	Autonomy Press Release, at 1 Autonomy Technology Whitepaper, at AUT00069-70 Autonomy Agentware User Guide, at AUT00002 **MONTEBELLO** at 3, 4.
d) analyzing a document d to identify properties of the document;	JOACHIMS at 3: "What is the form of the knowledge required by WebWatcher? In general, its task is to suggest an appropriate link given an interest and Web page. In other words, it requires knowledge of the following target function: *LinkQuality : Page x Interest x Link* → [0; 1] The value of LinkQuality is interpreted as the probability that a user will select Link given the current Page and Interest. In the following we present three approaches to learning this target function from experience. The first approach uses previously given tours as a source of information to augment the internal representation of each selected hyperlink. The second approach is based on reinforcement learning. The idea is to find tours through the Web so that the amount of relevant information encountered over the trajectory is maximized. The third approach is the combined method that includes both of the first two approaches." *See generally* p. 4. *See also*: **CULLISS** at 2:43-46; 5:36-48; 10:47-52. **MLADENIC**, at 4, 12. **REFUAH** at 7:53 – 8:6; 20:31-37; 21:6-30. **SCHUETZE** at 3:40-44; 5:59-63; 10:20-31; Fig.1. *See generally* 11:42 – 15:19.

	WASFI at 61. Autonomy Press Release, at 1 Autonomy Technology Whitepaper, at AUT00068-69, AUT00071 Autonomy Agentware User Guide, at AUT00002, AUT00117 **MONTEBELLO** at 3, 4.
e) estimating a probability P(u|d) that an unseen document d is of interest to the user u, wherein the probability P(u|d) is estimated by applying the identified properties of the document to the learning machine having the parameters defined by the User Model; and	JOACHIMS at 3: "What is the form of the knowledge required by WebWatcher? In general, its task is to suggest an appropriate link given an interest and Web page. In other words, it requires knowledge of the following target function: *LinkQuality : Page x Interest x Link* → [0; 1] The value of LinkQuality is interpreted as the probability that a user will select Link given the current Page and Interest. In the following we present three approaches to learning this target function from experience. The first approach uses previously given tours as a source of information to augment the internal representation of each selected hyperlink. The second approach is based on reinforcement learning. The idea is to find tours through the Web so that the amount of relevant information encountered over the trajectory is maximized. The third approach is the combined method that includes both of the first two approaches." 3: "To suggest hyperlinks during a tour WebWatcher compares the current user's interest with the descriptions of all hyperlinks on the current page. WebWatcher suggests those hyperlinks which have a description sufficiently similar to the user's interest. The metric used to compute similarity between a user's stated interest and a hyperlink description is based on a technique from information retrieval [Salton, 1991]. Interests and hyperlink descriptions are represented by very high-dimensional feature vectors, each dimension representing a particular word in the English language. The elements (called word-weights) of a vector are calculated using the TFIDF heuristic [Salton, 1991]. Based on this vector representation similarity is calculated as the cosine between vectors."

	See generally p. 4. *See also*: **CULLISS** at 2:43-51; 5:36-62. **MLADENIC** at 5, 7, 10, 12, Table 2. **REFUAH** at 3:56 – 4:4; 17:32 – 18:4. **SCHUETZE** at 18:62 – 19:10. **WASFI** at 60, 61. Autonomy Press Release, at 1-2 Autonomy Technology Whitepaper, at AUT00069, AUT00071 Autonomy Agentware User Guide, at AUT00119 **MONTEBELLO** at 4.
f) using the estimated probability to provide automatic, personalized information services to the user.	JOACHIMS at 2: "We can tell that WebWatcher accompanies us from the additions it makes to the original page. A page augmented with WebWatcher's additions is shown in figure 1. Among others they include: ☼ A list of WebWatcher Commands is inserted above the original page. The user can invoke these commands to communicate with WebWatcher as he browses from page to page. ☼ Selected hyperlinks from the original page have now been highlighted by WebWatcher, to suggest directions relevant to our browsing interests. WebWatcher highlights these hyperlinks by inserting eyeball icons () around the hyperlink, as shown in the 'projects' link in figure 1 (recall in this example the user expressed an interest in 'intelligent agents'). The advice WebWatcher provides in

	this fashion is based on knowledge learned from previous tours." 2: "In addition to highlighting hyperlinks WebWatcher also provides other forms of assistance. In particular, it provides a keyword search using a variant of the Lycos search engine applied to the set of pages previously visited by WebWatcher."

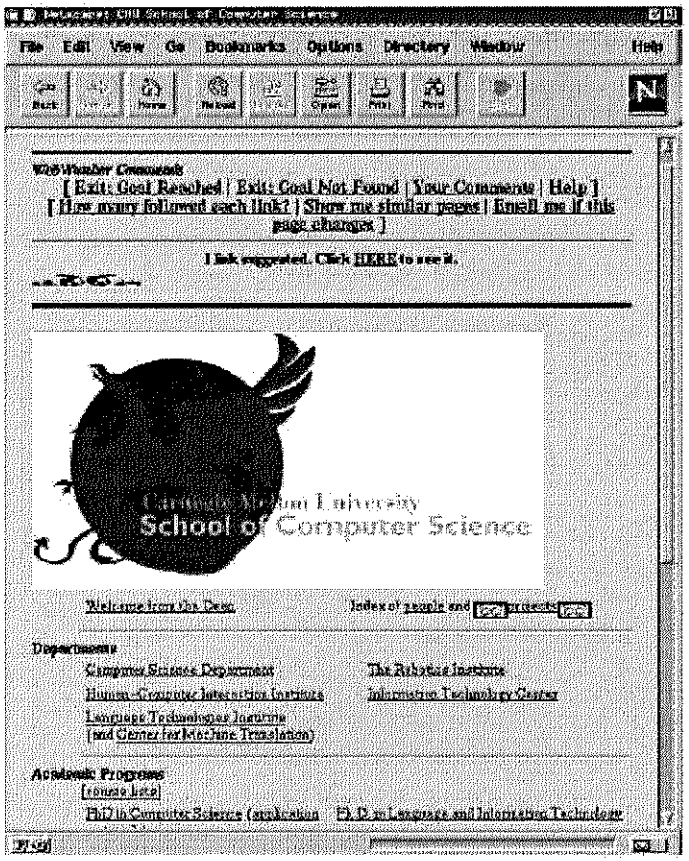

Figure 1: The Front-Door page with WebWatcher's additions.

See also:

CULLISS at 2:39-51.

	MLADENIC at 2, 12. **REFUAH** at Abstract, 2:63 – 3:11; 3:47-55; 17:21-43; 18:56-65; 23:11-28. **SCHUETZE** at 1:29-33; 7:54-60. **WASFI** at 61. Autonomy Press Release, at 1 Autonomy Technology Whitepaper, at AUT00068 Autonomy Agentware User Guide, at AUT00002, AUT00004, AUT00116-117 **MONTEBELLO** at 4.
Claim 11	
11. The method of claim **1** further comprising estimating a posterior probability P(u|d,q) that the document d is of interest to the user u, given a query q submitted by the user.	JOACHIMS at 2: "Selected hyperlinks from the original page have now been highlighted by WebWatcher, to suggest directions relevant to our browsing interests. WebWatcher highlights these hyperlinks by inserting eyeball icons () around the hyperlink, as shown in the 'projects' link in figure 1 (recall in this example the user expressed an interest in 'intelligent agents'). The advice WebWatcher provides in this fashion is based on knowledge learned from previous tours." 3: "What is the form of the knowledge required by WebWatcher? In general, its task is to suggest an appropriate link given an interest and Web page. In other words, it requires knowledge of the following target function: *LinkQuality : Page x Interest x Link* → [0; 1] The value of LinkQuality is interpreted as the probability that a user will select Link given the current Page and Interest. In the following we present three approaches to learning this target function from experience. The first approach uses previously given tours as a source of information to augment the internal representation of each selected hyperlink. The second approach is based on reinforcement learning. The idea is to find tours through the Web so that the amount of relevant

	information encountered over the trajectory is maximized. The third approach is the combined method that includes both of the first two approaches." *See also*: **CULLISS** at 2:39-51, 9:41-49, 10:1-7, 10:47-52. **MLADENIC** at 2. **REFUAH** at 17:21-43. **SCHUETZE** at 21:57 – 22:16, 22:31-48, 30:58 – 31:13.
Claim 22	
22. The method of claim **1** wherein the monitored user interactions include a sequence of interaction times	*See*: **MLADENIC** at 2. **REFUAH** at 5:57-58. **WASFI** at Abstract, 57..
Claim 32	
32. A program storage device accessible by a central computer, tangibly embodying a program of instructions executable by the central computer to perform method steps for providing automatic, personalized	*See* citations for claim 1[preamble].

information services to a user u, the method steps comprising:	
a) transparently monitoring user interactions with data while the user is engaged in normal use of a client computer in communication with the central computer;	*See* citations for claim 1[a].
b) updating user-specific data files, wherein the user-specific data files comprise the monitored user interactions with the data and a set of documents associated with the user;	*See* citations for claim 1[b].
c) estimating parameters of a learning machine, wherein the parameters define a User Model specific to the user and wherein the parameters are estimated in part from the user-specific data files;	*See* citations for claim 1[c].
d) analyzing a document d to identify properties of the document;	*See* citations for claim 1[d].
e) estimating a probability P(u|d) that an unseen document d is of interest to the user u, wherein the probability P(u|d) is	*See* citations for claim 1[e].

estimated by applying the identified properties of the document to the learning machine having the parameters defined by the User Model; and	
f) using the estimated probability to provide automatic, personalized information services to the user	*See* citations for claim 1[f].
Claim 34	
34. The program storage device of claim **32** wherein analyzing the document d provides for the analysis of documents having multiple distinct media types.	JOACHIMS at 1: "In particular we describe WebWatcher, a system that accompanies the user as he or she browses the Web." *See also*: **CULLISS** at 1:22-28, 2:19-24. **MLADENIC** at 2. **REFUAH** at 1:63 – 2:2. **SCHUETZE** at 4:12-35, 5:48-58. *See generally* "Summary of the Invention", 5:43 – 8:32. **WASFI** at 57, 58. Autonomy Press Release, at 2 Autonomy Agentware User Guide, at AUT00004-5

'276 Patent	
Claim 1	
A computer-implemented method for providing personalized information services to a user, the method comprising:	*See* citations for '040 Patent, claim 1[preamble].
a) transparently monitoring user interactions with data while the user is engaged in normal use of a browser program running on the computer;	*See* citations for '040 Patent, claim 1[a]. *See also*: **MLADENIC** at 2. **REFUAH** at 5:57-58. **WASFI** at Abstract, 57.
b) analyzing the monitored data to determine documents of interest to the user;	*See* citations for '040 Patent, claim 1[b].
c) estimating parameters of a user-specific learning machine based at least in part on the documents of interest to the user;	*See* citations for '040 Patent, claim 1[c].
d) receiving a search query from the user;	JOACHIMS at 2: "In this example, we enter the phrase 'intelligent agents' as our interest. WebWatcher now returns us to the initial page, prepared to guide our tour." 2: "In addition to highlighting hyperlinks WebWatcher also provides other forms of assistance. In

particular, it provides a keyword search using a variant of the Lycos search engine applied to the set of pages previously visited by WebWatcher."

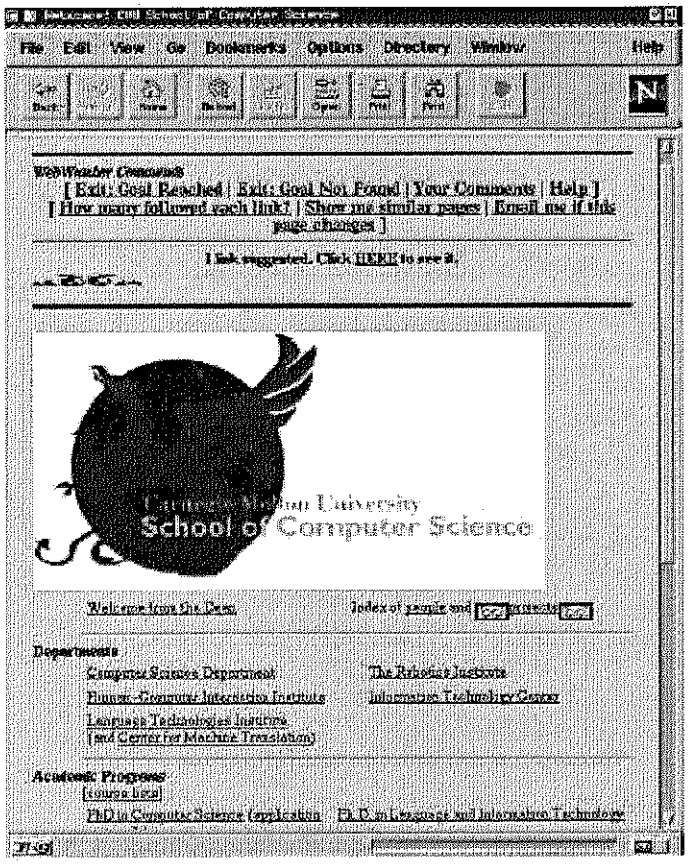

Figure 1: The Front-Door page with WebWatcher's additions.

See also:

	CULLISS at 2:39-51. **MLADENIC** at 1, 2. **REFUAH** at 1:63 – 2:2; 3:12-24. **SCHUETZE** at 21:57 – 22:16; 22:31-48. Autonomy Agentware User Guide, at AUT00002 **MONTEBELLO** at 3.
e) retrieving a plurality of documents based on the search query;	JOACHIMS at 2: "In this example, we enter the phrase 'intelligent agents' as our interest. WebWatcher now returns us to the initial page, prepared to guide our tour." 2: "In addition to highlighting hyperlinks WebWatcher also provides other forms of assistance. In particular, it provides a keyword search using a variant of the Lycos search engine applied to the set of pages previously visited by WebWatcher."

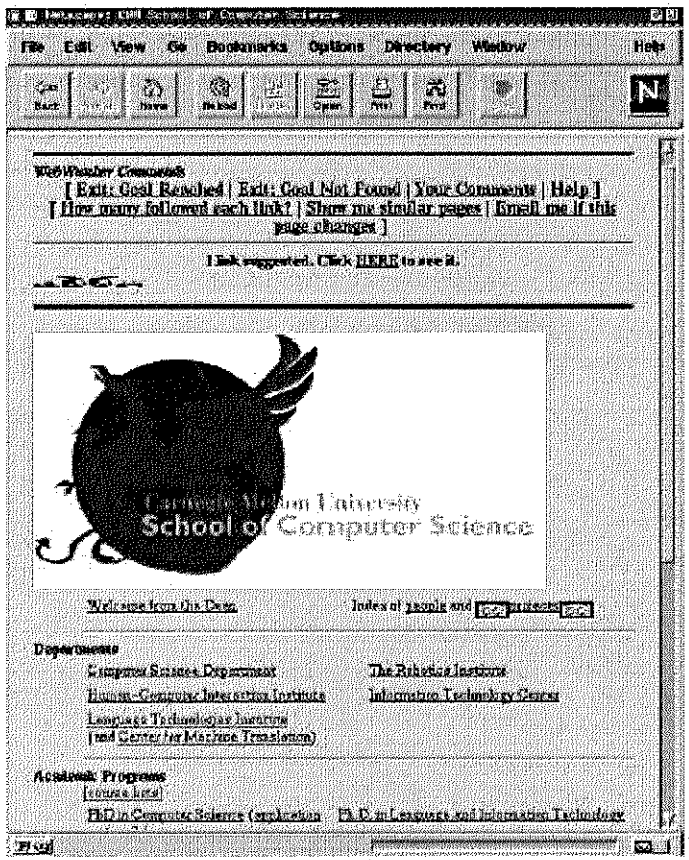

Figure 1: The Front-Door page with WebWatcher's additions.

See also:

CULLISS at 2:39-51.

	MLADENIC at 1. **REFUAH**, 1:63 – 2:2; 3:12-24. **SCHUETZE**, 21:57 – 22:16; 22:31-48. **MONTEBELLO** at 3.
f) for each retrieved document of said plurality of retrieved documents: identifying properties of the retrieved document, and applying the identified properties of the retrieved document to the user-specific learning machine to estimate a probability that the retrieved document is of interest to the user; and	*See* citations for '040 Patent, claim 1[d, e].
g) using the estimated probabilities for the respective plurality of retrieved documents to present at least a portion of the retrieved documents to the user.	*See* citations for '040 Patent, claim 1[f].
Claim 3	
3. The method of claim **1**, wherein transparently monitoring user interactions with data comprises	JOACHIMS at 2: "In this example, we enter the phrase 'intelligent agents' as our interest. WebWatcher now returns us to the initial page, prepared to guide our tour (figure 1). We are no longer browsing alone, but will be accompanied by WebWatcher on any sequence of hyperlinks we follow from this point forward."

monitoring user interactions with data during multiple different modes of user interaction with network data.	p. 2: "In general, the user may click on any hyperlink, recommended or not. Each time the user selects a hyperlink, WebWatcher accompanies the user to the next page, and logs this hyperlink selection as a training example for learning to improve future advice." p. 6: "Richer dialogs with users. One major shortcoming of the current WebWatcher is that it allows the user only to express a few keywords, and only at the beginning of the tour. A more flexible approach would involve an ongoing dialog with the user, much more like that a museum visitor might have with a human guide." *See also*: **CULLISS** at 2:43-46, 3:29-35, 3:46-56. **MLADENIC** at 3. **REFUAH** at 5:34-50. **SCHUETZE** at 18:11-17. **WASFI** at 58, 61. Autonomy Press Release, at 1
Claim 5	
5. The method of claim **1**, further comprising analyzing the monitored data to determine documents not of interest to the user, and wherein estimating parameters of a user-specific	p. 4:

learning machine further comprises estimating parameters of a user-specific learning machine based at least in part on the documents not of interest to the user.

Reinforcement Learning

Reinforcement learning allows agents to learn control strategies that select optimal actions in certain settings. Consider an agent navigating from state to state by performing actions. At each state s the agent receives a certain reward $R(s)$. The goodness of an action a can be expressed in terms of an evaluation function $Q(s, a)$, defined for all possible state-action pairs. The value of $Q(s, a)$ is the discounted sum of future rewards that will be obtained if the agent performs action a in state s and subsequently chooses optimal actions. If the agent can learn this function, then it will know how to act in any state. More precisely,

$$Q(s_t, a) = \sum_{i=0}^{\infty} \gamma^i \cdot R[s_{t+i+1}]$$

where s_t is the state the agent is in at time t, and where γ is a discount factor $0 \leq \gamma < 1$ that determines how severely to discount the value of rewards received further into the future. Under certain conditions, the Q function can be iteratively approximated by updating the estimate for $Q(s, a)$ repeatedly as follows [see [Watkins, 1989]]:

$$Q_{n+1}(s, a) = R(s') + \gamma \max_{a' \in \text{actions}_{s'}} [Q_n(s', a')]$$

s' is the state resulting from performing action a in state s. Once $Q(s, a)$ is known, the optimal control strategy for the agent is to repeatedly pick the action a that maximizes $Q(s, a)$ for its current state s.

Figure 4 gives an example. Boxes represent possible states of the agent. The edges represent actions that bring the agent from one state to another. The edges are annotated with values of the function $Q(s, a)$. In the rightmost state the agent receives a reward of 1. The reward is 0 in all other states. If the agent always follows the action with the highest Q value, it will get to the reward state in the smallest number of steps and thus maximize the discounted reward it receives.

	To the extent this reference does not teach this claim element, this reference in combination with the knowledge of one of ordinary skill in the art renders this claim element obvious. See, e.g.: **CULLISS** at 3:57-65, 4:61-64, 5:4-10. **EDWARDS** at 35-36. **GREEN** at 5, 8. **MLADENIC** at 8. *See generally* 10-11. **PAYNE** at 3. **PAYNE II** at 6. **REFUAH** at 22:6-14. **SCHUETZE** at 17:47-67. **WASFI** at 58. **TAN** at 4.2.1, 4.3.2. **HOFFERER** at 3, 4.2; Figs. 1, 4, 5, and 6. **KAMBA** at 5-7. Autonomy Press Release, at 1 Autonomy Technology Whitepaper, at AUT00070 Autonomy Agentware User Guide, at AUT00123

	LIEBERMAN, Autonomous Interface Agents, at 8 **LIEBERMAN**, Letizia: An Agent That Assists Web Browsing, at 4, 6 **ASNICAR** at 2, 5.
Claim 6	
6. The method of claim **1**, wherein monitoring user interactions with data for a document comprises monitoring at least one type of data selected from the group consisting of information about the document, whether the user viewed the document, information about the user's interaction with the document, context information, the user's degree of interest in the document, time spent by the user viewing the document, whether the user followed at least one link contained in the document, and a number of links in the document followed by the user.	JOACHIMS at 3: "What is the form of the knowledge required by WebWatcher? In general, its task is to suggest an appropriate link given an interest and Web page. In other words, it requires knowledge of the following target function: *LinkQuality : Page x Interest x Link* → [0; 1] The value of LinkQuality is interpreted as the probability that a user will select Link given the current Page and Interest. In the following we present three approaches to learning this target function from experience. The first approach uses previously given tours as a source of information to augment the internal representation of each selected hyperlink. The second approach is based on reinforcement learning. The idea is to find tours through the Web so that the amount of relevant information encountered over the trajectory is maximized. The third approach is the combined method that includes both of the first two approaches." *See generally* 3-4. *See also*: **CULLISS** at 2:43-46, 3:27-35. **MLADENIC** at 8. **REFUAH** at 5:34-50, 14:54-59. **SCHUETZE** at 5:36-40, 11:12-14, 18:11-17. **WASFI** at Abstract, 60, 61.

Claim 22	
22. The method of claim 1, wherein identifying properties of the retrieved document comprises identifying properties selected from the properties consisting of a topic associated with the retrieved document, at least one product feature extracted from the retrieved document, an author of the retrieved document, an age of the retrieved document, a list of documents linked to the retrieved document, a number of users who have accessed the retrieved document, and a number of users who have saved the retrieved document in a favorite document list.	JOACHIMS at 2: "Clicking on the command 'how many followed each link?' asks WebWatcher to display for each hyperlink the number of previous visitors who took that link." *See also*: **CULLISS** at 2:26-37. **MLADENIC** at 3, 4, 12. *See generally* 3-6. **REFUAH** at 7:53 – 8:6, 9:50-59, 20:19-30, 21:6-30. *See generally* 20:19- 21:36. **SCHUETZE** at 6:58 – 7:15, 10:40-56, Fig. 3. *See generally* 17:47 – 18:27. **WASFI** at 61.

Claim 7	
7. The method of claim **1**, wherein said plurality of retrieved documents correspond to a respective plurality of products.	JOACHIMS at 2: "In this example, we enter the phrase 'intelligent agents' as our interest. WebWatcher now returns us to the initial page, prepared to guide our tour (figure 1). We are no longer browsing alone, but will be accompanied by WebWatcher on any sequence of hyperlinks we follow from this point forward." *See also*: **CULLISS** at 9:55- 10:13. *See generally* 9:55 – 11:33. **MLADENIC** at 2, 8, Fig. 2. **REFUAH** at 1:63 – 2:2, 3:56 – 4:4, 7:24-32, 18:35-39, 18:40-55. **SCHUETZE** at 35:66 – 36:8.
Claim 21	
21. The method of claim **1**, wherein using the estimated probabilities for the respective plurality of retrieved documents to present at least a portion of the retrieved documents to the user comprises presenting to the user at least said portion of the retrieved documents based on the estimated probability that the retrieved document is of interest to the user and the relevance of the retrieved document to the search query.	*See* citations for claim 1[g].