



## Letters

The following letter was addressed to Daniel Bobrow, editor of *Artificial Intelligence*. Many of us felt that the issue raised is a very important one for the AAI and deserved wide exposure. It is printed here, along with Bobrow's reply, for your interest.

- Ed

Editor:

When Henry Oldenburg published the first issue of *Philosophical Transactions* in March 1665, he had done more than merely patch together a new magazine. He had invented the scientific journal. The intervening three centuries have proven Oldenburg's invention to be a priceless vehicle for the dissemination of knowledge. It is, therefore, ironic indeed that artificial intelligence, a field whose very essence is knowledge, has developed a literature that is extraordinarily difficult and inefficient to use.

Effective use of the literature of AI is frustrated by two fundamental deficiencies: (a) there is no central index to the field's published works, and (b) not only are far too many original works not published in journals, but a shockingly high percentage of these are "published" in sources that may be generously described as inaccessible.

As an example of a field that does not have these problems, consider medicine. *Index Medicus* is an index to the world's entire published medical literature, compiled since 1880 and available on-line since 1966. The *Science Citation Index* appeared in 1955 and is based on the clever idea of tabulating, for each article published, all subsequent papers that cite that article. It is the completeness of these tools that make them so valuable; thus, the anthologies of seminal papers or selected bibliographies that are common in AI are, at best, a poor substitute for them. The index published by Scientific DataLink (*The Scientific DataLink Index to Artificial Intelligence Research: 1954-1984*, New York, 1985) and the *CompuMath Citation Index* (*CMCI CompuMath*

*Citation Index*, Philadelphia: Institute for Scientific Information, annual) are a start, but both fall short in providing the complete coverage of the field needed. That a field of study as large as AI lacks a central index to its literature seems positively medieval.

The second deficiency of the AI literature is not only more subtle, but also requires some proof of its existence. An admittedly less-than-rigorous proof is found in table 1. It was constructed by examining the references cited in the main articles of the most recent issues of the leading journal in AI and in medicine. Each reference was classified as citing either a journal article, book selection, conference proceeding, technical report, or other source.

The table clearly shows that a large proportion of work cited by AI researchers is not to be found in journals. This has two important consequences. The first is that perhaps as much as 40% of the work cited by AI researchers (comprising the book, technical report, and unknown categories) has not been impartially peer-reviewed. The reasons for, and ramifications of this staggering observation are far beyond the scope of this letter, and so will not be discussed further. The second consequence of AI's low use of journals is that of impeded physical access to information. Not only do journals not leave the library, they also enter the library in a much more predictable way than other media. Books, conference proceedings, or technical reports relevant to one's research may or may not be ordered by one's library, and the presence or absence of each must be ascertained individually, perhaps involving trips

to widely separated campus libraries. By contrast, the presence or absence, and even physical location, of a journal need be established but once.

The fundamental differences in the nature of research results from medicine and AI will, of course, lead to unavoidable disparities in the optimal means by which they are communicated. Despite this, there is much that AI can do to more effectively reap the advantages of journals, for example:

Accelerate the trend to publish papers from workshops and single-theme conferences as a supplement to, or special issue of, an appropriate journal;

For larger conferences (for example, AAI & IJCAI), publish abstracts in a single [special] journal issue and the papers themselves in an appropriate journal soon afterwards. Pre-prints or reprints could be distributed at the conference itself;

Technical reports should be purely archival, for example restricted to system manuals. Novel aspects of a system would be journal reported.

Technical reports present a special problem, for they are the most inaccessible of references. This is clearly evident in a recent advertisement that stated "It took nearly 3 years, more than 600 hours, more than 120 transatlantic calls to 6 countries, and almost 1100 letters" to collect 14 years worth of AI reports from one university. (Direct mailing from Scientific DataLink, Summer 1987, promoting the collected AI reports of the University of Edinburgh.) If nothing else, this may explain why no central AI index exists: it is simply too difficult to construct one.

Although some may defend technical reports because they enable the rapid publication of results, this merely illustrates that there is an unfilled niche in the current population of AI journals. Since it is common in engineering fields to report work-in-progress and partial results in technical reports, there is no reason not to place this material in journals, thereby giving it wider distribution. An interesting possibility is that an analysis of what portions of technical reports are most frequently cited might lead to the development of a new form of published work

<i>New England Journal of Medicine</i> <sup>1</sup>	<i>Artificial Intelligence</i> <sup>2</sup>
Journal	41%
Book	22%
Conference	-
Proceedings	17%
Technical Reports <sup>3</sup>	15%
Other <sup>4</sup>	2%
Unknown <sup>5</sup>	2%

1 Six articles (totalling 212 references) were analyzed *New England J Med* 316 (1987) 1613-1642  
2 Seven articles (totalling 211 references) were analyzed. *Artificial Intelligence* 31 (1987) 271-395, and *Artificial Intelligence* 32 (1987) 1-146. Does not add to 100% because of round-off error  
3 Includes departmental reports, technical memoranda, and manuals.  
4 Includes dissertations and personal communications  
5 Includes those references for which inspection alone was insufficient to determine the type of reference

**Types of References Cited in Two Journals**

amenable, of course, to journal format.

To summarize, my suggestions are: (1) An AI organization, either extant or specially created, should take responsibility for the compilation and dissemination of a monthly index (collated and bound annually) to the entire literature of the field of artificial intelligence. I make no suggestions about how to finance this enterprise, and only note that computer science is not a financially impoverished discipline.

(2) Strenuous efforts should be made to publish more original AI research work in journals. This will no doubt require: (a) continued proliferation of journals; (b) frequent publication of supplements to, or special issues of, journals in conjunction with conferences, workshops, etc (c) de-emphasis of technical reports and, to a somewhat lesser degree, books as a vehicle for the publication of original research.

No matter what the field, progress is built upon a foundation of previous achievements. In artificial intelligence there is currently no simple and efficient means to retrieve these past lessons. Although it will clearly be nontrivial to implement the suggestions made in this letter, I humbly propose that the field of artificial intelligence devote some of its formidable energy to the problem of search—as applied to itself.

John G. Sotos  
Stanford University

Dr Sotos:

Your letter points up serious information access problems for the field of artificial intelligence. It is easy to understand this problem from a historical perspective. As a very young field, it was only 18 years ago that the first international journal, *Artificial Intelligence*, was started. The small number of people working in the field allowed informal communications to be the norm. Rapid computer-based access to current papers by most researchers, and rapid distribution to the invisible college on the mailing list sufficed to keep the few centers of excellence and many outposts reasonably informed. The few hundreds gathered at the biannual international conference could hear the latest results, and have a chance to meet the most relevant of their colleagues face to face. But the recent explosive growth in the field has been matched by a corresponding growth in the number of conferences each year. Many new journals have started, particularly in the last two years, and may provide some answer to your third shortfall.

A number of organizations have recognized the information access problem. At the New Mexico State University, Yorick Wilks has organized *Artificial Intelligence Abstracts*, which is still in its infancy. It publishes a bound index that covers the major conference proceedings and University Report Series as well as the major journals in artificial intelligence. It claims a wide coverage of papers in

the field. Its first issue appeared in 1987, and abstracts start with papers published in 1985. The service collects abstracts, and also accepts abstracts submitted electronically. The editor of the abstracts journal, Yorick Wilks, has a CS-Net address (yorick@nmsu.cs.net), which makes communication easier, particularly for persons in the United States, and somewhat less for members of the overseas community.

The Turing Institute, a nonprofit foundation set up in Glasgow, Scotland (36 North Hanover Street, Glasgow G1 2AD), is an older service that also publishes a periodical, *The Turing Institute Abstracts in Artificial Intelligence*, with its first volume starting in 1986. The Turing Institute has an estimated 30,000 abstracts on-line and this database is growing at an estimated 1000 entries per month. The Turing Institute provides some interesting additional services to institutional subscribers. Subscribers may connect via phone to a Turing Institute computer to conduct an on-line search of the abstract database. Full text match of free vocabulary is used. Facilities for interactive modification of the retrieval formula are supported. More important from some points of view, the library at the Turing Institute provides, on demand for its subscribers, copies of the full article for any abstract retrieved.

These two services are a partial answer for your first two shortfalls—indexed search, and document access—for all referenced material. Although each service provides only current coverage, and historical coverage would be useful, the future will keep getting better. But there are problems with both of these services. For communication reasons at this time, the Turing Institute service is only available reasonably in Europe. The New Mexico venture is only available in hardcopy form.

The AAI is attempting to make an arrangement to allow our membership some computer access to databases like these. The executive Council of the AAI approved an initial investment of funds to make something happen this year, and I have been delegated to explore the possibilities. In addition to possibly getting access to

the existing databases, the AAI wants to explore creative ways to have the community augment such databases on line. We also hope to find ways AI researchers can build more intelligent access and retrieval mechanisms to turn such databases into real "knowledge bases". People who are interested in this project are encouraged to contact me, electronically or by mail.

Daniel G. Bobrow  
AAAI President-Elect  
Palo Alto, California

#### Letovsky, Part Two

Editor.

You appear to have stirred up something of a hornet's nest with the publication of the Letovsky article!

Personally, I found the article to be stimulating reading and highly amusing to boot, and found the comments published in the Letters column somewhat reminiscent of Sam, The American Eagle on the Muppet Show, full of bombastic indignation.

Anyone who thinks there is no room for humor in the world of artificial intelligence should have heard Patrick Winston's keynote speech at AAI-87!

Keep up the good work, and by all means, let us have an outrageous article from time to time!

David J. Steele  
Singapore

Editor:

I just read the comments on Stan Letovsky's report entitled, "Ecclesiastes: A Report From the Battlefields of the Mind-Body Problem." I am astounded that anyone would think that this report would hurt *AI Magazine* in any way. The review is everything an *AI Magazine* article should be (and often isn't)—well written, informative, and entertaining. It is neither "sloppy" nor sophomoric, and is certainly not "unprofessional, offensive tripe." I feel that the report ranks with the best articles that have appeared in the *AI Magazine* since its inception.

Of course, the report is not a "neu-

tral" account of the conference, but I do not believe that conference reports should necessarily be neutral. Colorful reports are often much more informative than the dry conference and workshop reports that generally show up in AI publications. Instead of

"[establishing] editorial policy to reject this kind of material in the future" (as Leonard R. Kasday suggests), I suggest that the *AI Magazine* actively solicit such reports for conferences and workshops—especially conferences and workshops on controversial top-



## ASSISTANT PROJECT MANAGER

**Pioneering work in artificial intelligence technology for demonstration projects. AI systems for Space Station, Lunar Base and/or other Humans-in-Space systems automation. Development and integration methodologies for AI knowledge acquisition, representation, and implementation into these systems.**

**NASA-Ames is a multidisciplinary research center chartered to provide National leadership in aeronautics, astronautics, and space research. Join an elite team at the cutting edge of technology, daring what others dream.**

**Requires advanced degree with extensive experience in project management and artificial intelligence. U.S. citizenship required for federal civil service position. NASA is an equal opportunity employer. Send resume to:**

**E. Harris  
Box RIS, Mail Stop 241-6  
NASA-Ames Research Center  
Moffett Field, CA 94035**

**For free information circle no. 154**

# Prentice Hall's Expert Systems!

## Artificial Intelligence Through PROLOG

Neil C Rowe  
1988, 368 pp, cloth 013-048679-5, \$37 33

## Common LISP: A Tutorial

Wendy L. Milner  
1988, 384 pp, paper 013-152844-0, \$26 67

## The Scheme Programming Language

R Kent Dybvig  
1987, 304 pp, paper 013-791864-X, \$22 00

Available at better bookstores or direct from  
Prentice Hall at (201) 767-5937. For  
quantity orders call (201) 592-2498

Order your copy of the **Prentice Hall Professional/Technical/Reference Catalog: Books for Computer Scientists, Computer/Electrical Engineers and Electronic Technicians** for only \$2 00 and receive \$5 00 off your first purchase from the catalog (013-622804-6). Prices subject to change without notice.

## The XLISP Primer

Bonnie J Fladung  
1987, 144 pp, paper 013-972084-7, \$25 67

## The T Programming Language

Stephen Slade  
1987, 448 pp, paper 013-881905-X, \$22 00

## A Guide to Commercial Artificial Intelligence: Fundamentals and Real-World Applications

Wendy B Rauch-Hindin  
1988, 592 pp, cloth 013-368770-8, \$28 95



**PRENTICE HALL**

Simon & Schuster  
Higher Education Group  
Englewood Cliffs, NJ 07632

For free information, circle no. 26

ics. Such a policy would let everyone interested in AI see the healthy (and heated) discussion that is present in certain subfields of AI, but which so often is restricted to the participants in small workshops.

I also take issue with Bob Engelmores weak support of the article. *The AI Magazine* is not a scholarly journal, and should be the precise place where controversial and even outrageous articles appear. Such articles make a welcome change from the usual fare that appears in it.

Similarly, calling the report "biased" does it a great disservice. Although biased can simply mean "inclined in some direction," the meanings "unduly or unfairly influenced" or "prejudiced" (all taken from the *Oxford English Dictionary*) are much more common, and the use of "biased" in Bob Engelmores response clearly allows these stronger meanings. If any apologies are called for, it is this use

of "biased" that needs apology.

Peter F. Patel-Schneider  
Palo Alto, California

Editor:

I offer this note as a philosophical footnote to the Letovsky article and subsequent reactions. Perhaps you can attribute my slowness of response to my being a philosopher.

While I both enjoyed and had problems with Letovsky's "report" and subsequent exchanges, I would like to briefly address the apparent assumption that dualism is a religious position. I would assert that mind-body dualism is purely a philosophical theory, and not a religious one.

There is, of course, a spirit-body dualism that is a religious position and this has been confused historically with mind-body dualism. But where these have been clearly understood, a thinker who is both a spirit-body dualist and a mind-body dualist will actu-

ally be a body-soul-spirit tripartist. The question here is whether soul (mind) and spirit are reducible.

Mind-body dualism is a different issue. Secular workers in AI should recognize it as such and not dismiss it out of hand since they are not religious.

There are in fact many kinds of dualism. The most easily caricatured is the ghost-in-the-machine variety commonly attributed to Descartes and lambasted by Ryle. There are also linguistically based dualisms: We simply say different things of the brain than we do of the mind. (The brain weighs approximately five pounds, but how much does the mind weigh?) Between these, are such positions as phenomenalism, and the form of dualism Eccles has been espousing, whatever that is.

These forms of dualism are generally supported not on religious grounds, but as attempts to explain such secu-

# Sierra™ OPS5

Create sophisticated expert systems on your PC with the expert systems language.

Sierra OPS5 is a fast, sophisticated, absolutely 100% complete implementation of OPS5 designed specifically for the PC

**FULL:** We left nothing out — 32 bit integers, real numbers, files, 'build', external functions, complete!

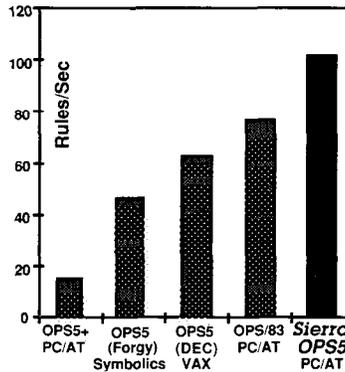
**FAST:** Running the same program used to benchmark other expert system languages, Sierra OPS5 topped them all. No other data-driven production system language, interpreted or compiled, matched our performance!

**FLEXIBLE:** DEVELOP your expert systems using the fully integrated workbench environment — multiple windows, multi-buffer editor, incremental compilation, full tracing and debugging. COMPILE your knowledge base to create a stand-alone executable OPS5 program. EMBED multiple independent knowledge bases in your own 'C' programs and call them when and as often as YOU want!

**RESOURCE EFFICIENT:** The runtime libraries require as little as 40K!

Requires an IBM PC/XT/AT or compatible with 384K RAM, DOS 2.0 or later, Microsoft C V4.0 V5.0 or QuickC required for external functions or stand-alone program. Trademarks: OPS5+/Computer Thought Corp. OPS/83/Production Systems Tech. IBM/International Business Machines. Microsoft/Microsoft Corp. Symbolics/Symbolics Inc. DEC/Digital Equip. Corp.

NASA  
"Monkeys & Bananas"  
Benchmark



**Inference  
Engine  
Technologies**

## PERSONAL VERSION \$129.00

- 250 rules
- OPS5 Workbench
- Compiler
- Runtime Libraries
- External Functions
- Full Manual & OPS5 textbook
- Non-Commercial License

## COMMERCIAL VERSION \$495.00

- Personal Version plus:
- 1000 - 1500 rules
- Full Embedability within 'C' programs
- Commercial Royalty-free runtime license
- 1 year of updates
- Telephone support

Shipping & Handling \$5.00, \$20.00 outside US and Canada. Visa/MC/POs/Drafts on U.S. banks. MA residents add 5% sales tax.

**1-800-255-0625**

in MA 923-0998  
1430 Mass Ave., Suite 306-I  
Cambridge, MA 02138

For free information circle no. 60

lar things as purposive action, consciousness, and intentionality. Some forms of dualism are more alive than others, and that's because the phenomena they are meant to explain are more current. Consciousness, for example, is a notion constantly just below the surface, and sometimes on the surface, in specifically AI literature.

Writing in AI has always had a philosophical flavor, and it seems to be increasing. Clark Glymour commented that people who 10 or 20 years ago would have gone into philosophy (logic), are now going into computer science (AI). I would appeal, at least to these cross-over scholars, not to ignore the history of the issues they address. While you are, possibly of necessity, reductionists, don't simply dismiss long-standing philosophical positions, but help solve the problems that led to these positions

Bruce Boyer  
University of California

### Letovsky Replies

Editor:

My article on the "AI and the Human Mind" conference (*AI Magazine*, Fall 1987, p. 63) elicited a spate of critical letters (see Winter 1987 issue) accusing me of a range of offenses, including unprofessionalism, shoddy journalism, and poor note taking. The letter writers seem to have been outraged that I did not write the standard format, "objective," journalistic conference report. To this charge I plead nolo contendere: the article in question was never intended as a journalistic piece. I wrote it as a personal diary of an unusual experience; it circulated privately among friends and acquaintances for a year and received enthusiastic responses. I sent it to *AI Magazine* with a note explaining this, saying the article was not originally intended for publication, but that its content might be of interest to the *AI*

*Magazine* readership and that I could edit it into publishable form. In his reply, Bob Engelmores not only expressed interest in publishing it, but wrote that, "the article as it stands has a certain spontaneity that I like, and I wouldn't want you to grind and polish that away." On this basis, the article was printed in its original diaristic form with only minor typographical corrections. Granting my own bias in the matter, I think the decision to publish the article in this form was entirely appropriate. *AI Magazine* is not an exclusively technical journal; it is a wide-spectrum magazine that publishes articles of general interest to the AAAI community. This broad mandate, it seems to me, leaves room for an occasional piece that treats the subjective and personal aspects of the scientific experience.

Stanley Letovsky  
Yale University