Letizia: An Agent That Assists Web Browsing

Henry Lieberman
Media Laboratory
Massachusetts Institute of Technology
Cambridge, MA, USA
lieber@media.mit.edu

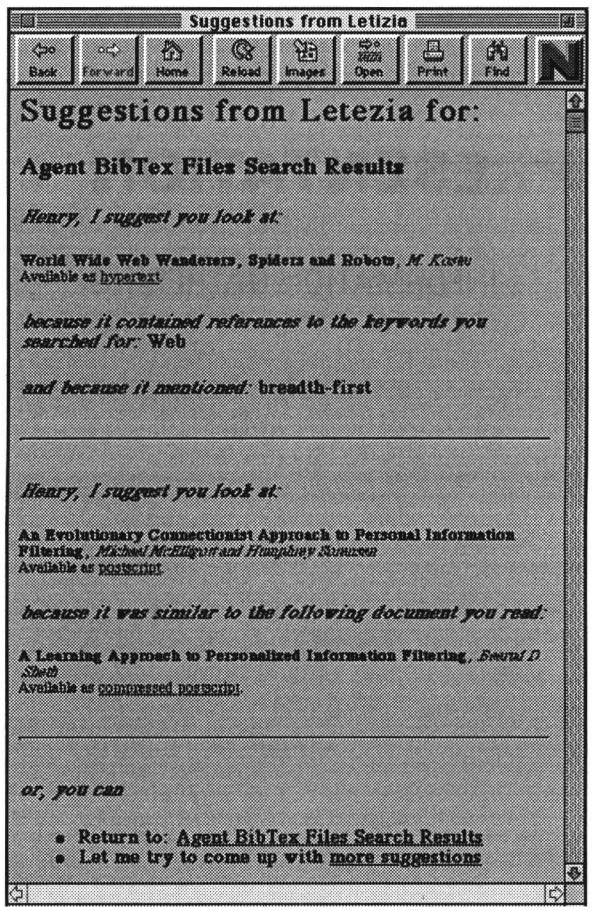

Abstract

Letizia is a user interface agent that assists a user browsing the World Wide Web. As the user operates a conventional Web browser such as Netscape, the agent tracks user behavior and attempts to anticipate items of interest by doing concurrent, autonomous exploration of links from the user's current position. The agent automates a browsing strategy consisting of a best-first search augmented by heuristics inferring user interest from browsing behavior.

1 Introduction

"Letizia Álvarez de Toledo has observed that this vast library is useless: rigorously speaking, a *single volume* would be sufficient, a volume of ordinary format, printed in nine or ten point type, containing an infinite number of infinitely thin leaves."

- Jorge Luis Borges, *The Library of Babel*

The recent explosive growth of the World Wide Web and other on-line information sources has made critical the need for some sort of intelligent assistance to a user who is browsing for interesting information.

Past solutions have included automated searching programs such as WAIS or Web crawlers that respond to explicit user queries. Among the problems of such solutions are that the user must explicitly decide to invoke them, interrupting the normal browsing process, and the user must remain idle waiting for the search results.

This paper introduces an agent, *Letizia*, which operates in tandem with a conventional Web browser such as Mosaic or Netscape. The agent tracks the user's browsing behavior -- following links, initiating searches, requests for help -- and tries to anticipate what items may be of interest to the user. It uses a simple set of heuristics to model what the user's browsing behavior might be. Upon request, it can display a page containing its current recommendations, which the user can choose either to follow or to return to the conventional browsing activity.

2 Interleaving browsing with automated search

The model adopted by Letizia is that the search for information is a cooperative venture between the human user and an intelligent software agent. Letizia and the user both browse the same search space of linked Web documents, looking for "interesting" ones. No goals are predefined in advance. The difference between the user's search and Letizia's is that the user's search has a reliable static evaluation function, but that Letizia can explore search alternatives faster than the user can. Letizia uses the past behavior of the user to anticipate a rough approximation of the user's interests.

Critical to Letizia's design is its control structure, in which the user can manually browse documents and conduct searches, without interruption from Letizia. Letizia's role during user interaction is merely to observe and make inferences from observation of the user's actions that will be relevant to future requests.

In parallel with the user's browsing, Letizia conducts a resource-limited search to anticipate the possible future needs of the user. At any time, the user may request a set of recommendations from Letizia based on the current state of the user's browsing and Letizia's search. Such recommendations are dynamically recomputed when anything changes or at the user's request.

Letizia is in the tradition of *behavior-based* interface agents [Maes 94], [Lashkari, Metral, and Maes 94]. Rather than rely on a preprogrammed knowledge representation structure to make decisions, the knowledge about the domain is incrementally acquired as a result of inferences from the user's concrete actions.

Letizia adopts a strategy that is midway between the conventional perspectives of *information retrieval* and *information filtering* [Sheth and Maes 93]. Information retrieval suggests the image of a user actively querying a base of [mostly irrelevant] knowledge in the hopes of extracting a small amount of relevant material. Information filtering paints the user as the passive target of a stream of [mostly relevant] material, where the task is to remove or de-emphasize less relevant material. Letizia can interleave both retrieval and filtering behavior initiated either by the user or by the agent.

3 Modeling the user's browsing process

The user's browsing process is typically to examine the current HTML document in the Web browser, decide which, if any, links to follow, or to return to a document previously encountered in the history, or to return to a document explicitly recorded in a *hot list*, or to add the current document to the hot list.

The goal of the Letizia agent is to automatically perform some of the exploration that the user would have done while the user is browsing these or other documents, and to evaluate the results from what it can determine to be the user's perspective. Upon request, Letizia provides recommendations for further action on the user's part, usually in the form of following links to other documents.

Letizia's leverage comes from overlapping search and evaluation with the "idle time" during which the user is reading a document. Since the user is almost always a better judge of the relevance of a document than the system, it is usually not worth making the user wait for the result of an automated retrieval if that would interrupt the browsing process. The best use of Letizia's recommendations is when the user is unsure of what to do next. Letizia never takes control of the user interface, but just provides suggestions.

Because Letizia can assume to be operating in a situation where the user has invited its assistance, its simulation of the user's intent need not be extremely accurate for it to be useful. Its guesses only need be better than no guess at all, and so even weak heuristics can be employed.

4 Inferences from the user's browsing behavior

Observation of the user's browsing behavior can tell the system much about the user's interests. Each of these heuristics is weak by itself, but each can contribute to a judgment about the document's interest.

One of the strongest behaviors is for the user to save a reference to a document, explicitly indicating interest. Following a link can indicate one of several things. First, the decision to follow a link can indicate interest in the topic of the link. However, because the user does not know what is referenced by the link at the time the decision to follow it has been made, that indication of interest is tentative, at best. If the user returns immediately without having either saved the target document, or followed further links, an indication of disinterest can be assumed. Letizia saves the user considerable time that would be wasted exploring those "dead-end" links.

Following a link is, however, a good indicator of interest in the document *containing* the link. Pages that contain lots of links that the user finds worth following are interesting. Repeatedly returning to a document also connotes interest, as would spending a lot of time browsing it [relative to its length], if we tracked dwell time.

Since there is a tendency to browse links in a top-to-bottom, left-to-right manner, a link that has been "passed over" can be assumed to be less interesting. A link is passed over if it remains unchosen while the user chooses other links that appear later in the document. Later choice of that link can reverse the indication.

Letizia does not have natural language understanding capability, so its content model of a document is simply as a list of keywords. Partial natural language capabilities that can extract some grammatical and semantic information quickly, even though they do not perform full natural language understanding [Lehnert 93] could greatly improve its accuracy.

Letizia uses an extensible object-oriented architecture to facilitate the incorporation of new heuristics to determine interest in a document, dependent on the user's actions, history, and the current interactive context as well as the content of the document.

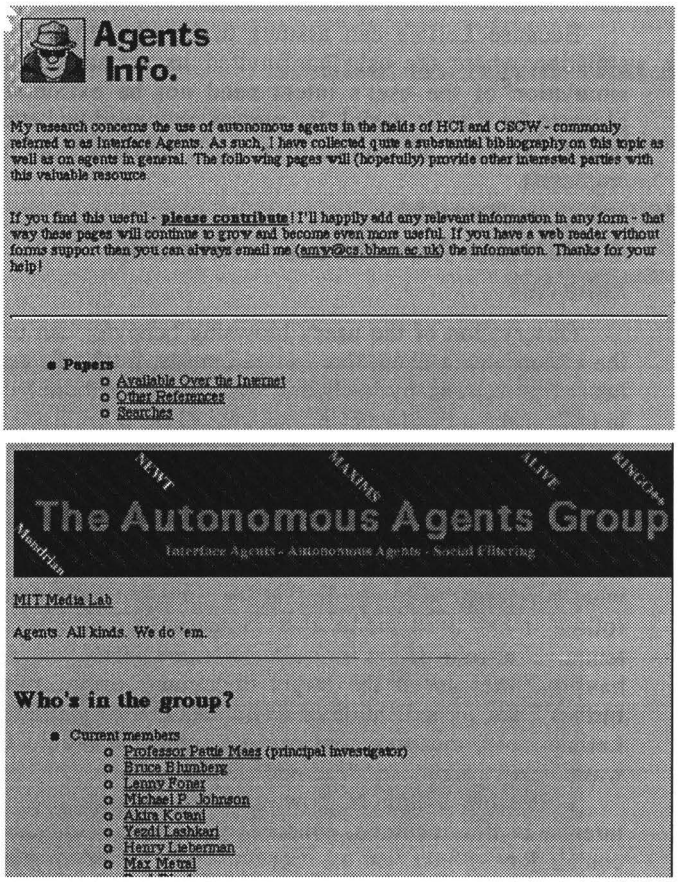

User browses many pages having to do with "Agents". System infers interest in the topic "Agent".

An important aspect of Letizia's judgment of "interest" in a document is that it is not trying to determine some measure of how interesting the document is in the abstract, but instead, a *preference* ordering of interest among a set of links. If almost every link is found to have high interest, then an agent that recommends them all isn't much help, and if very few links are interesting, then the agent's recommendation isn't of much consequence. At each moment, the primary problem the user is faced with in the browser interface is "which link should I choose next?", And so it is Letizia's job to recommend which of the several possibilities available is most likely to satisfy the user. Letizia sets as its goal to recommend a certain *percentage* [settable by the user] of the links currently available.

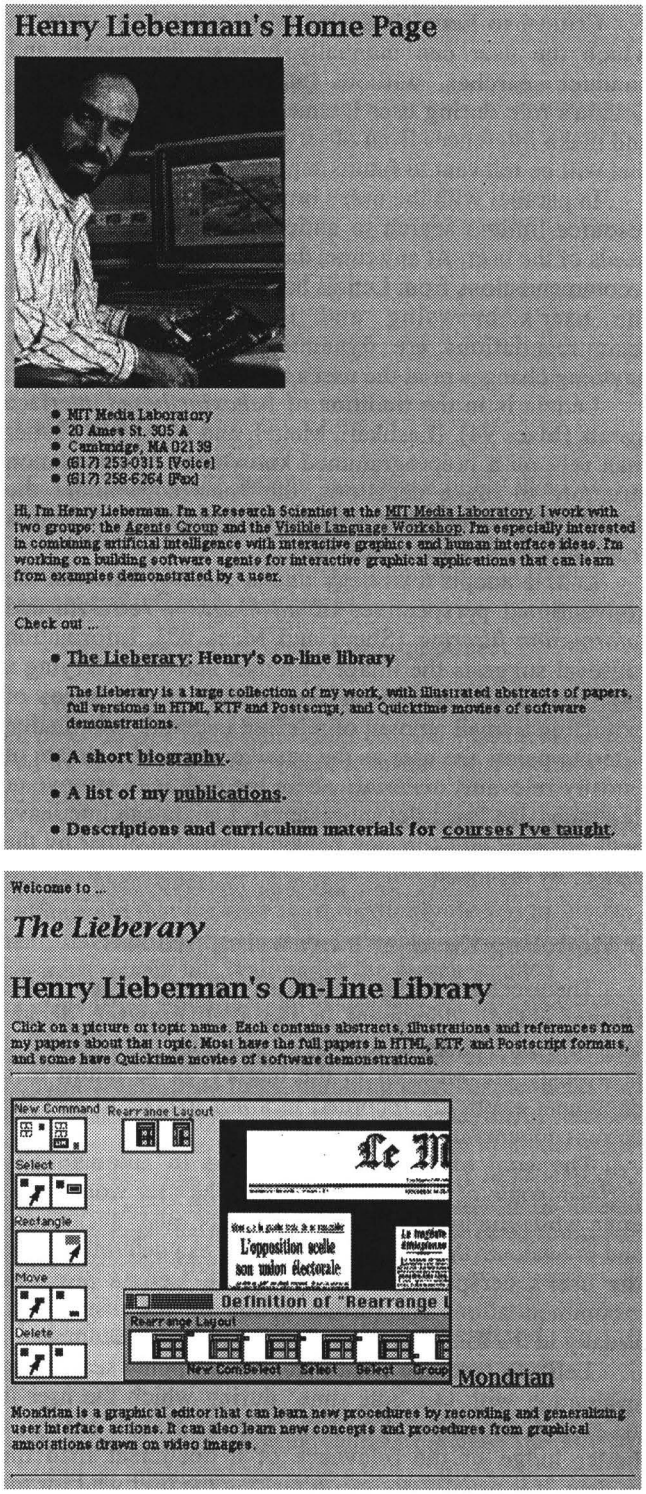

Later, the user independently browses a personal Web page, with a publications list. Letizia recommends articles having to do with "Agents".

5 An example

In the example, the user starts out by browsing home pages for various general topics such as Artificial Intelligence. Our user is particularly interested in topics involving Agents, so he or she zeros in on pages that treat that topic, such as the general Agent Info page, above. Many pages will have the word Agent in the name, the user may search for the word Agent in a search page, etc. and so the system can infer an interest in the topic of Agents from the browsing behavior.

At a later time, the user is browsing personal home pages, perhaps reached through an entirely different route. A personal home page for an author may contain a list of that author's publications. As the user is browsing through some of the publications, Letizia can concurrently be scanning the list of publications to find which ones may have relevance to a topic for which interest was previously inferred, in this case the topic Agents. Those papers in the publication list dealing with agents are suggested by Letizia.

Letizia can also explain why it has chosen that document. In many instances, this represents not the only reason for having chosen it, but it selects one of the stronger reasons to establish plausibility. In this case, it noticed a keyword from a previous exploration, and in the other case, a comparison was made to a document that also appeared in the list returned by the bibliography search.

6 Persistence of interest

One of the most compelling reasons to adopt a Letizia-like agent is the phenomenon of *persistence of interest*. When the user indicates interest by following a link or performing a search on a keyword, their interest in that topic rarely ends with the returning of results for that particular search.

Though the user typically continues to be interested in the topic, he or she often cannot take the time to restate interest at every opportunity, when another link or search opportunity arises with the same or related subject. Thus the agent serves the role of remembering and looking out for interests that were expressed with past actions.

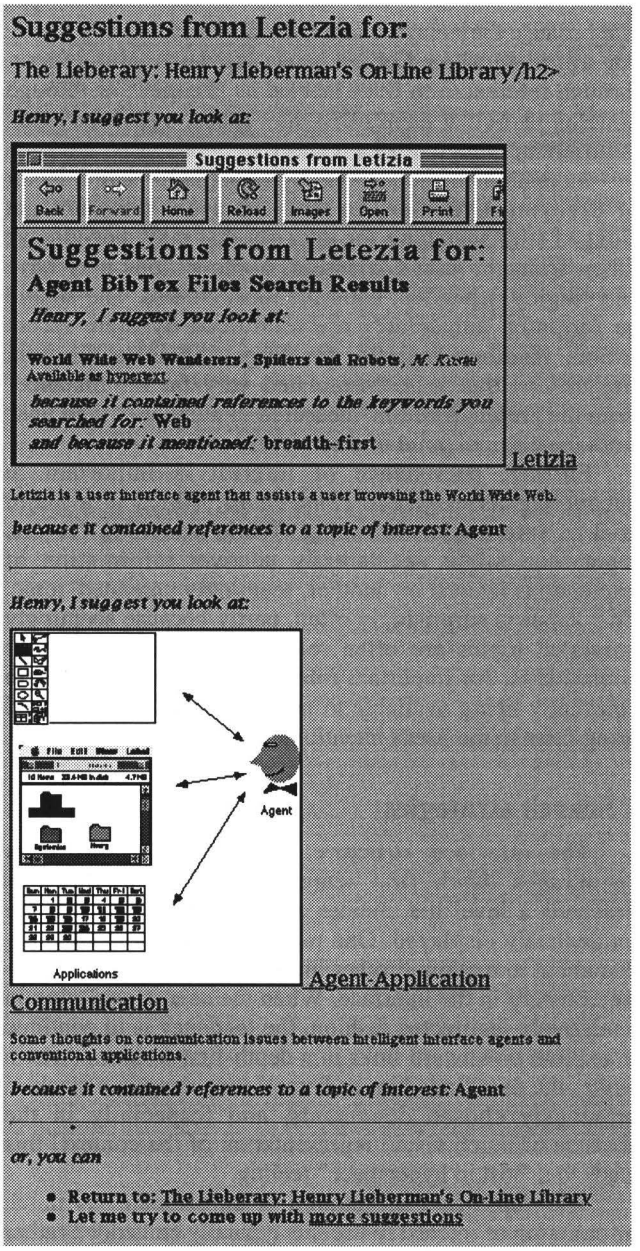

Persistence of interest is also valuable in capturing users' preferred personal strategies for finding information. Many Web nodes have both subject-oriented and person-oriented indices. The Web page for a university or company department typically contains links to the major topics of the department's activity, and also links to the home pages of the department's personnel. A particular piece of work may be linked to by both the subject and the author.

Some users may habitually prefer to trace through personal links rather than subject links, because they may already have friends in the organization or in the field, or

just because they may be more socially oriented in general. An agent such as Letizia picks up such preferences, through references to links labeled as "People", or through noticing particular names that may appear again and again in different, though related, contexts.

Indications of interest probably ought to have a factor of decaying over time so that the agent does not get clogged with searching for interests that may indeed have fallen from the user's attention. Some actions may have been highly dependent upon the local context, and should be forgotten unless they are reinforced by more recent action. Another heuristic for forgetting is to discount suggestions that were formulated very far in "distance" from the present position, measured in number of web links from the original point of discovery.

Further, persistence of interest is important in uncovering *serendipitous connections*, which is a major goal of information browsing. While searching for one topic, one might accidentally uncover information of tremendous interest on another, seemingly unrelated, topic. This happens surprisingly often, partly because seemingly unrelated topics are often related through non-obvious connections. An important role for the agent to play is in constantly being available to notice such connections and bring them to the user's attention.

7 Search strategies

The interface structure of many Web browsers encourages depth first search, since every time one descends a level the choices at the next lower level are immediately displayed. One must return to the containing document to explore brother links at the same level, a two-step process in the interface. When the user is exploring in a relatively undirected fashion, the tendency is to continue to explore downward links in a depth-first fashion. After a while, the user finds him or herself very deep in a stack of previously chosen documents, and [especially in the absence of much visual representation of the context] this leads to a "lost in hyperspace" feeling.

The depth-first orientation is unfortunate, as much information of interest to users is typically embedded rather shallowly in the Web hierarchy. Letizia compensates for this by employing a breadth-first search. It achieves utility in part by reminding users of neighboring links that might escape notice. It makes user exploration more efficient by automatically eliding many of the "dead-end" links that waste users' time.

The depth of Letizia's search is also limited in practice by the effects of user interaction. Web pages tend to be of relatively similar size in terms of amount of text and number of links per page, and users tend to move from one Web node to another at relatively constant intervals. Each user movement immediately refocuses the search, which prevents it from getting too far afield.

The search is still potentially combinatorially explosive, so we put a resource limitation on search activity. This limit is expressed as a maximum number of accesses to non-local Web nodes per minute. After that number is reached, Letizia remains quiescent until the next user-initiated interaction.

Letizia will not initiate further searches when it reaches a page that contains a search form, even though it could benefit enormously by doing so, in part because there is as yet no agreed-upon Web convention for time-bounding the search effort. Letizia will, however, recommend that a user go to a page containing a search form.

In practice, the pacing of user interaction and Letizia's internal processing time tends to keep resource consumption manageable. Like all autonomous Web searching "robots", there exists the potential for overloading the net with robot-generated communication activity. We intend to adhere to conventions for "robot exclusion" and other "robot ethics" principles as they are agreed upon by the network community.

8 Related work

Work on intelligent agents for information browsing is still in its infancy. The closest work to this is [Armstrong, et. al. 95], especially in the interface aspects of annotating documents that are being browsed independently by the user. Letizia differs in that it does not require the user to state a goal at the outset, instead trying to infer "goals" implicitly from the user's browsing behavior. Also quite relevant is [Balabonovic and Shoham 95], which requires the user to explicitly evaluate pages. Again, we try to infer evaluations from user actions. Both explicit statements of goals and explicit evaluations of the results of browsing actions do have the effect of speeding up the learning algorithm and making it more predictable, at the cost of additional user interaction.

[Etzioni and Weld 94], [Knoblock and Arens 93], and [Perkowitz and Etzioni 95] are examples of a knowledge-intensive approach, where the agent is pre-programmed with an extensive model of what resources are available on the network and how to access them. The knowledge-based approach is complementary to the relatively pure behavior-based approach here, and they could be used together.

Automated "Web crawlers" [Koster 94] have neither the knowledge-based approach nor the interactive learning approach,. They use more conventional search and indexing techniques. They tend to assume a more conventional question-and-answer interface mode, where the user delegates a task to the agent, and then waits for the result. They don't have any provision for making use of concurrent browsing activity or learning from the user's browsing behavior.

Laura Robin [Robin 90] explored using an interactive, resource-limited, interest-dependent best-first search in a browser for a linked multimedia environment. Some of the ideas about control structure were also explored in a different context in [Lieberman 89].

9 Implementation

Letizia is implemented in Macintosh Common Lisp. It uses Netscape as a Web browser and user interface. The agent runs as a separate process, and communication between Lisp and Netscape takes place using AppleEvents and AppleScript interprocess communication. Currently, we are severely limited by the extent to which Netscape is programmable via AppleEvents. HTML is parsed using the Zebu parser-generator [Laubsch 94].

Acknowledgments

This work has been supported in part by research grants to the MIT Media Laboratory from Alenia, ARPA/JNIDS, Apple Computer, the National Science Foundation, and other Media Lab sponsors.

References

[Armstrong, Freytag, Joachims and Mitchell 95] Robert Armstrong, Dayne Freitag, Thorsten Joachims and Tom Mitchell, WebWatcher: A Learning Apprentice for the World Wide Web, in *AAAI Spring Symposium on Information Gathering*, Stanford, CA, March 1995.

[Balabonovic and Shoham 95] Marko Balabanovic and Yoav Shoham, Learning Information Retrieval Agents: Experiments with Automated Web Browsing, in *AAAI Spring Symposium on Information Gathering*, Stanford, CA, March 1995.

[Etzioni and Weld 94] Oren Etzioni and Daniel Weld, A Softbot-Based Interface to the Internet, *Communications of the ACM*, July 1994.

[Knoblock and Arens 93] Craig Knoblock and Yigal Arens, An Architecture for Information Retrieval Agents, *AAAI Symposium on Software Agents*, Stanford, CA, March 1993.

[Koster 94] M. Koster, World Wide Web Wanderers, Spiders and Robots, http:// web.nexor.co.uk/ mak/doc/robots/robots.html

[Lashkari, Metral, and Maes 94] Yezdi Lashkari, Max Metral, Pattie Maes, Collaborative Interface Agents, *Conference of the American Association for Artificial Intelligence*, Seattle, August 1994.

[Laubsch 94] Zebu: A Tool for Specifying Reversible LALR(1) Parsers, Hewlett-Packard Laboratories, 1994.

[Lehnert and Sundheim 91] Wendy Lehnert and B. Sundheim, A Performance Evaluation of Text-Analysis Technologies," *AI Magazine*, p. 81-94., 1991

[Lieberman 89] Henry Lieberman, Parallelism in Interpreters for Knowledge Representation Systems, in *Concepts and Characteristics of Knowledge-Based Systems*, M. Tokoro, Y. Anzai, A. Yonezawa, eds., North-Holland, 1989.

[Maes 94] Pattie Maes, Agents that Reduce Work and Information Overload, *Communications of the ACM,* July 1994.

[Perkowitz and Etzioni 95] Mike Perkowitz and Oren Etzioni, Category Translation: Learning to Understand Information on the Internet, *International Joint Conference on Artificial Intelligence*, Montréal, August 1995.

[Robin 90] Robin, Laura, Personalizing Hypermedia: The Role of Adaptive Multimedia Scripts, MS Thesis, Massachusetts Institute of Technology, 1990.

[Sheth and Maes 93] Beerud Sheth and Pattie Maes, Evolving Agents for Personalized Information Filtering, *IEEE Conference on Artificial Intelligence for Applications*, 1993.