

# Adaptive Web Site Agents

Michael J. Pazzani and Daniel Billsus

Department of Information and Computer Science  
University of California, Irvine  
Irvine, CA 92697  
(949) 824-5888

pazzani@ics.uci.edu, dbillsus@ics.uci.edu

## ABSTRACT

We discuss the design of a class of agents that we call adaptive web site agents. The goal of such an agent is to help a user find information at a particular web site, adapting its behavior in response to the actions of the individual user and the actions of other visitors to the web site. The agent recommends related documents to visitors. It integrates and coordinates among different reasons for making recommendations including user preference for subject area, similarity between documents, frequency of citation, frequency of access, and patterns of access by visitors to the web site.

## Keywords

Intelligent Agents; Web Logs; Document Similarity; Learning

## 1. INTRODUCTION

Designers of web sites publish information on the World Wide Web so that visitors to the site may access it. As the site contains more information, it becomes harder for a visitor to locate any item of interest or to explore all the parts of the site that contain pertinent information. In this paper, we argue that a web site should be augmented with an intelligent agent to help the visitors explore the site. We further argue that such agents should learn from the visitors to the web site. Such learning happens by both analyzing web logs to identify common access patterns of the site and by analyzing the visitor's actions to infer the visitor's interests. From the visitor's viewpoint, the agent should help the user make sure that useful information is not overlooked. The web site designer also wishes to increase the amount of useful information accessed by the visitor for a variety of reasons ranging from authors who want their work to have an impact to merchants who want products or advertisements to be considered.

A variety of systems have been proposed to help adapt a web site to visitors. Perkowski and Etzioni, [4] have argued user access patterns can be exploited to synthesize new "index pages" which group together pages commonly accessed together. The Footprints system [8] allows a visitor to a web site to visualize the paths through a web site that are commonly traversed. AVANTI [2] uses a set of adaptation rules that customize the appearance of a web site for "stereotypical" subgroups of users.

Each of these systems uses a single strategy to influence how recommendations are made. Here, we advocate the use of an agent that has multiple strategies for making recommendations. For example, the agent can recommend a document that is on a topic similar to that of other items seen by the user. Alternatively, a document could be recommended that has been accessed frequently by other visitors in combination with documents seen by the user. Furthermore, we argue that an animated agent is a more flexible way of presenting such information than annotating links [4] or synthesizing web pages that summarize this information [5]. We are influenced in our work by Bickmore, Cook, Churchill & Sullivan [1] who have advocated animated autonomous personal representatives, i.e., scripted synthetic characters. "Document Avatars" that interact with readers of a document are one such application of such a representative. However, in our work, we desire for the agent to learn about its user and to learn from patterns of use.

We have deployed two agents with quite different topics and audiences and monitored how users interact with the agents. One at <http://www.ics.uci.edu/~pazzani/Publications> recommends scholarly publications to visitors. The other agent at <http://www.ics.uci.edu/~pazzani/4H> presents information on raising goats and other livestock. The agents share the same engines for analyzing web logs, determining similarity between documents, making recommendations, and learning about users.

Figure 1 shows an example of the adaptive web page agent. Here, the user has viewed a paper, and when the agent detects the user returning to the browser, it recommends that the user download another paper. The agent displays its recommendation in the lower frame, scrolls the upper frame to the appropriate location, and if the user has installed and activated the free 3D Assistant by 3D Planet ([www.3dplanet.com](http://www.3dplanet.com)), the agent gestures to the recommended document and makes a spoken recommendation via a text-to-speech converter.

## 2. ADAPTIVE WEB SITE AGENTS

The goal of the Adaptive Web Site Agent is to assist the user with navigating the web site. When the user views a document (either a web page in HTML or a paper in postscript or PDF), the user can receive recommendations for related documents. In this section we discuss how documents may be related, how a personalized profile is created for a user, and how this profile is used to make recommendations.

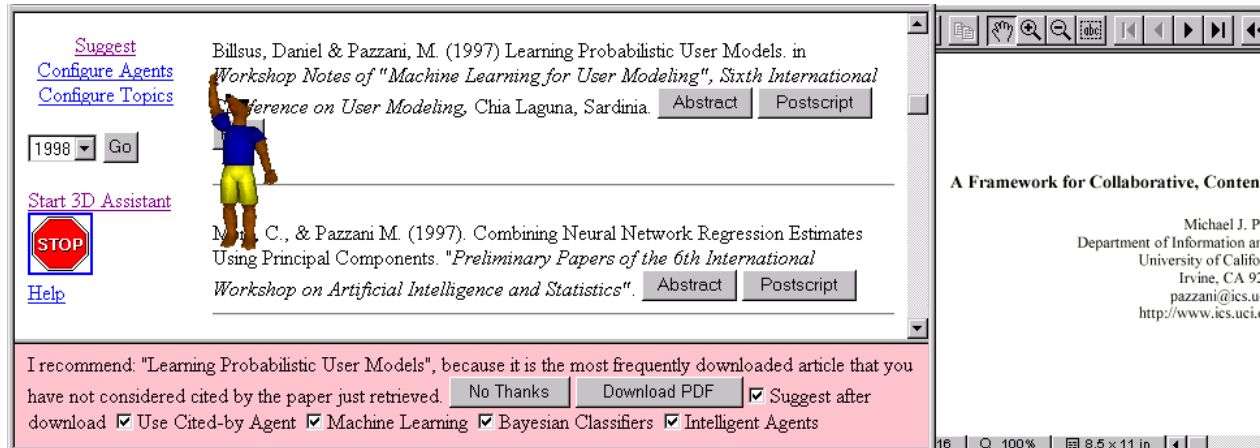


Figure 1. The adaptive web page agent recommending scholarly publications.

There are many ways that two documents may be related. The agent knows about four reasons that one document can be related to another.

1. Similarity: The document is similar to another as determined by comparing the TF-IDF representation [7] of the documents using the cosine similarity metric.
2. Referenced: The document contains a hypertext link (or citation) to another.
3. Referenced-by: There is a hypertext link from another document to this one (or the document is cited by another document). Since the agent has knowledge of the inverse of every link, it can recommend documents that link to the current document.
4. Downloaded-with: The document is frequently accessed in combination with another during the same session. This information is obtained from web logs.

Through analysis of web logs and the HTML of web pages, the agent has access to knowledge of the popularity of various documents in terms of the number of downloads or the number of links (or self-citations) to a document. The agent uses this information to decide upon the "strength" of a recommendation. One of the key issues the agent must deal with is the integration and coordination of possible recommendations for different reasons.

## 2.1 Making recommendations

The agent makes a recommendation to the user when explicitly requested by the user or when focus is returned to the web browser after the user has activated another program or window (e.g., to display a Postscript file) and then returns to the web browser.

To recommend a document to a user, documents that the user has already visited or that have already been recommended are eliminated. The single best document for each of the four reasons is proposed as a possible recommendation. The agent selects one of these by stochastically choosing a single recommendation with probability proportional to the overall strength of the recommendation.

When making a recommendation, the agent can be viewed as a personal representative (cf. [1,6]) of the web site author. The agent infers from the user's actions whether the user followed the

recommendation and increases or decreases the probability that future recommendations are made for that same reason for that user [3]. This allows the system to learn whether a particular user prefers to have documents recommended for a specific reason corresponding to one of the four types of relationships.

The agent can be viewed a success in that the average number of documents viewed on a single visit has increased by 16% in the goat domain and 68% in the publication domain since its installation.

## 3. ACKNOWLEDGMENTS

This research was funded in part by the National Science Foundation grant IRI-9713990 and a gift from Sun Microsystems.

## 4. REFERENCES

- [1] Bickmore, T., Cook, L., Churchill, E., & Sullivan, J. (1998). Animated Autonomous Personal Representatives. In Proceedings of the Second International Conference on Autonomous Agents (Minneapolis, MN), pp 8-15.
- [2] Fink, J., Kobsa, A., & Nill, A. (1996). User-oriented Adaptivity and Adaptability in the AVANTI Project. Designing for the Web: Empirical Studies, Microsoft Usability Group, Redmond (WA).
- [3] Hebb, D.O. (1949). *The organization of behavior*. New York: Wiley.
- [4] Pazzani, M. & Billsus, D. (1997). Learning and Revising User Profiles: The identification of interesting web sites. *Machine Learning*, 27, 313-331.
- [5] Perkowitz, M. & Etzioni, O. (1998). Adaptive Web Sites: Automatically Synthesizing Web Pages. In Proceedings of AAAI98.
- [6] Rist, T., André, E., & Müller, J. (1997). Adding Animated Presentation Agents to the Interface. In Proceedings of the 1997 International Conference on Intelligent User Interfaces, pp 79-86, Orlando, Florida.
- [7] Salton, G. (1989). *Automatic Text Processing*. Addison-Wesley.
- [8] Wexelblat, A. & Maes, P. (1997). Using History to Assist Information Browsing. *RIA0'97: Computer-Assisted Information Retrieval on the Internet*, Montreal.