

## LETTERS

## The Next Knowledge Medium

## Editor:

In his article "The Next Knowledge Medium" (*AI Magazine*, 7 (1):34-46), Mark Stefik contrasts the goal of building a "knowledge medium" with AI's long-standing goal of building an autonomous, intelligent agent. It does indeed seem plausible, as Stefik suggests, that this knowledge medium will be the same sort of indispensable infrastructure of the new information age as the roads and railroads were of nineteenth century France. But I don't think a successful knowledge medium can exist without building autonomous agents first.

Knowledge comes in all sorts of forms and formats (just as railroads used to come in all sorts of gauges). Unlike railroads, I don't think we will be successful in standardizing knowledge (nor should we even try!). We can live with an infrastructure consisting of several knowledge gauges. A key part of that infrastructure will be *knowledge switchers*. In fact, we have a knowledge infrastructure already, and it is already immense. It consists of countless gigabytes of data, books, maps, reports, *etc.* The only problem (the problem that slows today's "isochronic waves") is that the knowledge switchers in this infrastructure are mainly humans—and we are pretty slow switchers! Imagine fully autonomous knowledge switchers—ones that can read books, scan data, answer messages, process memos, conduct surveys, and perform all the other tasks that human knowledge workers do today. It will be necessary for these knowledge switchers to be more than simply passive translators; they must be active agents with local goals and stores of knowledge—just as humans are. These autonomous knowledge processing agents will permit an expansion of Stefik's vision to one in which the "trains" travel throughout the network, automatically changing "gauge" as needed.

I think millions of autonomous, intelligent agents will be the all-important nodes of tomorrow's knowledge medium. In a previous article ("Artificial Intelligence Prepares for 2001," *AI Magazine* 4(4)), I suggested the goal of developing autonomous "computer individuals." Stefik's excellent article makes me even more convinced of the importance of this goal.

Nils Nilsson  
Dept. of Computer Science  
Stanford University  
Stanford, California 94305

## Stefik replies:

Do we need to have autonomous intelligent agents in order to have a successful knowledge medium (Stefik, M. "The next knowledge medium." *AI Magazine* 7(1): 34-46)? Nilsson argues that we must. I think that this is

an important question of emphasis, and that the answer turns on how we envision the medium.

Nilsson's knowledge switchers provide an intriguing vision of action in the medium. Many of his roles for automated knowledge workers—conducting surveys, processing memos, and answering messages—did not occur to me and do seem to require intelligent autonomous agents. Perhaps a hint of these can be seen in the Information Lens project (Malone, T. W., Grant, K. R., & Turbak, F. A. "The information lens: an intelligent system for information sharing in organizations." Tech. Report, Sloan School of Management, MIT) at MIT, where simpler agents are used to help process electronic mail.

But to what extent should fundamental questions about artificial intelligence be considered on the critical path of experimenting with knowledge media? Nobody should ascribe general intelligence to the agents in the information lens; they are small collections of rules for sorting, deleting, and forwarding mail. As computer programs go they are really quite simple. The interesting point is that they also appear to be quite useful. In contrast, if we require very general notions of agent autonomy, or that agents understand sophisticated motivations, or that agents carry out sophisticated planning—then we put the enterprise of building a knowledge medium off into the indefinite future.

So do we have to build autonomous agents first? Here Nilsson and I recommend different strategies. I recommended emphasizing the construction of experimental media, starting with the development of next generation shells for building expert systems. These shells would be used in small technical communities organized around common purposes. Specialized tools might be developed for collaboration, such as for supporting argumentation processes that range from semi-formal to formal representations. Experiments in the design of knowledge markets would be interesting when the communities become large enough. In such settings, many services would present themselves as opportunities for incremental automation. We would also find many opportunities for studying issues of knowledge additivity, which I consider critical.

Nilsson's point about the problems of standardizing knowledge are quite right. Indeed, that was the main point of the section titled "From Standardization to Interoperability." But I think that work on the issues raised in that section will be preceded by simpler attempts at finding ways to evolve standards.

In conclusion, building knowledge media stands as a long term goal. It stands on other work in computer science, and rests on the same core parts of AI of knowledge representation and problem solving. It doesn't replace work on building intelligent autonomous agents, but augments it. It rests on the same observation that has

served the most successful work on expert systems: that (today) knowledge comes (mostly) from people.

Mark Stefik  
Xerox PARC  
3333 Coyote Hill Road  
Palo Alto, California 94304

### Workshop on the Foundations of AI: An On-The-Spot Report

#### Editor:

The NSF and AAAI sponsored Workshop on the Foundations of AI (6-8 February 1986, Las Cruces, New Mexico) is over and, from my perspective at least, it was a very worthwhile event. I am preparing a report that I will send to you in due course. Meanwhile, the following is a somewhat terse and dubious summary that you may wish to publish in the trivia section (or wherever) of *The AI Magazine*. In addition, I noticed that John McCarthy was snapping freely with his camera at the workshop. He may have some amusing illustrations of the unlikely events recorded.†

Friday, February 7th - 9:35 am.

†Thanks to John McCarthy for the photograph reproduced herein  
—Ed

Yesterday we made some progress. Let me elaborate slightly. In his introductory overview Chandra[other name too long]sekarana noted that in AI (and thus I presume in RI also) power is more important than correctness. (The rest of the analysis was shouted in order to compensate for the possibility of minor errors of interpretation that may have crept in.)

Roger Needham, and others, showed that AI doesn't really exist: As soon as you successfully code up some aspect, it ceases to be AI. Mark Halpern exposed the real reason for Turing's self-destruction: acute embarrassment resulting from a realization that he had failed to appreciate a fundamental axiom of western intellectual thought, to wit, "machines can't think."

Jerry Fodor suggested that the mind/brain muddle might not be an object worthy of scientific study (like Tuesdays in whereveritwas Connecticut, even though the garbage man comes that day, *etc.*). Despite the early successes, brains and Tuesdays are just not the sorts of things that you can have a science about. He then proved that network representations (such as that of the brain) cannot possibly exhibit intelligence—tapes, as in Turing Machines, appear to be critical. Thus showing why Lashley failed to find the engram, *etc.*—he was looking in the wrong place! Fodor's contentions were hotly disputed by both the connectionists present, and his Auntie who had apparently spoken to him about it beforehand. The impli-

## A Fuzzy Logic Production System Language and Shell

**F  
L  
O  
P  
S**

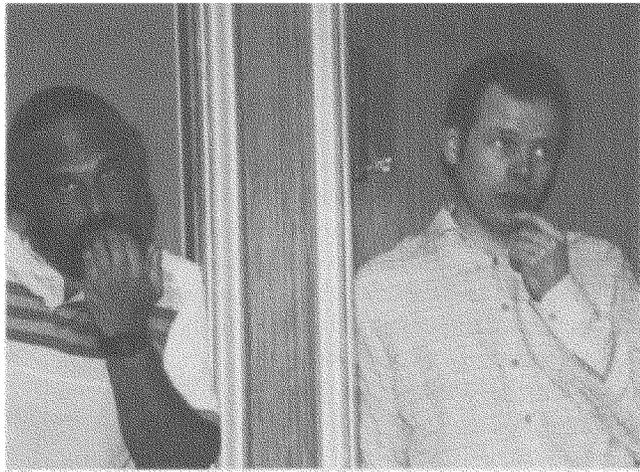
- *Successfully applied* to unsupervised of echocardiogram image classification (Computers in Cardiology 1985, IEEE Computer Society).
- *Seven Data Types*: string, integer, float, fuzzy number, fuzzy set, certainty factor and time tag. Employs Bellman-Zadeh fuzzy logic.
- *Complete set* of fuzzy and non-fuzzy numerical comparison operators.
- *Nested external FLOPS program files.*
- *Read/Write data files* for blackboard, DBMS or external program link.
- *Fully automatic backtracking* in sequential rule-firing version.
- *Parallel rule-firing version* emulates non-Von-Neumann parallel machine.
- *Contradictory conclusions or facts* easily stored in fuzzy sets.
- *External programs called* with command strings or argument list.
- *Rules can create other rules* to any depth.
- *300-page tutorial and reference manual*; over twenty sample programs.
- *PC/XT/AT and compatibles*, with and without 80x87 coprocessor.
- *Free upgrade to FLOPS 1.3* with purchase of FLOPS 1.2 for **\$195**.
- *VAX version with blackboard and IBM mainframe version* in works.

### \* \* For Both Deductive and Inductive Reasoning \* \*

- While FLOPS can be operated in sequential (deductive) mode, the combination of parallel rule firing, external data files as input to rule-generating rules and storage of contradictory information in fuzzy sets makes FLOPS *a powerful tool for inductive reasoning*.
- C program link permits *combining symbolic and numeric computation*.

• Call or write: Dr. William Siler (205) 226-6697  
Kemp-Carraway Heart Institute,  
Birmingham, AL 35234





Schank: "Are those the foundations of AI?"  
 Bundy: "When's lunch?"

cation is that the hitherto unrecognized, seminal paper is "Artificial Digestion" (AISBQ, no. 46, p. 30), even though it fails to draw attention to the (as we now know) crucial role of tape worms.

Bert Dreyfus completed a wonderful day of progress by revealing that he is really one of us. Everyone bad-mouthed expert systems, and that behavior was generally agreed to be a good thing.

In sum: AI doesn't exist; the brain is not the seat of intelligence; and Dreyfus has seen the light, at last.

*Saturday February 8th—10:10 am.*

Yesterday again progress continued apace, undermining the foundations of AI. On Thursday we saw working programs removed from the domain of AI. Yesterday we made great strides towards a further paring down of AI—more encumbering baggage was thrown away in the hopes of revealing a lean and vibrant discipline within this flabby body of folk-knowledge, folk-theories, etc. (or simply, folk-all, as Churchland said, if I heard him correctly)

Alan Bundy gave theories to Cognitive Science, and truth to anyone who wanted it. John McCarthy was, of course, right there to make a grab for it in the guise of the magic of nonmonotonous logic. Not being a logician myself, I can clearly see the merits of making logic less boring, but I'm dubious about the introduction of magic, even as a truth-preserving wheeze.

Somewhere about here the tide turned, and we began to hear about features that do contribute to a discipline of AI. Magic, I've already mentioned; and from John

# Intelligent Tutoring Comes of Age

LISP-ITS: the first commercially available intelligent tutoring system

## Teaching Effectiveness

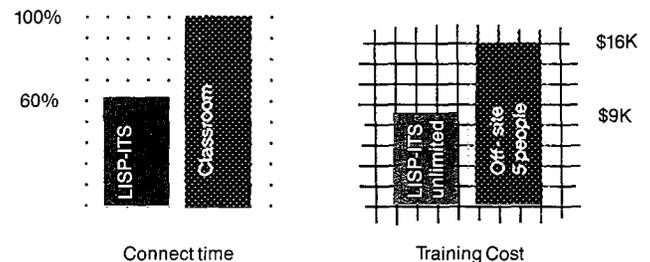
- tutors a full semester LISP course
- proven effective in industrial and university settings
- proven more efficient and effective than standard LISP instruction

## Cost Effectiveness

- less expensive than other training options
- perpetual license means unlimited use: meet your training needs at a moment's notice

## Uniqueness

- combines AI technology and cognitive research
- LISP-ITS is a powerful tutor because it *knows* LISP. It is not an "electronic page-turner"



**Advanced Computer Tutoring, Inc.**

- We also offer dial-in access and LISP-ITS based courses
- The course material that accompanies LISP-ITS will soon be available as a LISP textbook from Addison-Wesley
- Send for more information and stop by our booth at AAAI '86

**Advanced Computer Tutoring, Inc.**  
 701 Amberson Avenue, Pittsburgh, Pa 15232 412-621-5111



a personal, independent consultant  
for AI consumers and product developers

Avron Barr can help you...

- Select projects, people, vendors, tools and services
- Keep your knowledge system development efforts on course
- Inform your management and technical staff through custom briefings and seminars

**Aldo Ventures, 525 University Avenue, Suite 1206,  
Palo Alto, CA 94301 (415) 322-2233**

*Aldo Ventures is the AI consulting business of  
Avron Barr, co-editor of the **Handbook of Artificial  
Intelligence** and a co-founder of **Teknowledge** in 1981.*

Campbell we got the two R's of AI—Rational Reconstructions (roughly, failed attempts to reproduce famous AI programs). We also heard about a methodological constituent: the evaluation of programs *a posteriori*. In keeping with this trend to more bottom-up characterization of the field (and, I suspect, a desire to preserve some vestiges of civilized behavior), Karen Sparck-Jones advocated that the ultimate question of "AI, or not AI?" must be decided democratically, *i.e.*, by a show of hands.

At this point it became clear that we had heard, from a number of speakers, of an implementation-independent (indeed, notation-independent) goal of AI research: the accumulation of brownie points. Let me illustrate this:

    . a brownie point (to a first approximation)  
 ..  
 ..... an accumulation of brownie points  
 ...

This was the major insight of the workshop at this point, as best as I could tell. And, of course, Zenon Pylyshyn latched onto this realization in his talk, the last one of the day. It is true that he never actually mentioned these intriguing objects—but we could all read between the words.

In an effort to put theory into practice, I polled my family at breakfast this morning as to whether what I had drawn adequately represented brownie points. The

resultant three blank faces immediately told me that a lot of computation was being invested in this question (there were no answers just popping out; partial differential equations were perhaps being solved in those heads). In fact the answers are still not in [10:12 am] and with reaction times currently running at about 1hr 53min, I predict that there is a lot of structure in brownie points (they are only weakly equivalent to dots and periods). And when my three reaction times finally come in, I expect (after analysis of the data) to be well on the way to discovering some of the functional primitives of brownie points. Clearly, they are not cognitively impenetrable.

*Sunday February 9th—all day.*  
[No visible signs of life].

*Monday February 10th—10:00 am.*

On Saturday, after a game of "find the meeting room," Paul Churchland, flying in the face of reason, sided with Fodor's Auntie and insisted that study of the brain might contribute to our understanding of RI (and by implication, AI). He argued that this neuronal clump is a sort of surrealistic lunch box containing phase-space and state-space sandwiches. The inevitable Star Wars question was raised by Aravind Joshi; the fundamental problem seems to be in the use of bad language: nonstarred and ill-starred sentences.

Dave Rumelhart reiterated the view that AI is no place for theories. He then weighed into the networks-versus-symbol-processing battle with the suggestions that the microarchitecture of brownie points is likely to be connectionistic, and that this microarchitecture cannot be dismissed as mere implementation detail.

Aaron Sloman, thinly disguised as Marvin Minsky, argued against a dichotomization of systems into those with minds and those that don't mind. He illustrated his argument with a continuum from viruses through ants, gerbils, and chimps to Yoricks, but he could not be lured into a clear statement of the direction in which mindedness was increasing in the continuum presented.

Maggie Boden wound up the day's proceedings with an argument that the benefit of AI in psychology was to be found in the power of computers to draw researchers into such messy and complex theories that they never get out and bother anyone else again. Fortunately, prior perfect analysis *a la* Marr is unlikely when a dynamic researcher, itching to build a respectable heap of code, is given ready access to a powerful, interactive, program development environment—it's our only hope for a science of Tuesdays in whereveritwas Connecticut. At the banquet in the high-tack setting of a local eating house, Roger Schank reminded us about the significance of reminding and the nonsignificance of mathematics (which, incidently, he doesn't like). Frank Harary proposed a vote of thanks in which he handed out hand-crafted brownie points and thanked us for having the good sense to opt for a graph-

theoretic representation of these pointy objects (*i.e.*, the trivial graph).

—from the diary of A.I. Person

Derek Partridge  
Department of Computer Science  
New Mexico State University  
Las Cruces, New Mexico 88003

### The Profit Motive

#### Editor:

Robert Rich brings up an important issue in his letter in the spring issue of the *AI Magazine*. Rich states that "...those who promote the products of research are rewarded by promotion and salary increases, while those who promote commercial products are rewarded . . . by promotions and salary increases."

Exactly. It is high time we removed the profit motive from science. Too long have we suffered from authors who write papers, motivated only by financial and professional advancement.

I would like to propose a solution. *AI Magazine* should refuse to include the names of authors of articles. Instead, *AI Magazine* should give each author a "post office box," and interested readers can contact the authors indirectly through the offices of the *AI Magazine*.

If this idea catches on, there would be other benefits. There would be no more *ad hominem* attacks (who can you attack?). And if we extend the idea to conferences and workshops, fewer people would be willing to give long boring talks (to preserve anonymity, the speaker would be required to wear a mask—preferably uncomfortable).

Of course, the problem with identifying the product (either commercial or research) still exists. The names of projects would have to be deleted from articles, replaced by some artificial, GENSYM-ed name. We might even require that articles be (intentionally) misleading, to prevent any identification of the authors and projects by the content of the article. . . .

Signed,

G00001\*

*In fact, why not extend the principle to our editorial staff? It would then be useless to send letters to "the offices of the AI Magazine" (a rather abstract entity, I might add).—Ed.*

\*AKA Stan Lanning, Xerox PARC, 3333 Coyote Hill Road, Palo Alto, California 94304

## TOPSI — OPS5 FOR THE IBM PC

### THREE POWERFUL VERSIONS:

- 1.2 – Prototyping System. Full OPS5 compatibility at a modest price \$125.
- 2.1 – Production System. Full-sized rules and floating point procedures \$250.
- 3.0 – Professional System. Smart editor, windows, menus. A complete environment \$375.

### FULL SOFTWARE SUPPORT:

- Telephone Hot-Line
- Training and Seminars
- System Design Consulting
- Custom Software Modules

Discounts Available.  
Credit Cards and PO's accepted.



**Dynamic Master Systems, Inc.**

PO Box 566456  
Atlanta, Ga 30356  
(404) 565-0771

## Make Your Mac Run Like a LISP Machine with MacScheme™

■ **HIGH QUALITY BYTE CODE INTERPRETER**  
FAST as interpreted LISP on LISP machines  
COMPACT CODE for large programs  
RELIABLE error reporting and recovery  
Debugger, trace facility  
Now with QUICKDRAW GRAPHICS

■ **EASY TO USE**  
Edit and evaluate expressions in any window  
Editor matches parentheses, indents automatically  
Needs only one drive, 512K or more RAM  
Works with HFS and Macintosh Plus™  
Not copy protected

■ **CONFORMS TO 1985 SCHEME STANDARD** -  
Lexically scoped with block structure  
True first class procedures  
Generic arithmetic: fixnums, flonums, bignums  
Strings, vectors, arrays, ports, etc.



**Semantic Microsystems**



**\$125**

4470 S.W. Hall St.  
Suite 340  
Beaverton, OR 97005  
(503) 643-4539