Recursive Reviews

Artificial Intelligence, Libraries, and Information Retrieval By Martin Halbert In the science fiction short story "Anniversary" (Amazing, March 1959), Isaac Asimov described a computer system that combined advanced elements of artificial intelligence and information retrieval. Called "Multivac" in the story (I wonder if the name was inspired by the UNIVAC systems that were being marketed in the early fifties), Asimov's system is described as "a mile-long super-computer that was the repository of all the facts known to man; that guided man's economy; directed his scientific research; helped make his political decisions--and had millions of circuits left over to answer individual questions that did not violate the ethics of privacy." Multivac was capable of understanding and answering what we would now call natural language queries on any topic. The protagonists of the story typed in their questions on a terminal that worked much like a typewriter. Their questions required that the system not only collate information, but also draw conclusions. Imagine a system that could answer a query such as one made in the story: "Why is Trans-space Insurance conducting its Silver Queen search-project to which reference was made in the previous question?" Multivac was not only capable of answering this question, but it was also cognizant enough of privacy issues NOT to provide the information to unauthorized researchers. Multivac was the ideal combination of AI and IR technologies that functioned as an oracle of all recorded knowledge. While no system currently in the works has capabilities even remotely like Asimov's Multivac, computer scientists and librarians have at least begun to take first steps toward the concept. If Multivac is an ideal to aim for, what has been accomplished so far? Is AI technology really relevant to libraries in the nineties or is it just a science fiction dream? To help you answer this question for yourself, this column reviews both selected current articles about AI in libraries and a few basic guides to the field of artificial intelligence.

A mass of controversy and confusion surrounds artificial intelligence. To be able to make any kind of reasonable judgments about AI in libraries, one must understand the basic issues and history of the AI field. I have included three basic sources that constitute an excellent foundation in the literature of the field, both at the layman and intermediate levels of technical knowledge. I have also included some additional sources relating AI more specifically to libraries and information work.

Alberico, Ralph, and Mary Micco. Expert Systems for Reference and Information Retrieval. Westport, CT: Meckler, 1990. (ISBN: 0-88736-232-X)

This is the best general text on AI in libraries. An enjoyable mix of theory and practice, the book includes a history of artificial intelligence, many how-to chapters on the various ways of creating library expert systems, a good discussion of trends, and an excellent bibliography on the field. The examples in Prolog are interesting to work through and illustrate the concepts well. I give this source my highest recommendation.

Kurzweil, Raymond, ed. The Age of Intelligent Machines. Cambridge, MA: MIT Press, 1990. (ISBN: 0-262-11121-7)

This is currently the premiere overview text on artificial intelligence. I know of no better contemporary overview of the AI field. It provides both historical and topical overviews of all major realms of artificial intelligence. The book is both understandable to the layman and authoritative (it includes 23 articles by major AI figures such as Marvin Minsky). The many inserts, illustrations, and personal accounts make the book very enjoyable. Kurzweil conservatively speculates on what impacts AI may have in the near-term future on society, artistic expression, medicine, and other areas. The article by Edward Feigenbaum describes a hypothetical network of artificially intelligent electronic libraries.

+ Page 23 +

Graubard, Stephen R., ed. The Artificial Intelligence Debate: False Starts, Real Foundations. Cambridge, MA: MIT Press, 1988.

This collection of articles examines in detail some of the major controversies that have plagued AI research during its brief history. It is very important to understand the context in which many of the exaggerated claims of AI research have been made. Without this context, one may be tempted to unfairly write off the field as being composed of ivory tower dreamers or worse. AI researchers in many cases had justifiable reasons to expect truly dramatic strides in a short period of time and, being human, were also prone to blind allegiance to their own theories. Without the background of the symbol-processing versus "perceptron" approaches, for example, it is impossible to understand why researchers have such strong feelings for or against current research into neural networks. Some of the articles are fairly technical, but most are easily comprehensible, even to the non- specialist. The discussion provides the reader with balanced views on different sides of the current controversies.
--- "AI: Metamorphosis or Death?" (State of the Art section that includes 8 articles). BYTE 16 (January 1991): 236-301. (ISSN 0360-5280)
--- This special section of articles is geared to the typical PC enthusiast, and it should be intelligible to anyone who uses computers regularly (which certainly includes all PACS Review readers). These articles provide a valuable computer industry perspective on the changing expectations of people involved in artificial intelligence work. The old AI view was that researchers were on the verge of developing machines that would truly be as intelligent as human beings. This view has died. It has been supplanted by the more realistic view that computers can

accomplish amazing feats in very well-defined, limited technical domains that require sophisticated human training and effort. The nagging, unsatisfying thing about the realistic view, however, is that once an AI application is successfully implemented, one tends to discount it as involving any intelligence at all.

+ Page 24 +

For example, since 1979 Digital Equipment Corporation has successfully used an expert system called XCON for custom configuration of its VAX and PDP-11 computers. When human technicians did the configuration, analyzing and fitting a specified set of components into a cramped computer chassis, this seemed like a technically sophisticated process. Now, when the same job is accomplished by a computer system applying a set of rules to the problem, it seems like "just rote work." This points out the "moving target" nature of artificial intelligence; what was science fiction in the fifties (e.g., computers playing chess at the grand master level and analyzing blood diseases) now seem like routine programming. AI's biggest challenge may be to settle on a satisfactory focus (i.e., a defining vision of what AI is trying to accomplish). This cannot be the naive goal of early AI--to make computers intelligent--since it is now quite clear that we have no definition of exactly what constitutes intelligence! Given the identity crisis of AI in the nineties, how can it be relevant to libraries and librarians?
-- Bailey, Charles W., Jr. "Intelligent Library Systems: Artificial Intelligence Technology and Library Automation Systems." Advances in Library Automation and Networking 4 (1991): 1-23. (ISBN: 1-55938-188-4) -- In this paper, Charles Bailey provides a clear, systematic analysis of the opportunities and difficulties associated with applying AI to library work. He discusses barriers to the development of library AI systems, such as the high cost of AI development, the limitations of current technology, and the lack of AI expertise among the majority of librarians. He also identifies fruitful avenues for future progress. Any promising strategy requires finding tasks suitable for AI, rather than trying to force AI into roles that are best managed by human beings. This still leaves many intriguing possibilities for improving library systems through AI techniques. Bailey's examples focus on making databases more intelligent and capable of improved responsiveness to the user, whether the database is a library catalog or an intelligent computer-assisted instruction program. The bibliography is an excellent tool for exploring the literature further.

+ Page 25 +

Morris, A. "Expert Systems for Library and Information Services- -A Review." Information Processing & Management 27, no. 6 (1991): 713-724. (ISSN 0306-4573)

For an overview of specific library AI applications through the years, turn to Morris' article. Morris analyzes library AI applications in five areas: online information retrieval, cataloging, abstracting, reference work, and indexing/classification. The portrait he paints across the library landscape is a somewhat disappointing one of several decades of prototypes that rarely resulted in successful systems. The fundamental problem that researchers repeatedly encountered was the fact library work is fundamentally "messy." It requires common-sense background knowledge about the world and flexibility in applying rules and drawing conclusions--exactly where current AI systems are weak. The

marginal successes of the field occur primarily in online information retrieval, where the data to be operated on have already been extensively refined, indexed, and otherwise regularized by human beings. This points out that library AI researchers have been taking on unrealistically complicated problems given the current state of the art. This is in no way an indictment of library AI researchers. On the contrary, the first three sources reviewed indicate that this has been a common mistake in all areas of AI research. If an AI system is to successfully capture the "intelligence" that goes into solving a problem, the scope of a problem must be clearly understood, and its solution must be able to be formulated in rules. Morris concludes by saying that the work accomplished so far must be looked on as experimental, laying the foundation for future implementations that will be more productive. The experiments of the past should bear fruit in the future, enabling expert systems to handle many of the mundane features of library work.

+ Page 26 +

Bailey, Charles W., Jr., and Judy E. Myers, comps. Expert Systems in ARL Libraries. SPEC Kit 174 (Washington, DC: Association of Research Libraries, 1991). Also, ERIC, ED 337 178.

This ARL report offers a concise survey of current AI activity in libraries. The general findings are roughly as follows: very few ARL libraries (seven percent, or six libraries) have developed or are planning to develop expert systems, but most ARL libraries (72 percent) believe that expert systems will become an important technology in libraries during the nineties. Taken together, these statements are intriguing. Evidently, ARL libraries believe that turnkey systems will become widely available in this decade without much development effort from libraries. While this is possible, it does not seem likely. The development of turnkey integrated library automation systems was a gradual process that built up synergistically from circulation and cataloging experiments in many individual libraries to commercial endeavors. While this pattern may yet repeat itself in library AI systems, it is still too early to tell.

Sparck Jones, Karen. "The Role of Artificial Intelligence in Information Retrieval." Journal of the American Society for Information Science 42, no. 8 (1991): 558-565. (ISSN 0002-8231)

This article provides some very specific arguments concerning ways in which AI should NOT be used in information retrieval. Karen Sparck Jones emphasizes that the potential for AI should not be overestimated. Her points about the limitations of AI systems are similar to those made in the other works reviewed here. Sparck Jones' article is worth reading because of her detailed analysis of AI's limitations in information retrieval applications. She notes the more limited, focused applications that could benefit from artificial intelligence techniques. Sparck Jones also points out fundamental limitations in IR applications that AI is unlikely to ever overcome.

+ Page 27 +

Waterman, Donald A. A Guide to Expert Systems. Reading, MA: Addison-Wesley, 1986. (ISBN 0-201-08313-2)

I conclude with an excellent basic textbook on expert systems, since such systems seem to hold the most immediate promise and have been the AI technology most widely and successfully implemented both in libraries and in other industries. Although several years old now, Waterman's book is still one of the best introductions to the topic that I am aware of. He gives both the theory and practice of expert systems, with much discussion of actual working products. If you are seriously considering developing an expert system or even just trying to estimate what it would take to implement one, Waterman's book is a worthwhile resource. The effort to make computers more intelligent--however one interprets "intelligence"--will certainly continue. The approaches taken and the perspectives brought to this effort are changing, becoming more mature and realistic in their goals. One can only hope that expert systems WILL become an important library technology in the nineties, since this would mean that many of the problems with AI have been solved. The verdict will not be in for some time though, so stay tuned.

+ Page 28 +

About the Author Martin Halbert, Automation and Reference Librarian, Fondren Library, Rice University, Houston TX 77251-1892. Internet: HALBERT@RICEVM1.RICE.EDU.