2008 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology

Enabling Social Navigation on the Web

Athanasios Papagelis
University of Patras
Greece
papagel@ceid.upatras.gr

Manos Papagelis
University of Toronto
Canada
papaggel@cs.toronto.edu

Christos Zaroliagis
CTI & University of Patras
Greece
zaro@ceid.upatras.gr

Abstract

For a place that gathers millions of people the Web seems pretty lonely at times. This is mainly due to the current predominant browsing scenario; that of an individual participating in an autonomous surfing session. We believe that people should be seen as an integral part of the browsing and searching activity towards a concept known as social navigation. In this work, we extend the typical web browser's functionality so as to raise awareness of other people having similar web surfing goals at the current moment. We further present features and algorithms that facilitate online communication and collaboration towards common searching targets. The utility of our system is established by experimental studies. The extensions we present can be easily adopted in a typical web browser.

1 Introduction

Since its establishment, the World Wide Web (WWW) has experienced a prolific growth. It is estimated that the size of the WWW in June 2008 was 174 million websites [1]. Over the last couple of years, its growth has been largely driven by the increasing number of new forms of media on the web including blogs, social networks, video and photo sites, audio such as podcasts and much more.

The main aggregation point for accessing all these forms of media remains the *web browser*, the software application that enables a user to display and interact with information located on a website. Typically, websites contain links to other websites, and thus, web browsers allow a user to easily access information by traversing these links, a process known as *browsing* or *navigation*.

Browsing may be defined as opportunistic, reactive and unplanned information searching [17]. It is also commonly assumed that the browsing experience over the Internet is typically *passive*, in a sense that a user's search objective is not shared by anyone else and people cannot interact with web content, make personal notes, share comments and URLs [4].

However, traditionally, information search involves a series of interactions between the searcher and any available information source, including other people (e.g., in an academic library). Moreover, recent surveys of search strategies among knowledge workers [11] and in education [10] revealed search needs that are not supported by current search interfaces, such as the desire to collaborate.

We build on these observations and focus on enhancing the user browsing experience towards a process known as *social navigation* [5]. Social navigation describes the process where a number of people that share interests and searching goals decide to coordinate their efforts. As a design approach social navigation tries to raise *awareness* that social activities should be part of our information processing environments. Systems based on social navigation concepts typically make people more aware of each other and thus contribute to a more social experience of the information space. At the same time, awareness of others and their actions make a *space* feel more alive and turn it into something we might perceive as *place* [6, 8].

To this end, we introduce a system that aims to enable social navigation on the Web and make the following contributions:

- We extend the typical web browser's scope by providing means for connecting, communicating and sharing information with other users in a synchronous way.
- We present an intuitive user interface that is able to visualize awareness of others and their actions.
- We present algorithms that render the visualization tool scalable and evaluate the utility of the system by conducting user experimental studies.

In the rest of the paper, we first present an overview of the related work and provide a description of the system. Then, we formalize the problem of correlating temporal user navigational patterns and provide algorithmic solutions. We continue with an experimental evaluation of the system. Finally, we summarize our contribution.

978-0-7695-3496-1/08 $25.00 © 2008 IEEE
DOI 10.1109/WIIAT.2008.86

2 Related Work

Interfaces to databases have traditionally been designed as *single-user systems*. The existence of other users and their activities have been implicitly assumed to be an attribute of the system that should be hidden from end-users [14]. Similar design approach has been adopted for accessing information on the web.

In recent years the emergence of the field of *computer supported cooperative work* has highlighted the importance of collaborative approaches in many diverse activities. In regards of online collaboration, research on *annotating* online documents has emerged. Annotations are usually in the form of comments, notes or search trails attached to any online document [9]. Then, online navigation experience can be enhanced by providing to future users annotations from users in the past [18].

More recently, research on *social navigation* attracts a lot of attention. Social navigation is based on the social navigation theory introduced by Dourish and Chalmers [6]. It works by taking advantage of patterns of agreement and tastes between users [3]. These patterns are often consistent across a time period. The premise is that users who agreed with each other in the past are likely to agree in the future. Further, research in *collaborative search* aims at facilitating collaboration of *small groups* on performing search tasks, such as students working together on assignments, friends seeking information about recreational events, couples planning vacations and more [12].

Our work is complementary to work on annotation since we aim to offer real-time annotation and information sharing of online content and is mostly related to research in social navigation and collaborative search. It complements work on social navigation in that we aim to enable opportunistic collaboration between people that share *temporal* patterns of agreement and tastes. It also differs from work on collaborative search in that we aim to enable collaboration on performing search tasks *at-large*, with people not on one's contact list. Overall, we aim to facilitate collaboration by supporting awareness that more people exist in the same place at the same time and at providing means of real-time opportunistic communication with one another via the web browser.

3 System description

In this section we thoroughly describe our system that aims to enhance the user's browsing experience. Before elaborating on the actual browser extensions, we present a motivating example of the required functionality.

Motivating Example: Consider a searcher that tries to find a good diet book. A common searching scenario consists of

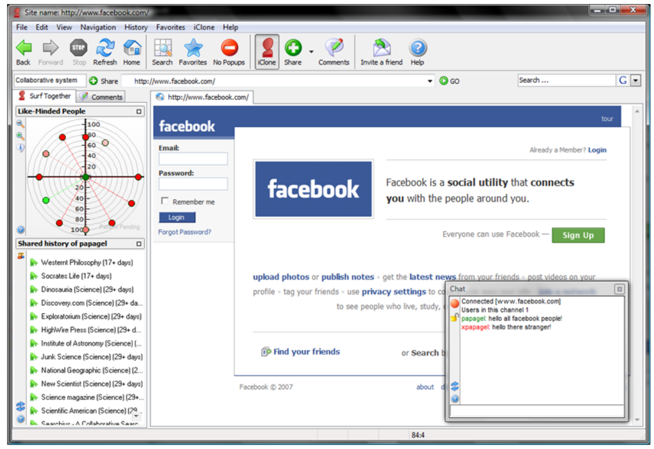

Figure 1. The web browser application. On the left, the *semantic neighborhood radar* (top) and the *shared history list* (down) of a user are presented. On the right, the *website-based chat* is presented that enables communication between users that co-exist in a website. On the top of the browser the *collaborative annotation tool* is presented that allows annotations to be assigned to shared websites (button "Share" near the address bar).

submitting free text queries in a search engine (e.g., "diet book", "diet book reviews") and visiting a few of the returned results. These results lead to other, potentially useful, websites that the searcher may decide to visit. These sites can lead to other sites and so on. The sequence of the recently visited websites can be used to represent a temporal user profile. This profile can then be compared to the temporal profiles of other currently online users to detect the most similar ones. Our system should identify these people, present them to the searcher in a comprehensive way and provide ways to communicate with them. As the searcher's interests shift through time people that appear in the interface shade away producing an up-to-date set of relevant users. In fact, each user's set of relevant users is dynamic and gets constantly updated due to its own and other user's browsing activity.

3.1 Browser Extensions

Our system is implemented as a stand-alone web browser application. More specifically, it consists of a number of tools, tightly integrated with a typical browser to extend its functionality. During the navigation process, information is collected and communicated to the main web server. Therefore, our system resembles the client-server architecture model, where many clients (web browsers) connect to a main server. Figure 1 presents a typical screenshot of this application from the end-user's point of view. In the remaining of the section we describe these extensions.

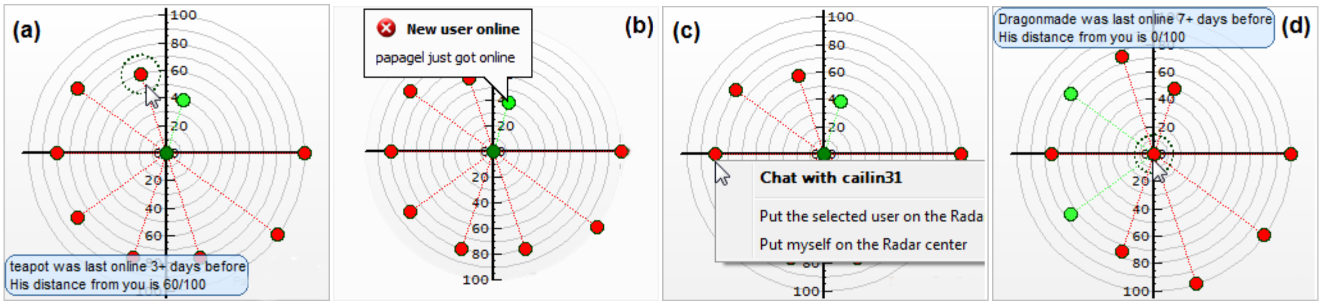

Figure 2. (a) The active user appears at the center of the radar. The distance of the dot at the center of the radar to any other dot represents the proximity of the users represented by these dots. The color of a dot represents the status of the associated user. By moving the mouse pointer over the radar we can discover more information for each user. (b) Activity indicators inform a user about the recent activities of other people in the neighborhood. (c) By right-clicking on a person, options to invite him to a private chat or to put him at the center of the radar are provided. (d) Choosing another person at the center of the radar reveals the users in his neighborhood.

3.1.1 The Semantic Neighborhood Radar

Visualizing a concept is often challenging since one needs to balance between informative, comprehensive and computationally feasible interfaces. We choose to represent the information of recently relevant users using a radar metaphor. A radar in the real world, that operates on an object x (the reflector), scans a wide area, measures the distance of other objects to x and presents these objects along with their distances from x on a display.

In our case, objects are users and distance is a metric of user to user proximity. This visualization enables users to instantly conceive the details of our conceptual model. It also has the nice property of displaying a lot of information in a restricted area.

The radar represents people as dots. The active user, for which the radar is defined, appears at the center of the radar, while the people most relevant to the active user are plotted on the radar in distances from the center that respect the computed proximity of each user to the active user (see Figure 2(a)). Conceptually, this represents a *semantic neighborhood* around the active user, with the captured semantic being the correlation of recent user navigational patterns.

Note, however, that when we are placing dots (users) on the radar we do not ask to respect all pair-wise similarities of all users. That case usually appears in the literature as the *ordination problem*, where we need to represent n-dimensional data by a small number of salient dimensions and thus be able to display multivariate data on the two-dimensional surface. The main tool for the ordination problem is the family of the *dimensionality reduction techniques*, including eigenvalue decomposition, multidimensional scaling, latent semantic analysis and more. Despite the wide adoption of these techniques in various problem areas, they suffer from limitations that render them inadequate in our case; they cannot be applied on a dynamic environment where the 2D layout needs to be regularly updated.

The radar metaphor evokes the proximity functionality we discussed, but also adds new features such as:

- **Representing the Time-axis on the Radar**: An essential aspect when representing users on the radar is to clearly indicate whether they are (recently) active or inactive. Active users are represented as green dots, while inactive users are represented as red dots. Furthermore, we would like to represent how recently a specific user has been relevant to the active user. To represent this information for each dot we use a spectrum of its color (either green or red) that spans from dark to light, with darker meaning more recently. The active user is therefore always represented as a dark green dot at the center of the radar (see Figure 2(a)).

- **Action indicators**: In addition to who is relevant and to what degree, we would like to also make available information of what people in our neighborhood are doing. To this end, we design action indicators that track the activity of people in the neighborhood and communicate the actions to the user (e.g., who got online, etc.). Action indicators are visualized in the form of a balloon assigned to a specific user. By this way, despite the stateless environment on which browsers operate, we are able to track the state of the radar information (see Figure 2(b)).

- **Private Chat**: One of the direct communication features that are provided by the system is the private chat. A user is able to start a private chat conversation (after invitation) with any of the users presented in the semantic neighborhood radar. (see Figure 2(c)).

- **Exploration of other Neighborhoods**: Another feature of the radar is that it provides the possibility to set another user at its center. By this way, one can explore the neighborhood (the relevant users) of another user. By traversing from a neighborhood to another, one can discover and communicate with more people (see Figure 2(c) and Figure 2(d)).

3.1.2 Website-based Chat

The website-based chat allows people that coexist in a website, during their navigation, to directly communicate with each other. Navigating a website automatically makes you a part of that website's *virtual public chat room*. Therefore communication between all users in a specific website is enabled (see Figure 1). Note, that our system enables the concurrent communication of users at any website and is not the same with the Web applications that allow to a website owner to directly communicate with its visitors. Interestingly, the website-based chat forms the foundation for transforming a *space* to *place* [8].

3.1.3 Sharing Information Spaces

Pointing out interesting information in a collaborative system is essential. Beyond (private or public) chatting, our system further supports exchange of information in the form of a shared history feature. During navigation, a user may share websites (along with tags) with the people that appear in his radar. A user receiving suggestions for websites would need to click on a user's dot at the radar to indicate his intention to see this list (see Figure 1).

3.1.4 Collaborative Annotation System

Our system provides the functionality to annotate a website with a set of keywords. By this way, a set of keywords is assigned to each website coming from different users at different times and therefore define a collaborative website annotation system (see Figure 1). These keywords can then form the basis for a number of applications.

4 Algorithms

With millions of users accessing billions of webpages everyday one would consider a visit of a user u to a website w at time t to be the primitive action that takes place in the overall web browsing activity. From a user's u perspective the chronologically ordered sequence of visits to a number of webpages defines a web history log H_u. Our system functions as the aggregation point of these individual history logs by defining a unified web history log H. In its most simple form, H consists of a chronologically ordered sequence of visit records of the form (u, w, t).

Such click-through information presents a challenging opportunity for analysis and mining with the goal of personalization and has been extensively used in research [15, 2, 16]. However, most existing approaches use the click-through data to devise similarity measures with little consideration of the *temporal factor*. At the same time, these data are often dynamic and contain rich *temporal information* [19]. In this section we present a *time-dependent* similarity model that exploits the temporal characteristics of historical click-through data. The intuition is that since the information needs of users change through time, user profiling algorithms should take into account the timestamps of the historical click-through data in order to identify regions of significant similarity that may be a consequence of functional relationship.

Formally, given the unified web history log H of all users and with respect to a temporal factor T, we would like to find a set of users that have similar web history in the time period bounded by T. In other words, we would like to identify users that recently navigated same websites. To express the temporal factor we use as a surrogate for time the number of the last T records in the unified log H and operate on the subset $H^T \subseteq H$ of the last T visits. In similar manner if we would like to constrain a specific user's u history log to its last S visits we write $H_u^S \subseteq H_u$.

Note that in order a web browser client to show the radar to a user u, it needs to communicate with the main server and retrieve the required information. This information consists of the set of users relevant to u along with their associated similarity values. Instead of computing this information whenever a client sends a request asking for the information to show on the radar, we perform all computations in a preprocessing phase that takes place in specific time intervals and cache the results. The premise of the preprocessing is that whenever a client requests information, this would be promptly available. Therefore we require that the client-server communication takes place in periodic time intervals that allow for the pre-computation phase to finish. Algorithm 1 describes the steps of the preprocessing phase.

Algorithm 1 takes as parameters the unified history log H, the temporal factor T, the temporal factor S and the maximum number of relevant users to be retrieved for each user k. First, it forms the set H^T of the last T records of the unified log H. Then, it identifies the set of unique users U in the set H^T and for each user $u \in U$ retrieves the set of its temporal history log H_u^S that corresponds to its last $|S|$ log records from H. For each pair of users u_i, u_j in U it computes their affinity by computing the overlap of their temporal history sets $H_{u_i}^S$ and $H_{u_j}^S$ and saves the score to array A. For each user $u \in U$ it computes the top-k users according to the scores in A and saves them in L. Finally, L is returned that keeps information about the most relevant users of each user along with their scores.

Algorithm 1 Finds Relevant Users

procedure COMPUTERELEVANTUSERS(H, T, S, k)
    L is a HashTable of the form $L < u, < Set >>$
    A is a HashTable of the form $A << i, j >, A_{i,j} >$
    H^S is a HashTable of the form $H^S < u, H^S_u >$
    $H^T = getLastRecords(H, T)$
    $U = getUniqueUsers(H^T)$
    for all $u \in U$ **do**
        $H^S_u = getLastRecords(H_u, S)$
    end for
    for all $u_i \in U$ **do**
        for all $u_j \in U$ **do**
            if $u_i \neq u_j$ && $i < j$ **then**
                $A_{i,j} = Affinity(H^S_{u_i}, H^S_{u_j}) = s(i, j)$
            end if
        end for
    end for
    for all $u \in U$ **do**
        $L_u = getTopK(A, k, u)$
    end for
    return L
end procedure

For computing the affinity between users i, j we employ a simple realization of the set-similarity algorithm. Given two sets $H^S_{u_i}$, $H^S_{u_j}$, with $i \neq j$, we can quantify the affinity of the sets by functions measuring their overlap, such as the *intersection* (I) or the *jaccard similarity coefficient* (J). For example, the similarity s between i and j, based on the intersection measure, is defined as:

$$s(i,j) = \frac{I(H^S_{u_i}, H^S_{u_j})}{max(|H^S_{u_i}|, |H^S_{u_j}|)} = \frac{|H^S_{u_i} \cap H^S_{u_j}|}{max(|H^S_{u_i}|, |H^S_{u_j}|)} \quad (1)$$

Other choices for quantifying user affinity are also possible. Our framework can easily incorporate any of these choices. Then, the distance d between i and j is defined as:

$$d(i,j) = 1 - s(i,j) \quad (2)$$

Note that this method does not take into account the popularity of the websites that participate in the similarity computation. However, it is reasonable to assume, that some websites are more important than others for defining similarity in our context. Indeed, we would like to weight websites in such a way that less popular websites are more significant than popular ones when determining similarities.

A way to naturally capture this property is to weight websites according to their inverted frequency in the unified history log H^T. More specifically, the frequency of a website w in H^T is given by:

$$f_w = \frac{n_w}{\sum_{i \in H^T} n_i}$$

where n_w is the number of times w appears in H^T. This count is normalized to prevent a bias towards longer logs and to give a measure of the importance of the website w within the particular log H^T. Then, a weight z_w of an element in the set is defined as follows:

$$z_w = \frac{1}{f_w}$$

We may now consider a weighted version of our set similarity algorithms, where there is a weight z_e associated with each set element e (i.e., with each website). Our approach is to convert a *weighted set* into an *unweighted bag* by making z_e copies of each element e. We use standard rounding techniques if weights z_e are non-integral. In that case, the similarity s between i and j is defined as:

$$s(i,j) = \frac{\sum_{w \in |H^S_{u_i} \cap H^S_{u_j}|} z_w}{max(\sum_{w \in |H^S_{u_i}|} z_w, \sum_{w \in |H^S_{u_i}|} z_w)} \quad (3)$$

We restrain our algorithmic description to the simplified assumption that the sets of users recent history logs fit in memory and that computation of set similarities between time intervals is feasible. Actually, we can enforce that condition by dynamically tuning the temporal-constraint according to system overload and available processing power.

5 Evaluation

Evaluating the utility of the proposed system is challenging on its own. Properties of the the social navigation paradigm differ substantially from ones in traditional systems, rendering many evaluation techniques obsolete. In our context, we are interested to evaluate the utility of the system in terms of *collaboration effectiveness* and *overall user satisfaction*. Consequently, we have to resort to user studies to address the evaluation issue. Given these considerations, we evaluate our system under two hypotheses:

- **Hypothesis 1**: Our system is able to raise awareness that other people have similar information needs at the current moment and to further identify and present these users in an informative way.
- **Hypothesis 2**: Our system serves as an online collaboration tool that helps people fulfill collaborative tasks more efficiently.

5.1 Data Sets

Information seeking on the web typically involves submitting queries to search engines supplemented by manual navigation [13]. Queries are usually classified according to their intent into three classes: navigational, informational and transactional. For the needs of our experimental evaluation we focus on informational queries (IQ), for which answers are assumed to be present on many web pages. These queries are the most likely to be benefited by social navigation tools, since the information seeking task requires to constructively select information from several sources. We further distinguish the informational queries into two categories; ambiguous (AIQ) and unambiguous (UIQ). AIQs do not require that a specific answer is sought, while UIQs usually look for a specific answer. For the various experimental scenarios we consider the queries of Table 1.

Table 1. Queries Data Set

Type	Query
AIQ_1	Find information and reviews about iPhone
AIQ_2	Find information about Ancient Rome
AIQ_3	Find reviews about Xbox
AIQ_4	Find information about Egyptian pyramids
AIQ_5	Collect information about USA presidents
AIQ_6	Collect information about space exploration
UIQ_1	Find the list of the current prime ministers of the world that are lawyers
UIQ_2	Find the list of paintings that have been sold for more than 1 million dollars

5.2 Experimental Studies

First Study: To test the first hypothesis we asked from 24 students (subjects), to participate in a study. The subjects were familiar with the Web searching task but not aware of our browser extensions.

First, subjects were asked to surf the web using our application without a specific objective for a time interval (*phase-1*). Then, each subject was assigned a search task from the set of the ambiguous informational queries of Table 1 ($AIQ_1 ... AIQ_6$). Each query was assigned to exactly four participants. Participants were not aware of each other's assigned search tasks. Subjects were asked to fulfill their assigned task for a predefined time period (*phase-2*). After this period has elapsed subjects were asked to surf the web without a specific objective for another time interval (*phase-3*).

The objective of this experiment is to evaluate how well the radar groups together subjects that have been assigned the same searching scenario. The intuition behind this experiment is that although subjects are free to perform the given searching task following their judgement, there is a high probability that some intersection between the subjects that have been assigned the same task will occur.

Second Study: To test the second hypothesis we asked from 16 students (subjects) to participate in a study. We separated the subjects into four groups of four subjects each. The first two groups (group1, group2) were assigned the query UIQ_1 and the other two groups (group3, group4) were assigned the query UIQ_2 from Table 1.

Then, subjects were asked to fulfill their assigned task for a predefined time period. The subjects of group1 and group3 were allowed to use the browser extensions provided by our application to communicate their findings, while the subjects of group2 and group4 were asked to use a typical web browser.

The objective of this experiment is to evaluate whether and to what degree direct collaboration, that is inherent to systems like ours, could improve the performance of an informational task. The intuition behind this experiment is that as people in group1 (group3) collaborate towards a common goal they will outperform group2 (group4) and will produce better results (either more or the same amount but in less time).

5.3 Results

First Study: A simple radar observation revealed the following: At the end of phase-1 subjects were far away from each other which indicates that they did not share the same interests. Then, at the end of phase-2, subjects that had the same search task approached each other. Finally, in phase 3, subjects drew away from each other.

For this experiment we used as evaluation metric the average distance $avgd$ between the set of users defined by each query. Formally, let the set of subjects that have been assigned the query AIQ_x be G_x. Then the average distance between subjects of G_x and G_y is defined as:

$$avgd(G_x, G_y) = \frac{\sum_{i \in G_x, j \in G_y} (d(i, j))}{n}, \forall i < j \quad (4)$$

where n is the number of distinct pairs of subjects between G_x and G_y and $d(i, j)$ is the distance of subjects i and j as defined in Equation 2. Table 2 shows the results for phase-2.

Table 2. First User Study Results

	G_1	G_2	G_3	G_4	G_5	G_6
G_1	0.77	0.99	0.99	0.95	0.97	0.98
G_2	-	0.75	0.99	0.97	0.99	0.97
G_3	-	-	0.73	0.99	0.98	0.96
G_4	-	-	-	0.89	0.95	0.99
G_5	-	-	-	-	0.78	0.98
G_6	-	-	-	-	-	0.77

The results indicate that our algorithm was able to capture the fact that some users had similar search interest. Note the diagonal of the Table 2, where it is clear that subjects that have been assigned the same search task came closer at the end of phase-2.

Second Study: For this experiment we need metrics to evaluate the performance of individuals and of groups. For the former, we define the *mean subject performance* to be the average number of distinct results found by the subjects of a group. For the latter, we define the *performance of a group* to be the number of distinct results found by all subjects in a group. Table 3 presents the results of this experiment.

Table 3. Second User Study Results

Query	Group	Average Subject Performance	Group Performance
UIQ_1	1	9	35
UIQ_1	2	9.75	13
UIQ_2	3	12	42
UIQ_2	4	33.5	37

Regarding UIQ_1, subjects in group1 and group2 had similar individual performance. However, group1 outperformed group2. Users of group1 coordinated their effort by communicating a simple strategy. They separated the workload (e.g., by continent) to four equal parts and assigned each part to a subject. On the other hand, all subjects in group2, which where not collaborating, followed a natural strategy of looking for prime ministers of popular countries causing their results to have a large overlap. Eventually, the results of the one were subsumed by the results of the other.

Regarding UIQ_2, subjects of group3 coordinated their effort by assigning only one member to the straightforward task of collecting the top paintings from available listings (e.g., Wikipedia). The rest of the subjects continued searching for expensive paintings by either submitting more sophisticated queries in search engines or by visiting the websites of known galleries and auction houses. On the other hand, subjects of group4 spent most of their time collecting the same paintings from listings of the top paintings.

6 Concluding Remarks

To enable social navigation on the web, we had to design a system that makes use of web history logs to encourage communication and collaboration among large groups of people. To support these features some personal information and a certain, limited amount of visibility of users' actions is required, which eventually infringes on user privacy. Privacy concerns could therefore serve as a major stumbling block towards acceptance of our system.

Erickson and Kellogg [7] introduced the concept of *social translucence* as an approach to designing systems that support social processes. According to this concept it is not only necessary to see other users, but to clearly communicate what information is disclosed and how it is used. We followed the same approach when designing our system and made sure that it entails a balance of visibility, awareness of others, and accountability.

The main idea of our system is to utilize temporal correlations between users' web history logs. Though intuitive the realization of such a system is not trivial since it poses a number of challenges spanning from technical to social aspects. Overall, the proposed application operates as a system for storing meta-data of web browsing activity, chat conversations and URL suggestions. By seamlessly merging the available collaboration data our system forms the foundation for a practical social navigation system.

References

[1] Web server survey (by www.netcraft.com), January 2008.

[2] E. Agichtein, E. Brill, S. T. Dumais, and R. Ragno. Learning user interaction models for predicting web search result preferences. In *SIGIR*, pages 3–10, 2006.

[3] J. Budzik, S. Bradshaw, X. Fu, and K. J. Hammond. Clustering for opportunistic communication. In *WWW*, 2002.

[4] S.-J. Chang and R. E. Rice. Browsing: A multidimensional framework. *Annual Review of Information Science and Technology*, 28, 1993.

[5] A. Dieberger. Supporting social navigation on the worldwide web. *Int. J. of Human-Computer Studies*, 46, 1997.

[6] P. Dourish and M. Chalmers. Running out of space: Models of information navigation. *Proceed. of HCI'94*.

[7] T. Erickson and W. A. Kellogg. Social translucence: an approach to designing systems that support social processes. *ACM Trans. CHI*, 7(1), 2000.

[8] S. Harrison and P. Dourish. Re-place-ing space: the roles of place and space in collaborative systems. In *CSCW*, 1996.

[9] J. Kahan and M.-R. Koivunen. Annotea: an open rdf infrastructure for shared web annotations. In *WWW*, 2001.

[10] A. Large, J. Beheshti, and T. Rahman. Gender differences in collaborative web searching behavior: An elementary school study. *Inf. Process. Manage.*, 38, 2002.

[11] M. R. Morris. A survey of collaborative web search practices. In *CHI*, 2008.

[12] M. R. Morris and E. Horvitz. Searchtogether: an interface for collaborative web search. In *UIST*, 2007.

[13] S. Pandit and C. Olston. Navigation-aided retrieval. In *WWW*, 2007.

[14] T. Rodden, J. Mariani, and G. Blair. Supporting cooperative applications. In *CSCW*, 1992.

[15] J.-T. Sun, H.-J. Zeng, H. Liu, Y. Lu, and Z. Chen. Cubesvd: a novel approach to personalized web search. In *WWW*, 2005.

[16] J. Teevan, S. T. Dumais, and E. Horvitz. Personalizing search via automated analysis of interests and activities. In *SIGIR*, 2005.

[17] M. B. Twidale, D. M. Nichols, and C. D. Paice. Browsing is a collaborative process. *Inf. Process. Manage.*, 33(6), 1997.

[18] A. Wexelblat and P. Maes. Footprints: history-rich tools for information foraging. In *CHI*, 1999.

[19] Q. Zhao and S. C. H. H. et al. Time-dependent semantic similarity measure of queries using historical click-through data. In *WWW '06*, pages 543–552.