers (few to none in each citation list) have a chance to be annotated. As a result the value of icons in the browsing process can be hardly appreciated by the user. In a different context reported in [2] where the users have many chances to see and to use social navigation cues during browsing process, browsing was used even more frequently than search and the user opinion about the importance of browsing cues was high.

4.3.6 Conclusion

In this trial we asked 30 students to complete a task using differing versions of a social support system. Our result shows that community wisdom collected thorough the search and browsing is an important support for the students while seeking information on the Web. In terms of retrieval users who used the

ASSIST engine retrieved more articles, and more importantly more relevant articles than their classmates who used the regular system. Secondly, we observed that the presence of social icons affected how users used the system, regularly mousing over the icons and following their associated links. It is here that the benefit of using search information in a browsing situation and vice versa is highlighted. Search information was utilized more often than browsing information when browsing and browsing information was utilized just as often as search information when searching. Thirdly, in a subjective analysis, students using the ASSIST engine reported a more useful system overall. Overall our results show us that a true integration of search and browsing enriched each component separately and the system as a whole

5. Discussion and Future Work

In this work we have presented a social support system –ASSIST which exploits a pool of information access wisdom to inform users of past interactions in the hope that its information is relevant to the users in their current tasks. The integration of two existing social wisdom driven systems was achieved at an interface level and integrated with the Communications of the ACM digital library. The result of the user study were positive and highlighted the potential gain for users when social search and social navigation wisdom is pooled and exploited by both technologies.

The ASSIST project has many opportunities for future work and future evaluations. The purpose of this small scale evaluation was to prove that the integration of the technologies was possible and beneficial. Currently researchers are looking into broadening the domain of the project outside the academic space. This will provide ample opportunities for user trials outside the classroom situation. We also wish to integrate the systems fully at a technology level to improve the overall performance of the system. As highlighted by the participants the use of icons for explanation purposed and social cues is beneficial but we need to rethink our strategy in order to use this method to a better potential.

6. ACKNOWLEDGMENTS

The authors thank Science Foundation Ireland for making this collaborative work possible through Grant No. 03/IN.3/I361, and through the E.T.S. Walton Award to the fourth author. This research is partially supported by the National Science Foundation under Grant No. 0447083, Vodafone Ireland, and the National Science Foundation graduate fellowship to the first author.

7. REFERENCES

- [1] Belkin, N. J., Kelly, D., Kim, G., Kim, J., Lee, H., Muresan, G., Tang, M., Yuan, X., and Cool, C. (2003). Query length in interactive information retrieval. In Proceedings of the 26th Annual international ACM SIGIR Conference on Research and Development in information Retrieval, Toronto, Canada, July 28 - August 01, 2003. (SIGIR '03).
- [2] Boydell O, and Smyth B. (2007) Enhancing Case-based, Collaborative Web Search. In Proceedings of the 7th International Conference on Case-based Reasoning (ICCBR '07), in press.

- [3] Brusilovsky, P., Chavan, G., and Farzan, R. (2004) Social adaptive navigation support for open corpus electronic textbooks. In: P. De Bra and W. Nejdl (eds.) Proceedings of Third International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH'2004), Eindhoven, the Netherlands, August 23-26, 2004, pp. 24-33
- [4] Brusilovsky, P., Farzan, R., and Ahn, J.W. (2005) Comprehensive personalized information access in an educational digital library. In: Proceedings of The 5th ACM/IEEE-CS Joint Conference on Digital Libraries, Denver, CO, June 7-11, 2005, ACM Press, pp. 9-18
- [5] Coyle, M. and Smyth, B. (2007a) On the Community-based Explanation of Search Results. In Proceedings of the 10th Annual Conference on Intelligent User Interfaces (IUI'07), Hawaii, U.S.A., (in press)
- [6] Coyle, M. and Smyth, B. (2007b) Supporting Intelligent Web Search. In ACM Transactions on Internet Technology Special Issue on Intelligent Techniques for Web Personalization, (in press).
- [7] Freyne J., Farzan R., Brusilovsky P., Smyth B., and Coyle M. (2007). Collecting Community Wisdom: Integrating Social Search & Social Navigation. In Proceedings of International Conference on Intelligent User Interfaces
- [8] Golbeck, J., and Wasser M. M. (2007). SocialBrowsing: Integrating Social Networks into Web Browsing. CHI 2007 Works-In-Progress. San Jose, California. April, 2007
- [9] Hill, W. C., Hollan, J. D., Wroblewski, D., and McCandless, T., (1992) Edit Wear and Read Wear. In: Proceedings of ACM Conference on Human Factors in Computing Systems, CHI'92. ACM Press, New York City, New York. 3-9
- [10] iProspect.com, inc. (2006) iProspect Search Engine User Behavior Study (April 2006). http://www.iprospect.com/premiumPDFs/WhitePaper_2006_SearchEngineUserBehavior.pdf
- [11] Jansen, B. J., Spink, A. and Saracevic, T. (2000) Real Life, Real Users, and Real Needs: a study and analysis of user queries on the Web. Information Processing and Management 36(2), 207-227.
- [12] Keller, R. M., Wolfe, S. R., Chen, J. R., Rabinowitz, J. L., and Mathe, N. (1997) A bookmarking service for organizing and sharing URLs. In: Proceedings of Sixth International World Wide Web Conference, Santa Clara, CA, April, 1997, pp. 1103 - 1114.
- [13] Larsen, R. L. (1997) Relaxing Assumptions . . . Stretching the Vision: A Modest View of Some Technical Issues. D-Lib Magazine 3, (April)
- [14] Lieberman, H., Van Dyke, N., Vivacqua, A.: Let's Browse: A collaborative Web browsing agent. In Maybury, M., ed.: IUI99: International Conference on Intelligent User Interfaces. ACM, New York (1999) 65-68
- [15] Millen, D. R., Feinberg, J., Kerr B. (2006). Dogear: Social bookmarking in the enterprise. Proceedings of the SIGCHI conference on Human Factors in Computing Systems, Montréal, Québec, Canada April 22 - 27, 2006. ACM Press, pp. 111-120.
- [16] Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., and Riedl, J. (1994) GroupLens: An open architecture for collaborative filtering of netnews. In: Proceedings of ACM conference on Computer supported cooperative work, Chapel Hill, North Carolina, ACM Press, pp. 175-186.
- [17] Shen, R., Vemuri, N.S., Fan, W., da S. Torres, R. and Fox, E.A. (2006). Exploring Digital Libraries: Integrating Browsing, Searching and Visualization. In Proceedings of the 6th ACM/IEEE-CS Joint Conference on Libraries (JCDL '06), Chapel Hill, North Carolina, ACM Press, pp. 1-10.
- [18] Smyth, B., Balfe, E., 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 and User-Adapted Interaction, 14(5). 383-423, 2004.
- [19] Torres, R., McNee, S. M., Abel, M., Konstan, J., and Riedl, J. (2004) Enhancing Digital Libraries with TechLens. In: Proceedings of The 4th ACM/IEEE-CS Joint Conference on Digital Libraries, Tuscon, AZ, June 7-11, 2004, pp. 228-236