



Adaptive Web Site Agents

MICHAEL J. PAZZANI AND DANIEL BILLSUS

{pazzani, dbillsus}@ics.uci.edu

Department of Information and Computer Science, University of California, Irvine, Irvine, CA 92697

Abstract. We discuss the design and evaluation of a class of agents that we call adaptive web site agents. The goal of such an agent is to help a user find additional 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 and we show that these recommendations result in increased information read at the site. 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. We argue that this information is best used not to change the structure or content of the web site but rather to change the behavior of an animated agent that assists the user.

Keywords: recommendation systems, personalization

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 [11] to help the visitors explore the site. We further argue that such agents should learn from the visitors to the web site. Our contribution is the combination of methods that operate 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 their products or advertisements to be considered.

A variety of systems have been proposed to help adapt a web site to visitors. Some commercial web sites (e.g., Amazon.com) recommend additional purchases to the customer based upon the purchasing patterns of previous customers. WebWatcher [1, 9] assists a user browsing the Web by learning to suggest which links to traverse. Perkowitz and Etzioni, [15, 16, 17] have argued for adaptive web sites and illustrated this concept by creating a system that learns from user access patterns to synthesize new “index pages” which group together pages commonly accessed together. The

Footprints system [23] allows a visitor to a web site to visualize the paths through a web site that are commonly traversed. AVANTI [7] uses a set of adaptation rules which 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 or a document 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 (e.g., [1, 14]) or synthesizing web pages that summarize this information [17]. We are influenced in our work by Bickmore, Cook, Churchill, and Sullivan [3] who have advocated animated autonomous personal representatives, i.e., scripted synthetic characters. “Document Avatars” who interact with readers of a document are one application of such a representative. However, in our work, we desire for the agent to learn about each visitor and to learn from patterns of use.

We have deployed two adaptive web site agents with quite different topics and audiences, and we have monitored (via web logs) how users interact with the agents. One agent located 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. Although deployed in different domains, the agents share the same engines for analyzing web logs, determining similarity between documents, making recommendations and learning about user preferences.

Figure 1 shows one example of the adaptive web page agent. In this scenario, the user has downloaded a PDF version of a paper and then returned to the browser window. The agent detects the user returning to the browser and 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), the agent gestures to the recommended document and makes a spoken recommendation via a text-to-speech converter.

One goal of this paper is to argue that an adaptive agent that serves as a guide to a web site avoids problems associated with adapting the structure and content of the web site. A second goal is to explore how multiple strategies might be integrated in a single recommendation system. In the remainder of this paper we first discuss the architecture of the Adaptive Web Site Agent describing the individual components and how they are integrated into the system. We demonstrate the effectiveness of the agent by showing that adding the agent to the web sites increased the average number of documents accessed on each visit to a web site. Furthermore, we compared the Adaptive Web Site Agent to a strawman agent that makes random recommendations to show that it is the quality of the recommendations and not just the act of recommending that increases the amount of information accessed. Next, we discuss the relationship between Adaptive Web Site Agents and Adaptive Web Sites. Finally, we discuss some of the implementation details and issues that we faced.

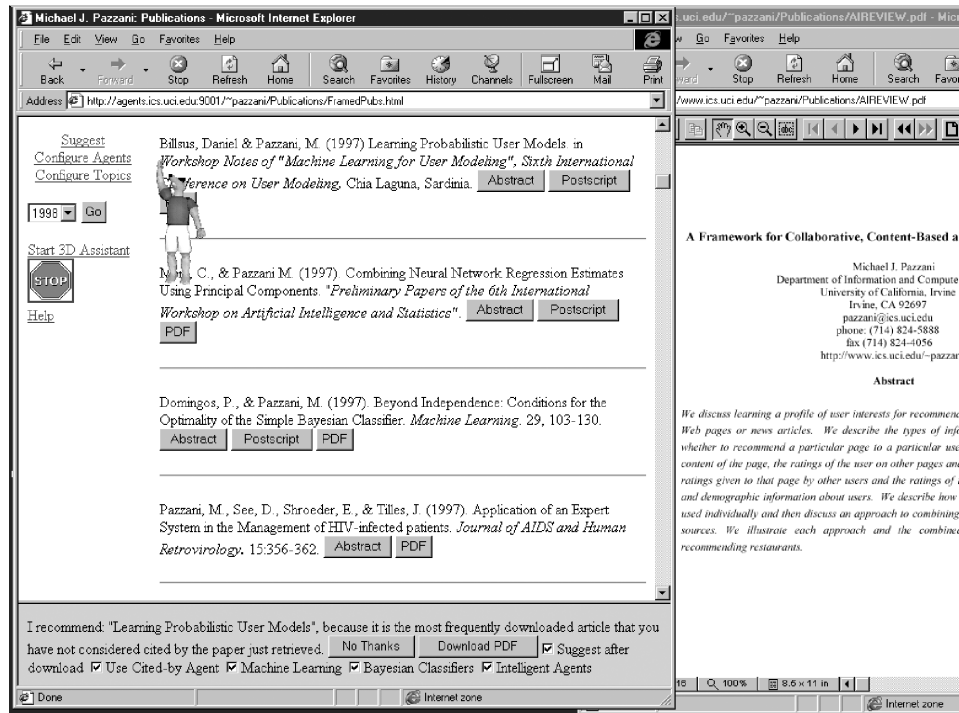


Figure 1. The adaptive web page agent recommending scholarly publications.

2. Adaptive Web Site Agents

The goal of the Adaptive Web Site Agent is to assist the user with navigating a web site. When the user views a document (either a web page in HTML or a paper in postscript or PDF), the user can ask for recommendations of 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.

2.1. Relationships between documents

There are many ways that two documents may be related. Some of these relationships may be stronger than others. For example, D_1 may be very similar to D_2 using the cosine similarity metric while D_3 may be a frequently downloaded document with a hypertext link from D_1 . To allow the agent to reason about the strength of each recommendation, they are normalized to a $[0, 1]$ scale where the strongest relationship of each type in the site is given a weight of 1. The notation $s_r(D_1, D_2)$ is used to represent the strength for each relationship r between document D_1 and document D_2 . 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 [20] of the documents using the cosine similarity metric. The strength of the relationship between D_1 and D_2 is given by the cosine similarity metric between D_1 and D_2 divided by the maximum cosine similarity at the site. Although cosine similarity is already on a $[0, 1]$ scale, the maximum is typically much less than 1. For example, it is 0.47 at the goat site.
2. *Referenced*: The document contains a hypertext link (or citation) to another. When there is hypertext link from the current document D_1 to another D_2 , the strength of the relation is given by the number of hyperlinks (or citations) to D_2 divided by the number of hyperlinks to the most frequently linked document at the site. The idea here in the publication domain is to favor the most frequently cited article,¹ although the intuition applies equally well to other domains.
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. When there is hypertext link to the current document D_1 from another D_2 , the strength of the relation is given by the frequency with which the document D_2 has been downloaded divided by the frequency of the most downloaded document. This same type of information is used by the PageRank algorithm of Google [2] to measure the relative importance of web pages.
4. *Downloaded-with*: The document is frequently accessed in combination with another during the same session. This information is obtained from web logs. We follow Perkowitz and Etzioni [17] by estimating $P(D_2|D_1)$ from the log data (i.e., the probability that document D_2 is downloaded given that D_1 has been downloaded). The strength of the relationship is calculated by $P(D_2|D_1)$ divided by the maximum $P(D_i|D_1)$ for all documents.

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. Finally, the agent optionally knows about the topic of a document by the author choosing any number of subject terms from a set defined by the author. For example, in the publication site, the topics are items such as “machine learning” and “intelligent agents.” In the goat site, the topics include “birth” and “goat health.”

The agent uses this information to decide upon a document (or documents) to recommend. One of the key issues the agent must deal with is the integration and coordination of possible recommendations for different reasons. The solution implemented in the Adaptive Web Site Agent is described in Section 2.2

2.2. Learning personalized profiles

The Adaptive Web Site Agent allows for the possibility that different users may have different preferences for sources of recommendations. For example, in the publication domain some users may prefer to have the agent recommend papers that cite

a recently downloaded paper. Others might prefer to have the agent recommend papers that are cited by a recently downloaded paper. To permit this sort of customization to the individual user, we associate a weight, w_r with each type of relationship r . The weight of each type of relationship is combined with the strength of the individual relationships to come up with a combined overall weight $o(r, D_1, D_2)$ for recommending that document D_2 be viewed because it is related to document D_1 by relationship r .

$$o(r, D_1, D_2) = w_r s_r(D_1, D_2) \quad (1)$$

The strength component $s_r(D_1, D_2)$ is identical for all users while the weight w_r is personalized for each user. The weight is learned by a combined Hebbian (Hebb, 1949) and anti-Hebbian (e.g., [21]) learning rule: each time a recommendation is followed, the weight is increased by a constant factor, and when a recommendation is not followed it is reduced by a constant factor. In particular, when a recommendation is followed, the new weight becomes $(1 + \epsilon)w_r$ and when ignored $w_r = (1 - \epsilon)w_r$. We use a value 0.25 for ϵ in the Adaptive Web Site Agent.

The Adaptive Web Site Agent also allows for the possibility that different users may have preferences for different subjects. For example, some visitors to the publication page might be interested in medical applications of artificial intelligence and others might be interested in intelligent agents. The designer of a site using the Adaptive Web Site Agent may optionally tag documents with a set of subject terms defined by the site designer. Weights for these subject terms are learned in exactly the same manner as the relationship weights: when a recommended document is visited, the subject weights are increased, and when a recommendation is ignored, the weights are reduced.² Empirically, we have found a wide variation in weights across users in the publication domain, suggesting that the learning is helping people with different interests to find different documents.

The profile for a user consists of two sets of weights, one indicates the preferences for the source of recommendations and the other for subjects. These are stored in a “cookie” so that training on one visit affects the next visit. This cookie also stores the names of pages that the user has downloaded. The publication agent uses this information to recommend new publications related to ones that the user has downloaded and to avoid recommending something that the user has already seen.

2.3. Making recommendations

First, we concentrate on recommending a single document to a user. This process first eliminates a number of documents from further consideration and then looks for relationships between the most recently visited document and the other documents. Three classes of documents are eliminated from consideration:

- Documents that have already been recommended to the user in the current visit.
- Documents already visited by the user. In the publication domain, we assume that a single user will only want to download a paper once. In the goat domain,

users frequently visit for reference or entertainment purposes and the agent only keeps track of the documents visited in any one visit.

- Documents on subjects that are not interesting to the user. A document is eliminated unless it has at least one subject term with a weight above a threshold (0.5). If there are no documents that meet this criterion (or no subject terms are provided), this constraint is ignored.

In the current implementation, the single best document for each of the four reasons is proposed as a possible recommendation.³ The agent selects one of these by stochastically choosing a single recommendation with weight proportional to the overall strength of the recommendation. In particular, the probability of recommending document D_j recommended for reason r is:

$$p(D_j) = \frac{o(r, D_j, D_1)}{\sum_i o(r_i, D_i, D_1)} \quad (2)$$

The selection process allows a trade-off between always accepting the recommendation with the highest overall strength and having the agent learn about the user by presenting alternatives. This strategy prevents the agent from making the same recommendations to every new user and strongly biasing future recommendations to be only the most popular previous items. However, by taking the overall strength of the recommendation into account, the agent avoids making frequent recommendations that are likely to be ignored by the user.

There are three situations in which the agent may make a recommendation to the user:

1. When explicitly requested by the user.
2. When focus is returned to the web browser. This occurs when the user has activated another program or window (e.g., to display a Postscript file) and then returns to the web browser. Currently, this capability is used in the publication domain but not the goat domain, since in the goat domain there is no use of an auxiliary program to display information from the web site.
3. When the visitor has been idle for a few minutes. If two “idle” recommendations in a row are ignored, the agent turns off this capability, assuming that the user is no longer using the computer.

Finally, there is a special case in the publication domain for recommending a publication when no publication has been downloaded yet. First, a check is made to determine whether any new papers have been added since the user’s last visit. This is accomplished by storing the date of visit in a cookie. If a paper has been added that is cited by a paper previously downloaded by the user, then the most frequently accessed new paper cited by a previously downloaded paper is recommended. Otherwise, the agent stochastically selects among the most frequently accessed paper or the most frequently referenced paper as determined by analysis of the web log. In the goat domain, this special case is not needed since the agent treats the goat home page like any other page, while the publication home page differs from other publications.

When making a recommendation, the agent can be viewed as a personal representative (cf. [3, 19]) of the web site author. The visitor requests recommendations and the agent makes them. If they are not accepted, the agent makes a different recommendation (cf. [6]) and learns about the preferences of the user.

So far, we have discussed the agent making a single recommendation. The information collected by the agent may also be used to present alternatives to the user. In particular, the overall strength can be used to order possible recommendations. In Figure 2, there is a selection list that the visitor may use to jump to other pages. The items in this selection box are reordered by the agent by the overall strength so that the most related ones are first. This reordering happens automatically as each web page in the site is visited.

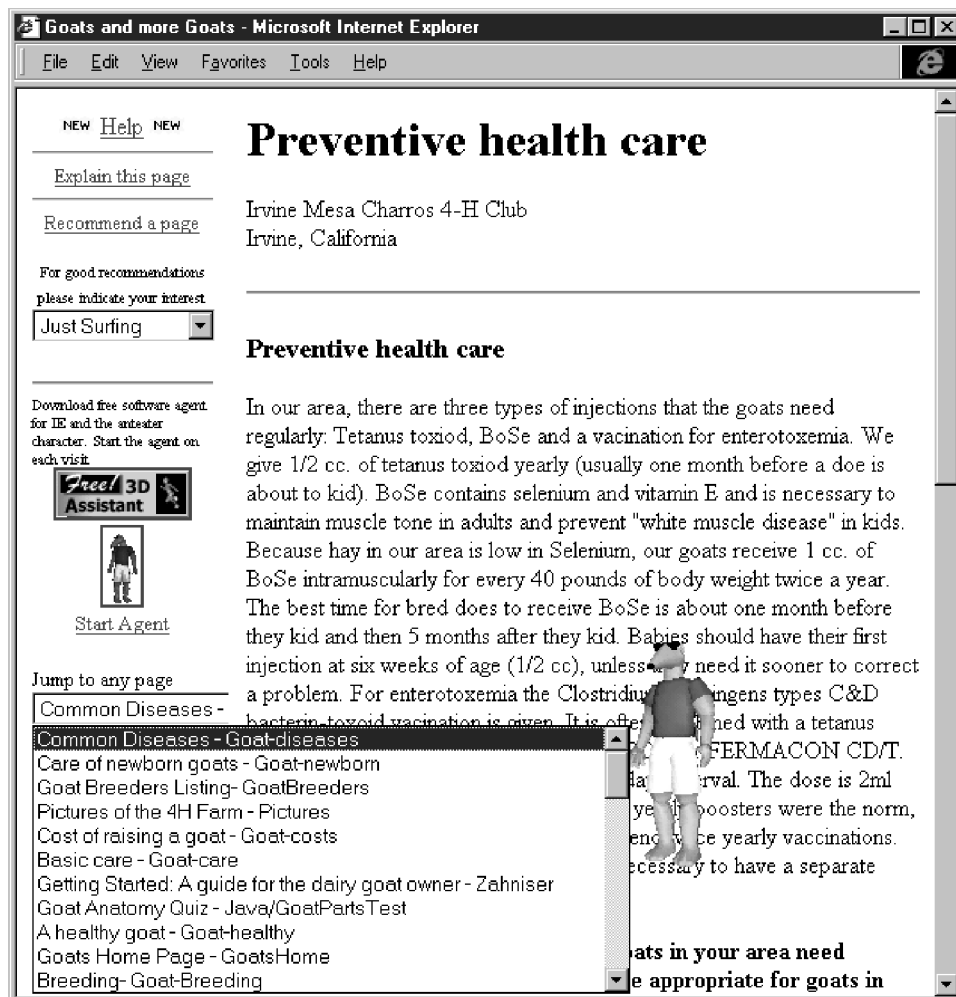


Figure 2. The adaptive web page agent recommends a single page or reorders a selection box that allows the user to jump to other related pages.

2.4. Evaluation

The Adaptive Web Page Agent is intended to find information at a web site. It recommends additional information to the visitor and gives the visitor the option of seeing that information. We have indirect evidence that the agent is beneficial. In particular, analysis of the log data determine whether the addition of the agent to the web site increases the amount of information accessed by the visitor. The idea is that the agent makes recommendations for related pages, and if more information is accessed, we can infer that the average user follows the recommendations. This measure would be of interest to an author interested in having their work read more often, an advertiser interested in having more pages with ads viewed, or a merchant interested in having more products considered. We collected data on the number of items viewed by users for 60 days prior to adding the intelligent agent to each site and 120 days after adding each agent. Table 1 shows the mean number of pages visited by a user to each web site before and after deploying the agent. A two-tailed *t*-test, a commonly used to test to determine if the mean of two sample conditions is different, indicates the difference in each domain is significant at the .0001 level. Although both results are statistically significant, the effect of the agent is more substantial in the publication domain. We speculate that this occurs because the agent is more “aggressive” in the publication domain, and makes recommendations when focus is returned to the web browser after downloading a publication. In contrast, in the goat domain the agents must wait for a user to ask for a recommendation because there is no easy way to detect when the user is finished with the current web page.

The previous experiment demonstrates that the Adaptive Web Site Agent increases the amount of information accessed by users. However, it could be the case that the people accept any recommendation or play with the system to test its abilities. To rule out this possibility, we modified the publication agent so that on odd days it recommended a randomly selected publication that had not been read by the user and on even days the agent as described in this paper was used. On all days, the explanations for the recommendation was changed to a generic explanations (“because it is related to the paper just downloaded”), so that users could not tell from the explanation whether a sensible or random recommendation was made.⁴ Note that in this domain (in contrast to a larger domain such as recommending books at Amazon) even the random recommendations are somewhat related since they are publications by the same author. We collected data for 64 days and compared the number of articles downloaded on odd and even days. On odd days users downloaded an average of 3.05 articles and on even days

Table 1. Average number of documents downloaded by each visitor to web site before the agent was installed and the same web site with the adaptive agent

Site	Before agent	With agent
Publications	2.41	4.77
Goats	2.28	2.81

users downloaded 4.48. The difference is significant at the 0.01 level. Note that the average number of articles downloaded with random recommendations is slightly higher than the earlier experiment with no recommendations. Nonetheless, there is a significant and substantial increase between the random and adaptive recommendations. This suggests that it is the quality of the recommendation and not just the act of recommending that accounts for much of the increase in items downloaded.

3. Discussion

Care has been taken in designing the Adaptive Web Site Agent to make sure that the algorithms could scale to larger problems. The learning algorithm used is incremental and does not require re-analysis of previous recommendation when new recommendations are made. The similarity between documents uses an approach that is also used in many WWW search engines. The analysis of web logs is also done with a single scan through the web log to find the probability that one document is downloaded given that another has been downloaded in the same session. However, it would be difficult for some of the recommendation strategies to be used in a system that accesses many web sites since the web log analysis is done on a single site and for efficiency purposes, document similarity is precomputed.

Currently, the analysis of web logs is done to find the probability that one document is downloaded given that another document has been. Our future plans include treating this as a “market basket analysis” problem (e.g., [12]) where the system would recommend a new document based upon all documents downloaded during this visit. Such a technique requires estimating joint probabilities from data, and the current web logs we have collected are not quite large enough (approximately 50,000 visits to the goat site and 500 to the publication site) to allow for reliable estimation of these probabilities.

The profile of a user is relatively small so that it may be stored in a cookie. It consists of a unique identifier for each document the user has downloaded, a set of weights representing the relative importance of each reason for making a recommendation, a set of weights for subject terms. The simplicity of the model may cause problems if the users have different goals at different times. For example, a new graduate student taking a course in machine learning may prefer to follow recommendations for papers on that topic and may prefer to follow references to find earlier foundation documents. The next semester, the same visitor taking an AI and medicine class might prefer papers on that topic and may prefer using the referenced-by relation to find the latest developments. In spite of this potential problem, we have shown that agents does indeed satisfy its goal of recommending papers that are then downloaded by visitors.

The Adaptive Web Site Agent proposes a solution to a problem raised by Terveen and Hill [22]: making a recommendation when coordinating among multiple recommendation strategies. Here, we give a strength to the recommendation made by each strategy and then pick the recommendations with a probability proportional to the strength of the recommendation. Rather than having a single predefined weight

for each strategy, we have used a simple neural network learning rule to determine the relative usefulness of each strategy for the individual user. This is some ways like the combination function learned when combining multiple hypotheses in ensemble learning algorithms (e.g., [13]) with the exception that we incrementally learn the weights rather than analyzing a collection of examples at one time.

4. Related work

The Adaptive Web Site Agent is addressing a similar problem to Adaptive Web Sites, i.e., “web sites that automatically improve their organization and presentation from user access patterns” [15, p. 727]. Perkowitz and Etzioni [17] consider one subproblem of adaptive web sites: synthesizing new index pages. The idea is that the designers of a web site cannot always anticipate what groups of pages users will want to access. For example, in a recipe site, the designers might have organized information one way (e.g., by type of cuisine), while visitors might also like to see information organized in other ways (e.g., low-fat main dishes, vegetarian foods, or spicy foods). By finding clusters of items commonly accessed together, the PageGather algorithm proposes items to put on new index pages. While this is a useful start, it leaves a number of questions unanswered. In particular, there is no mechanism for automatically integrating the new index pages into the web site, or labeling the index page so that visitors will find it and understand its purpose. Even if these tasks are done manually,⁵ there are some limitations of PageGather. In particular, the decision as to whether or not to include an item on the index page is made by comparing the probability that one page is downloaded with another to a single global threshold. It does not allow for degrees of category membership (e.g., spicy foods). Furthermore, it is unlikely that a single threshold would find all pages that are conceptually grouped together on different dimensions (e.g., all vegetarian recipes or all desserts).

The Adaptive Web Site Agent does not attempt to improve the site structure. Instead, the information culled from web logs is used to improve its ability to recommend related documents to the visitor. Rather than statically grouping related items together, it dynamically recommends which document to consider next or it sorts documents according to how related they are to the current document. As a consequence, it allows for varying degrees of relatedness between documents and avoids the issues concerned with finding a label for a group of documents and integrating a new page into the existing structure of a web site. Instead, it serves as a personal representative of the web site author guiding the user on a tour of the web site. Users are likely to assume that an index page at an Adaptive Web Site such as one listing sports utility vehicles contains all such vehicles and only such vehicles. Unfortunately, learning technology provides no such guarantee. In contrast, the Adaptive Web Site Agent need not be recommend every relevant article or only relevant articles to be considered useful by user who are familiar with human recommenders such as concierges, friend and spouses who are usually right.

The Adaptive Web Site Agent is in some ways similar to WebWatcher [1, 9] a system that learns what links to traverse to help a user achieve an explicitly stated goal. WebWatcher follows a user as the user browses the entire web and recommends links to follow to help achieve the user's goal. By restricting its attention to a single site, the Adaptive Web Site Agent has the ability to use sources of information unavailable to WebWatcher to make recommendations. In particular, the Adaptive Web Site Agent will make recommendations based upon similarity between pages at a site and the analysis of web logs at the site.

Adaptive Web Site Agent also shares some goals with the Remembrance Agent [18] and Watson [4]. Both of those systems recommend other information sources related to the current document focus of the user. Both systems use the words near the cursor of a user editing a document as query terms for a similarity-based search engine. The Remembrance Agent automatically locates documents within the file system of the user related to the current document (e.g., an email message sent to the user related to the paper that the user is revising). Watson constructs a query to search engines to find documents anywhere on the WWW. Neither of these systems personalizes its recommendations to the user based on the user's interactions with the system.

5. Implementation and interface

The Adaptive Web Site Agent uses 3D Assistant by 3D Planet to animate a character that makes recommendations to the visitor. The 3D Assistant has the ability to highlight locations on the screen by gesturing at them and to make verbal recommendations via a text to speech converter. Because the 3D Assistant cannot reference the location of specific buttons on the screen, we scroll the recommended document to the top of the web page and gesture at this location. This scrolling action is also useful for those users who do not have the required software to run the agent. To use 3D Assistant, a visitor must download and install a free program from the 3D Planet web site. The program is available only under Windows 95/98/NT for Internet Explorer and Netscape Navigator. However, users access these sites from a variety of platforms including Unix (publications) and WebTV (goats). As a consequence, many users miss out on the animated character but instead interact with the agent via dialog boxes or frames that display information for visitors. Nonetheless, we claim that the animated character presents a useful programming metaphor for developing intelligent agents. The design of the Adaptive Web Site Agent was guided by considering what an intelligent person who knew the content, structure, and access patterns of the web site would recommend to the user. This design allows for both improvement of the agent from analyzing logs of all web users and adaptation of the agent to the individual user.

The adaptation of the agent to individual users happens primarily via implicit feedback over whether the user accepts the recommendation of the agent. However, in the publication domain, where we expect the average user to be more technologically sophisticated than the user in the goat domain, we also allow for

customization by turning on or off various reasons for recommending documents (see Figure 3) or various document subjects (see Figure 4).

An additional opportunity for configuration presents itself when the agent makes a recommendation. In this case, the reasons that entered into this particular recommendation are displayed to the user together with the recommendation (see Figure 5). This allows for immediate customization to prevent this type of recommendation in the future.

Although these check boxes provide customization of the Agents behavior by the user, few users bother to configure the agent using these check boxes and accept the defaults. Since this customization functionality is rarely used, it might better be moved to a configuration page rather than being included with each recommendation.

The goat web site exploits another capability of the 3D Assistant. It contains an "Explain This Page" button which causes the agent to read a short overview of the page. Although it would be possible to do this on the publication web site, the feature was not implemented because we assumed that visitors interested in scholarly publications would not want an animated anteater (the mascot of UCI) to read the abstract of papers.




Figure 3. Optionally, the visitor may customize the reasons a document is recommended.




Figure 4. Optionally, the visitor may configure the subjects of documents recommended.




Figure 5. Making a recommendation is an opportunity to allow for configuring the reasons that recommendation occurred.

6. Conclusion

The Adaptive Web Site Agent serves as a personal representative of the web site designer. The goal of the agent is to help the visitor find information at the web site by recommending related documents to the user. The primary contribution of the research is the combination of recommendation methods that exploit a variety of ways in which documents may be related. Experience with the Adaptive Web Site Agent suggests that it helps users find personally relevant information in that it increases the amount of information accessed by the average visitor to a web site.

Notes

1. Since the agent collects information from a single site, it counts the frequency of self-citations to determine whether to recommend a paper. Integration with an agent such as CiteSeer [5] would allow the agent to count the frequency of all citations.
2. Note that the Adaptive Web Site Agent does not require explicit feedback on whether or not the user likes a particular document. Like Letizia [10], it infers from the user's actions whether the user follows its recommendations. This is accomplished by keeping track of recommendations and comparing them to the documents accessed. The "No Thanks" button on Figure 1 is used to clear the recommendation frame. Negative feedback occurs whenever the user ignores the recommendation of the agent. Viewing the abstract of the paper is not treated as evidence for or against the recommendation of the agent.
3. Since there is some correlation between the various sources of recommendations (e.g., an article is often more similar to an article it cites than other articles), a document may be proposed for more than one reason. In this case, its likelihood of being selected is increased.
4. We decided to restrict this experiment to the publication domain because we felt the users of this site would be more tolerant of random recommendations than the users of the goat site who have less computer experience.
5. Manual integration of discovered clusters into the web structure has a drawback in that few web site designers would be willing to continuously update the web page in response to changes in site access patterns.

References

1. R. Armstrong, D. Freitag, T. Joachims, and T. Mitchell, "WebWatcher: A learning apprentice for the World Wide Web," in *Proc. 1995 AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments*, Stanford, AAAI Press, 1995.
2. S. Brin and L. Page, "The anatomy of a large-scale hypertextual Web search engine," in *7th Int. World Wide Web Conference*, Brisbane, Australia, 1998, pp. 107–117.
3. T. Bickmore, L. Cook, E. Churchill, and J. Sullivan, "Animated autonomous personal representatives," in *Proc. Second Int. Conf. Autonomous Agents*, Minneapolis, MN, 1998, pp. 8–15.
4. J. Budzik, K. J. Hammond, C. Marlow, and A. Scheinkman, "Anticipating information needs: Everyday applications as interfaces to Internet information sources," in *Proc. 1998 World Conference on the WWW, Internet, and Intranet*, AACE Press, 1998.
5. K. Bollacker, S. Lawrence, and C. L. Giles, "CiteSeer: An autonomous Web Agent for automatic retrieval and identification of interesting publications," in *Proc. 2nd Int. ACM Conf. Autonomous Agents*, 1998, pp. 116–123.
6. R. Burke, K. Hammond, and B. Young, "The FindMe approach to assisted browsing," *IEEE Expert/Intelligent Systems & Their Applications*, vol. 12, no. 4, pp. 32–40, 1997.

7. J. Fink, A. Kobsa, and A. Nill, "User-oriented adaptivity and adaptability in the AVANTI project," *Designing for the Web: Empirical Studies*, Microsoft Usability Group: Redmond, WA, 1996.
8. D. O. Hebb, *The Organization of Behavior*, New York: Wiley, 1949.
9. T. Joachims, D. Freitag, and T. Mitchell, "WebWatcher: A tour guide for the World Wide Web," in *Fifteenth Int. Joint Conf. Artificial Intelligence*, 1997.
10. H. Lieberman, "Letizia: An agent that assists web browsing," in *Proc. Int. Joint Conf. Artif. Intell.* Montreal, 1995.
11. P. Maes, "Agents that reduce work and information overload," *Commun. ACM*, vol. 37, no. 7, pp. 31–40, 1994.
12. N. Megiddo and R. Srikant, "Discovering predictive association rules," in *Proc. Fourth Int. Conf. Knowledge Discovery in Databases and Data Mining*, New York, 1998.
13. C. Merz and M. Pazzani, "A principal components approach to combining regression estimates," *Machine Learning*, vol. 36, pp. 9–32, 1999.
14. M. Pazzani and D. Billsus, "Learning and revising user profiles: The identification of interesting web sites," *Machine Learning*, vol. 27, pp. 313–331, 1997.
15. M. Perkowitz and O. Etzioni, "Adaptive web sites: An AI challenge," in *Proc. IJCAI97*, 1997a.
16. M. Perkowitz and O. Etzioni, "Adaptive web sites: Automatically learning from user access patterns," in *Proc. WWW6*, 1997b.
17. M. Perkowitz and O. Etzioni, "Adaptive web sites: Automatically synthesizing web pages," in *Proc. AAAI98*, 1998.
18. B. Rhodes and T. Starner, "Remembrance agent: A continuously running automated information retrieval system," in *Proc. First Int. Conf. The Practical Application of Intelligent Agents and Multi Agent Technology*, 1996, pp. 487–495.
19. T. Rist, E. André, and J. Müller, "Adding animated presentation agents to the interface," in *Proc. 1997 Int. Conf. Intell. User Interfaces*, Orlando, FL, 1997, pp. 79–86.
20. G. Salton, *Automatic Text Processing*, Addison-Wesley, 1989.
21. N. Schraudolph and T. J. Sejnowski, "Competitive anti-Hebbian learning of invariants," in *Advances in Neural Information Processing Systems*, vol. 4, pp. 1017–1024, Morgan Kaufmann: San Mateo, 1992.
22. L. G. Terveen and W. Hill, "Human-computer collaboration in recommender systems," in J. Carroll (ed.), *HCI in the New Millennium*, Addison Wesley, in press.
23. A. Wexelblat and P. Maes, *Using History to Assist Information Browsing*, RIAO'97: Computer-Assisted Information Retrieval on the Internet, Montreal, 1997.