Let’s Browse: A Collaborative Web Browsing Agent

Henry Lieberman, Neil W. Van Dyke, and Adriana S. Vivacqua
Media Laboratory
Massachusetts Institute of Technology
20 Ames Street #E15-305
Cambridge, MA, 02139 USA
{lieber, nwv, adriana}@media.mit.edu

ABSTRACT
Web browsing, like most of today’s desktop applications, is usually a solitary activity. Other forms of media, such as watching television, are often done by groups of people, such as families or friends. What would it be like to do collaborative Web browsing? Could the computer provide assistance to group browsing by trying to help find mutual interests among the participants? *Let's Browse* is an experiment in building an agent to assist a group of people in browsing, by suggesting new material likely to be of common interest. It is built as an extension to the single-user Web browsing agent Letizia. Let’s Browse features automatic detection of the presence of users, automated “channel surfing” browsing, and dynamic display of the user profiles and explanation of recommendations.

Keywords
Browsing, collaboration, agents, user profiles

COLLABORATIVE BROWSING
Increasingly, Web browsing will be performed in collaborative settings, such as a family at home or in a business meeting. For example, WebTV estimates that the average number of people who are watching during a session with its service is two, indicating that multi-user browsing is the norm rather than the exception. In most such situations, one person has control of the remote or the keyboard and mouse, and the others present are relatively passive. Yet the browsing session can't be considered successful if the interests of others present are not taken into account.

Collaborative browsing can take many forms. A group of people may be searching for specific information, exploring previously unexplored territory to see what's interesting, or some combination of the two. What links a person chooses to view and how a person reacts to what appears can also serve as a meaningful form of communication between the participants. Participants can learn about each other as well as learn about the content of the Web pages.

Fig. 1: Let’s Browse screen display

Increasingly, also, we believe that Web browsing will be assisted by intelligent agent software, which can keep track of user’s interests, inferring interest from observing user actions, and autonomously looking for items that satisfy interests. We have had extensive experience with such a single-user agent, *Letizia*, [8, 9] which performs reconnaissance on Web pages. Letizia proactively fetches links from the page currently viewed, and chooses those pages that best match a user profile learned by watching the user’s choices. Letizia presents its recommendations in a separate, “channel surfing” window that continuously displays recommendations.

We were interested in extending the channel surfing metaphor to situations where, as in TV, more than one person may be watching. Even if only person “has the remote control,” the agent can be cast in the role of representing the interests of the other participants, without requiring negotiation at every step, which may be disruptive. The job of the agent is to choose, from the links reachable from the current page, those that are likely to best satisfy the interests of all the participants. We call the resulting system *Let's Browse.*

EXPERIENCE WITH THE PROTOTYPE
The first experiment with the Let's Browse prototype system was set up on the occasion of the Media Lab's Digital Life Consortium meeting in October 1997. The idea

was to experiment with an automated browsing system that surfed on behalf of the common interests of a set of users who were physically present at a single screen. The environment was originally intended for use with the users present before a wall-sized display, but was actually deployed with a single 19" CRT monitor (Fig. 1).

The goal is that users can walk up to the screen, and their presence is automatically sensed by Let's Browse, without any explicit action on their part. The screen would continually show a selection of Web pages that the agent determined might be of common interest to the participants, together with explanation of its choices.

Sensing the Presence of Users

Presence of the users was detected by use of "Meme Tags," active badges worn by the participants (Fig. 2), prepared by the Media Lab's Learning and Epistemology group. These badges were also being used at the same event, in another experiment to track the flow of *memes*, short phrases that were communicated from person to person in the group [1]. These electronic badges communicated the user's identity, transmitted through an infrared link. The range of transmission is good to a few meters in line of sight of the receiver, mounted on the Let's Browse display screen.

Interest Profiles of the Users

The identities of the users currently in front of the screen are used to index into a set of user profiles, in order to construct an interest profile for the group. We used several heuristics to try to determine user interests (semi-) automatically, rather than present the users with an explicit interest questionnaire.

In advance of the event, we determined interest profiles for each of the 475 attendees, by running an off-line Web crawler that scanned a breadth-first search around each attendee's page. The crawl started from either the home page of the attendee or the home page of the attendee's organization if their personal page was not known.

For each page, we performed a TFIDF (term frequency times inverse document frequency) keyword frequency analysis to extract terms approximating the "subject matter" of each page, a common technique in Web search engines. The crawler operated with a time limit that typically allowed scanning between 10 and 50 pages per

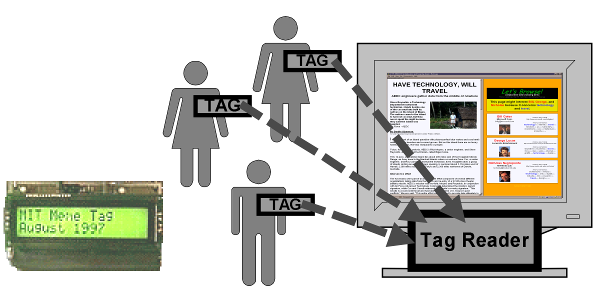

Fig 2: A Meme Tag and the kiosk configuration

user. A user profile is essentially a set of keyword-weight pairs, similar to the user profiling techniques used in the single-user agent Letizia [8].

The purpose of the off-line compilation of user interests was to "prime the pump", since when the user walks up to the screen, the agent typically does not have enough time to begin the process of determining the user's interests before the user expects to see some meaningful results. This is in contrast to Letizia, which does not start out with any precomputed profiles. We also found that, in contrast to a single user, who often has the patience to wait a while until the system learns his or her preferences, groups are far less patient, and it is important that the system be able to demonstrate its capability to make good choices quickly. However, if the Web access were fast enough, all the interest profiles for Let's Browse could also be computed on the fly.

With the current setup, indexing into precomputed user profiles gives the agent some immediate fodder to work on. Extra compute time during operation can be used to continue the search, deepening the breadth-first search for each user, interleaving this with the search for pages to recommend to the user. If a user appears whose profile is not already stored in the database, the system can start crawling around his or her web page to dynamically compute a profile, getting as good an approximation as is possible in the time allotted.

An interesting result from the Let's Browse experiment concerned the number of terms stored in the user profile. In Letizia, only a small number of terms were retained from the TFIDF computation on each page, typically about 10. This was justified by results from information retrieval that show that retaining a large number of terms per document does not improve the indication of the content of the document by very much.

However, for Let's Browse, this was not sufficient, and we needed to jack the number of terms up to about 50 to get acceptable results. The reason is that low-frequency terms, while by themselves not very good indicators of content, become very significant if they are shared by your other collaborators, even if they are also low-frequency terms for them. Their commonality increases the value of that term for both of you. Keeping a larger number of terms therefore increases the chance for a larger intersection of interests between collaborators.

Selection of Recommended Pages

The left half of the Let's Browse display is a browser containing the pages currently recommended by Let's Browse (Fig. 3). Whenever the list of current users changes, because a users walks up to or walks away from the workstation, the browsing restarts from a common page, in our case the Media Lab's home. The browsing proceeds by a breadth-first scan from the initial page, filtering through the user profiles.

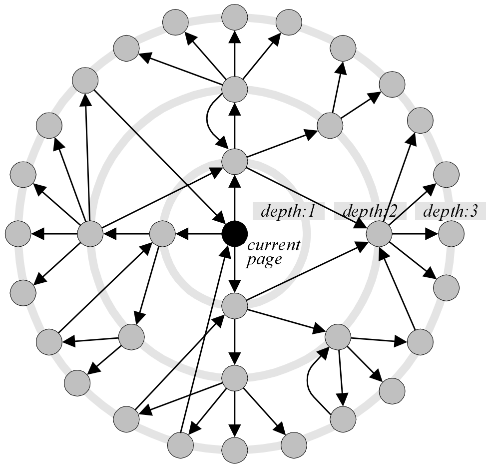

Fig. 3: Breadth-first search of graph of linked Web pages

The kind of browsing performed by Let's Browse might be termed *local reconnaissance*, as opposed to a *global reconnaissance* that we might get if instead we called up a search engine such as AltaVista on the terms of the user profiles and presented the results. Two pages are linked on the Web because some human being thought that someone who looked at one page might want also to look at the other. The "neighborhood" of pages a short distance from the current page constitutes a good approximation to the semantic neighborhood of the current page. This is in stark contrast to the set of pages retrieved by a search engine on a given query, which have no connectivity information whatsoever. Let's Browse shows just those pages in the semantic neighborhood of the current page that show good matches to the profiles of the current set of users.

If the page currently being searched matches the profiles of the users above a certain threshold, it is chosen for display to the user. We also compare the page being considered to the page currently being recommended, if any, so that we don't show the user a page that is significantly worse than the page he or she is seeing already. Recommendations typically came at the rate of between a few a minute and one every few minutes. We buffer up the recommendations and release them at fixed intervals, to create a smoother flow of recommendations to the user.

We used a simple linear combination of the profiles of each user, so that the recommendation was the page that scored the best in the combined profile. There are other possibilities for a recommendation policy. Instead, one could compute the score for each participant independently, then look for the recommendation that scored the best for any of the participants, so that a strong interest on the part of one participant is not discouraged by low interest on the part of the others. Alternatively, participant profiles could "take turns" influencing the recommendations, so as not to allow one participant's interest, or even an average, to dominate the group over a period of time.

Participant Browsing Input

In our experimental setup, we did not have any suitable means for interactive input from the participants. Thus we operated Let's Browse in an entirely passive "automatic channel surfing" mode. The next stage would be to consider interactive browsing input from the participants.

The simplest form of input would be a selection of a link from a currently displayed page. In Letizia, every link selection results in adding the page selected to the user profile. In Let's Browse, the closest analogy would be to add the page to the profile of each user, counting that selection as a joint decision between the participants, to be credited equally among them. If we had multiple pointing devices, we could credit the decision to the person wielding the pointing device, with perhaps a weaker "passive assent" credit on the part of the other participants.

Another possible configuration for a Let's Browse system would be to run independent Let's Browse browsers for each participant, and thus each participant would have control over their own input and profile. The browsers could then get out of sync with each other, but recommendations could still be displayed based on the common profiles. Finally, the most independent possible configuration would be to run completely independent copies of Let's Browse for each participant.

This last possibility would result in another interesting mode of use. By running someone else's profile on your own browsing activity, you are essentially "browsing with another person's eyes." Thus the agent is giving you an idea of what another person might have found interesting. In contrast to another person giving you specific Web sites to look at, this kind of browsing allows you to examine Web sites as the other person might, even if they have never seen the page before. For example, if you are not an expert on cars and are in the market to purchase one, it might be helpful to browse car manufacturer's sites with a profile of a friend who is an expert in the kind of cars you are looking for. This is a new kind of expertise sharing that deserves further exploration.

Visualization of the Recommendation Process

Earlier Web agents such as Letizia merely showed the recommendations without detailed explanation of why recommendations were made. In Let's Browse, we provide a visualization that explains why the agent made the choices it did. This consists of a section that presents the commonalities between the users, followed by a section for each individual user (Fig. 4).

The names of the users are highlighted, and the top few common terms that led to the selection of the paper are noted. For each user, their name, address, e-mail, and a picture gleaned from their home page are displayed. The top few terms from their profiles are displayed, with the terms that are in common with the other participants highlighted.

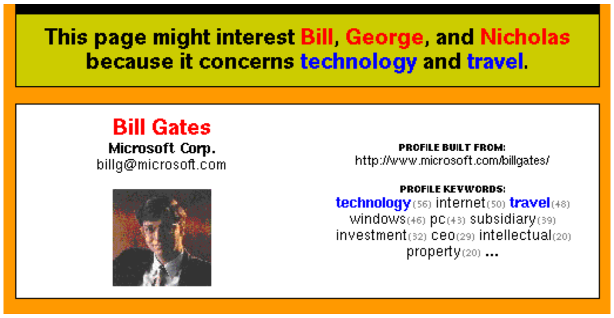

Fig. 4: Exposing recommendation and profiles

RESULTS

While we did not conduct controlled experiments, the participants reacted favorably to our initial experiment. In this conference-like setting, the participants could reasonably be expected to have some commonality of interest, though the exact interest match in any small group was usually not known. The group consisted of a few representatives each from a large number of companies, so large cliques that already knew each other well were rare. Since they were browsing the pages of the Media Lab, which they all had come to visit, some commonality was assured, but which aspects served to attract attention differed considerably from group to group. Thus, the agent served mainly in the function of icebreaker. How it would work for enhancing long-term collaboration between people who already knew each other well remains to be seen.

RELATED WORK

Several projects have been done that allow multiple users to collaborate in browsing manually, either by coupling the user interface of two or more browsers synchronously, or asynchronously by recording and playback of browsing histories. This may be augmented by audio- or video-conferencing, or by textual chat among the participants. ARIADNE [12], GroupWeb [3], and recent features in Lotus Notes™ are all examples of this approach. Webhound [7] used collaborative filtering techniques similar to those used in [4], [13], and Firefly to recommend Web pages that were chosen by other users whose overall tastes in Web pages matched yours, though explicit rating of pages. None of these agents has the kind of incremental, real-time, autonomous operation found in Let's Browse.

Other kind of agents do perform autonomous exploration and incremental learning, but in service of a single user only. A good example is WebWatcher [6], and our own Letizia [8, 9]. Other examples of agents that assist browsing for a single user can be found in [2] and [5].

Silhouettell [10] shares goals with Let's Browse, but differs in approach. Silhouettell uses a Web search engine query based on user-provided keywords to recommend single pages of common interest, whereas Let's Browse uses profiles extracted from users' Web pages to navigate through the Web. An important component of Silhouettell is to use machine vision for identifying users.

ACKNOWLEDGEMENTS

Support for this research was provided in part by grants through the Media Laboratory's Digital Life Consortium, News in the Future Consortium, British Telecom, IBM, and other sponsors of the MIT Media Lab.

REFERENCES

1. Borovoy, R., Martin, F., Vemuri, S., Resnick, M., Silverman, B., Hancock, C., Meme Tags and Community Mirrors: Moving from Conferences to Collaboration, in Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work (CSCW-98), Seattle, November 1998.
2. Bradshaw, J., ed., Software Agents, MIT Press, Cambridge, MA 1997.
3. Greenberg, S. and Roseman, M. GroupWeb: A Web Browser as Real-Time Groupware. ACM Conference on Computers and Human Interface (CHI-96 Conference Companion), 1996.
4. Hill, W., Stead., L., Rosenstein, M., Furnas, G., Recommending and Evaluating Choices in a Virtual Community of Use, ACM Conference on Computers and Human Interaction (CHI-95), 1995.
5. Huhns, M., ed. Readings in Software Agents, Morgan Kauffman, 1997.
6. Joachims, T., D. Freitag, T. Mitchell, WebWatcher: A Tour Guide for the World Wide Web, Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-97), Nagoya, Japan, August 1997.
7. Lashkari, Y. Feature-Guided Automated Collaborative Filtering, MS Thesis, Massachusetts Institute of Technology, Media Arts and Sciences, Sept. 1995.
8. Lieberman, H., Letizia, An Agent for Web Browsing, International Joint Conference on Artificial Intelligence (IJCAI-95), Montreal, August 1995.
9. Lieberman, H., Autonomous Interface Agents, ACM Conference on Computers and Human Interaction (CHI-97), Atlanta, May 1997.
10. Okamoto, M., Nakanishi, H., Ishida, T., Silhouettell: Awareness Support for Real-World Encounter, Community Computing and Support Systems, Lecture Notes in Computer Science, Springer-Verlag, 1998.
11. Sidler, G., Scott, A. and Wolf, H. Collaborative Browsing in the World Wide Web, Eighth Joint European Networking Conference, Edinburgh, May 1997.
12. Twidale, M. and Nichols, D. Collaborative Browsing and Visualization of the Search Process, Electronic Library and Visual Information Research Conference (ELVIRA-96), Milton Keynes, England, May 1996.
13. Wittenburg, K., Das, D., Hill, W., Stead, L., Group Asynchronous Web Browsing on the World Wide Web, Proceedings of the Fourth World Wide Web Conference, December 1995.