

Editorial Standards

Editor:

I should like to lodge a complaint about your editorial standards in the article "An Assessment of Tools for Building Large KB Systems," by William Mettrey, in the winter 1987 (volume 9 number 2) *AI Magazine*.

As a primary architect of CRL-Ops and a former KnowledgeCraft class instructor, I had to deal with the general public's misconceptions about forward versus backward chaining systems. Mr. Mettrey's article, in my opinion, is the type which generates the confusion that forward chaining rule systems cannot "backwards chain." This nonsensical view was held by the vast majority of our customers in the KC class.

The section on Rule-Based inference implies that backward chaining is done only by Prolog in KC with its statement "by contrast, KnowledgeCraft implements backward chaining by supporting a version of Prolog." Any forward chaining rules system can efficiently implement constrained backward chaining by simply using a goal structure to search for the required knowledge.

Additionally, to the best of my knowledge, which I will gladly admit is at least a year out of date, Art does not really "implement[s] backward chaining by allowing a single backward-chaining goal pattern as one of the conditions in a forward-chaining rule." More realistically, Art allows default backward chaining to occur when no forward chaining matches are available by automatically creating a goal to search for a match which would instantiate some partially matched rule.

The section on debugging claimed that "Art and Kee provide graphic tracing of rule activity during inferencing" by failed to mention CRL-Ops's facility for this.

Finally, the statements in the

"Requirements for Delivery Environments" regarding parallelism in production systems are particularly misleading. The results quoted are early work of Gupta and Forgy. Gupta's thesis (CMU CS Technical Report 86-122), which, I'm sure, must significantly predate the submission of this article, clearly shows that it is not true that "Production systems in particular appear to lend themselves to parallel execution." A large amount of work is still required to demonstrate if productions will lend themselves to parallel implementations.

Only one highly optimized general purpose parallel production system compiler exists ("Results of Parallel Implementation of Ops5 on the Encore Multiprocessor," Gupta, Forgy, Kalp, Newell, and Tambe Technical report CMU-CS-87-146) and this achieved only reasonable speed ups. The amount of parallelism in large production system programs so far demonstrated by this implementation is limited to an optimistic 13 in a large but highly match intensive program.

What is the review procedure? Was this article passed around in Inference, IntelliCorp, and Carnegie Group? Was it reviewed by experts in the knowledge engineering / production systems research areas?

Brian G. Milnes
Computer Science Department
Carnegie Mellon University
Pittsburgh, Pennsylvania

All articles are reviewed by at least one knowledgeable reviewer. The reviewer of Mettrey's article was not

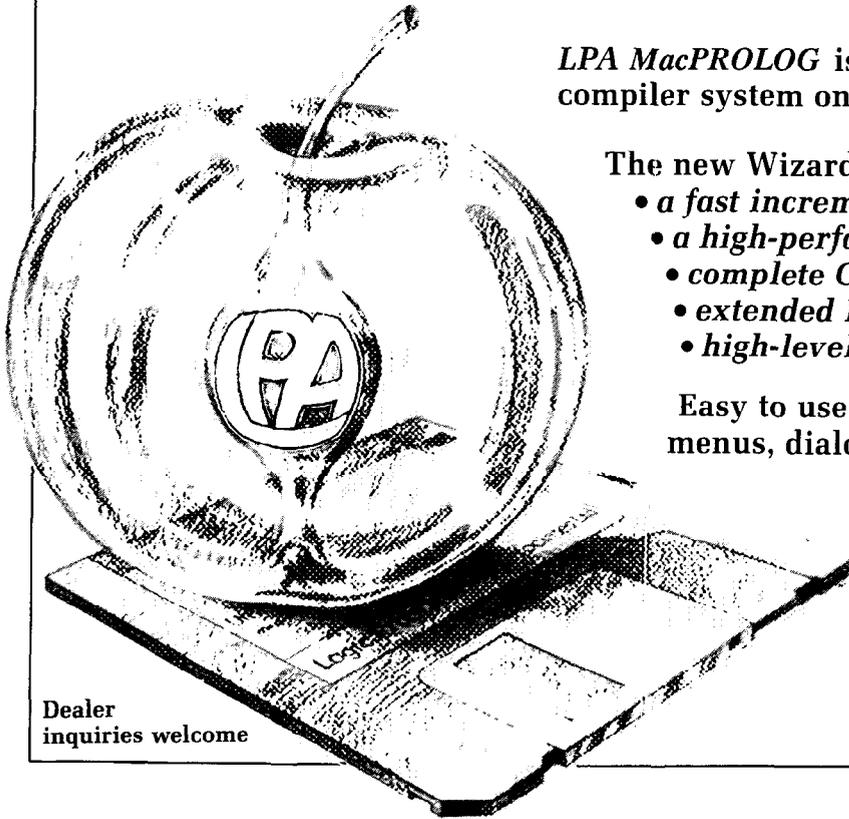
an employee or associate of any of the vendors discussed in the article.
—Ed.

Mettrey Responds

Milnes' reactions to the statement "By contrast, KnowledgeCraft implements backward chaining by supporting a version of Prolog" were that it "generates the confusion that forward chaining rule systems can not backward chain" and "implies that backward chaining is done only by Prolog in KC." Milnes reads more into the statement than was intended. The intent was to state that KnowledgeCraft directly supports backward chaining by supplying a version of Prolog. Milnes' point that a forward chaining tool can effect backward chaining is both correct and well documented (L. Brownston, R. Farrel, E. Kant, and N. Martin, *Programming Expert Systems in OPS5*, Reading, Mass.: Addison Wesley, 1985). This does not, however, eliminate the need for direct support of backward chaining. In my opinion, support of this nature is needed for some problems if only for perspicuity. Finally, I refer Milnes to the Carnegie Group literature (*KnowledgeCraft Overview, Version 3.1*. Carnegie Group.) that states "Using CRL-PROLOG and CRL-OPS together, blackboard control architectures can be designed that integrate goal-driven and data-driven problem solving strategies."

In addition, Milnes is incorrect in his statement that "ART allows default backward chaining to occur when no forward chaining matches are available by automatically creating a goal to search for a match which would instantiate some partially matched rule." ART does not automatically backward chain when it encounters any condition that is unknown. ART generates a goal only when there is a backward-chaining rule that can assist in locating matches for a particular pattern (ART Reference Manual, Version 2.0. Inference Corporation). Expressing a backward chaining rule in forward format is a matter of adding exactly one goal pattern among the conditions on the left-hand side of a rule.

Mac made intelligent



Dealer inquiries welcome

LPA MacPROLOG is the most advanced PROLOG compiler system on the Mac.

The new Wizard Edition includes:

- a fast incremental compiler
- a high-performance optimizing compiler
- complete C/Pascal interfaces
- extended Edinburgh syntax
- high-level graphics system

Easy to use access to the ToolBox for menus, dialogs and windows.

 **MacPROLOG**

For your *free* demo disc for the Mac
or MS/DOS phone:
LPA PLS
+441 871 2016 (800) AI LOGIC

For free information circle no. 74

For example:
(defrule RECOGNIZE-FAULT-IN-
CONNECTION
 (goal (FAULT (CONNECTION
?PIN-A ?PIN-B)))
 (VOLTAGE ?PIN-A ?VOLTAGE-A)
 (VOLTAGE ?PIN-B ~?VOLTAGE-A)
=>
(assert (FAULT (CONNECTION ?PIN-
A ?PIN-B))))

Milnes charges that I failed "to mention CRL-OPS's facility for graphic tracing of rule activity." Although I can understand Milnes' particular interest in CRL-OPS, the examples used in the article were not intended to be exhaustive for every feature of every tool. The goal of the article as stated was "not to compare the relative strengths and weaknesses of tools that were evaluated but rather to assess the current state of commercial knowledge engineering tools and estimate features that will be required in

the next generation." In particular, the article was not intended to serve as a marketing aid for any vendor or tool.

As for Milnes's statement, "parallelism in production systems are particularly misleading," I consider a potential "13 fold" improvement significant.

Finally, Milnes asks "What is the review procedure..." As in any review procedure, the author is unaware of the identity of the reviewers. I can only state that the comments that I received as a result of the reviews were, in my opinion, knowledgeable and served to improve the quality of the article. In addition, I would have objected strenuously if the article had been "passed around" to the various vendors. The article was written from the viewpoint of an (hopefully) unbiased user who had no desire to be subjected to vendor bias or marketing goals.

William Mettrey
Knowledge Systems Corporation
2000 Regency Parkway, Suite 212
Cary, North Carolina 27511

Use of Determinatives in Natural Language Processing

Editor:

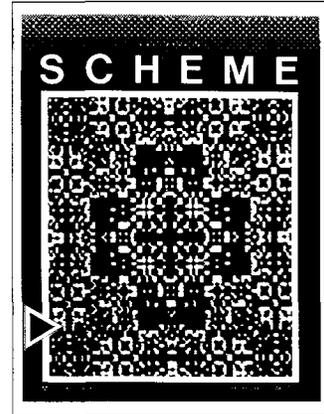
It is not surprising that many AI problems of today find parallels in the concerns of the ancients, but it is remarkable that some old solutions still hold important lessons for us. Consider the problems of language ambiguity and definition of context, both of which also exercised the ancients. Current research in natural language processing is, at one level, a research for methods to determine the ever-changing contexts and relationships in text. The problem is so hard because both context and relationships might change at the level of the

ANNOUNCING

Programming in SCHEME

Programming in SCHEME, written by Michael Eisenberg and edited by Harold Abelson (both at MIT Laboratory for Computer Science), 320 pages, paperbound, 1988, \$32.50. This book provides an introduction to the SCHEME programming language, using Texas Instruments' PC SCHEME as a source of examples. It assumes no previous programming experience on the part of the reader and covers all the basics of the language, some advanced topics, and a number of runnable sample projects.

CONTENTS: Programming in Scheme, Sample Project 1, Expressions & Procedures, Making Choices, Recursion, Debugging, Sample Project 2, Pairs/Lists/Symbols, Subprocedures, Sample Project 3, Environments, Procedures, Altering Bindings/Altering Objects, Debugging 2, Sample Project 4, Answers to Selected Exercises



Commercial Edition of PC SCHEME Software, by Texas Instruments, \$95.00.

Student Edition of PC SCHEME Software, by Texas Instruments, for academic classroom use, contains the full functionality of the Commercial Edition but without expanded or extended memory and external language interface support, \$37.50.

 **The Scientific Press** 507 Seaport Court • Redwood City, CA 94063
In California (415) 366-2577 • Outside California (800) 451-5409

For free information circle no. 26

sentence itself, as by metaphors, or the relationships may not be completely defined, and left to be inferred from the context.

One observes a progressive increase of inherent ambiguity in going from texts for children to those for adults. On the other hand, legal texts, while less ambiguous, are extremely complex. The writings of the ancient Egyptians have interesting parallels—not only across the same groups at a given age but also across ages. Consider the most ancient Egyptian writing. It was phonetic, though only consonantal. How the consonants were to be vocalized was to be inferred from the context. This inference was easy to make. At the same time they introduced logograms, symbols that directly represented objects, such as a rectangle with an opening for a house, or the picture of a lute or goose and so on. For words that could

not be drawn, they substituted words with similar sounds that could easily be drawn. Thus to write *nôfer*, good, they used *nefer*, the lute; *sa*, the son, was replaced by *sa*, the goose; and for *per*, to go out, *per*, the house was substituted.

Many short words, that occurred infrequently by themselves, lost their meaning and became syllabic signs that could be employed in any word with the syllables as constituent.

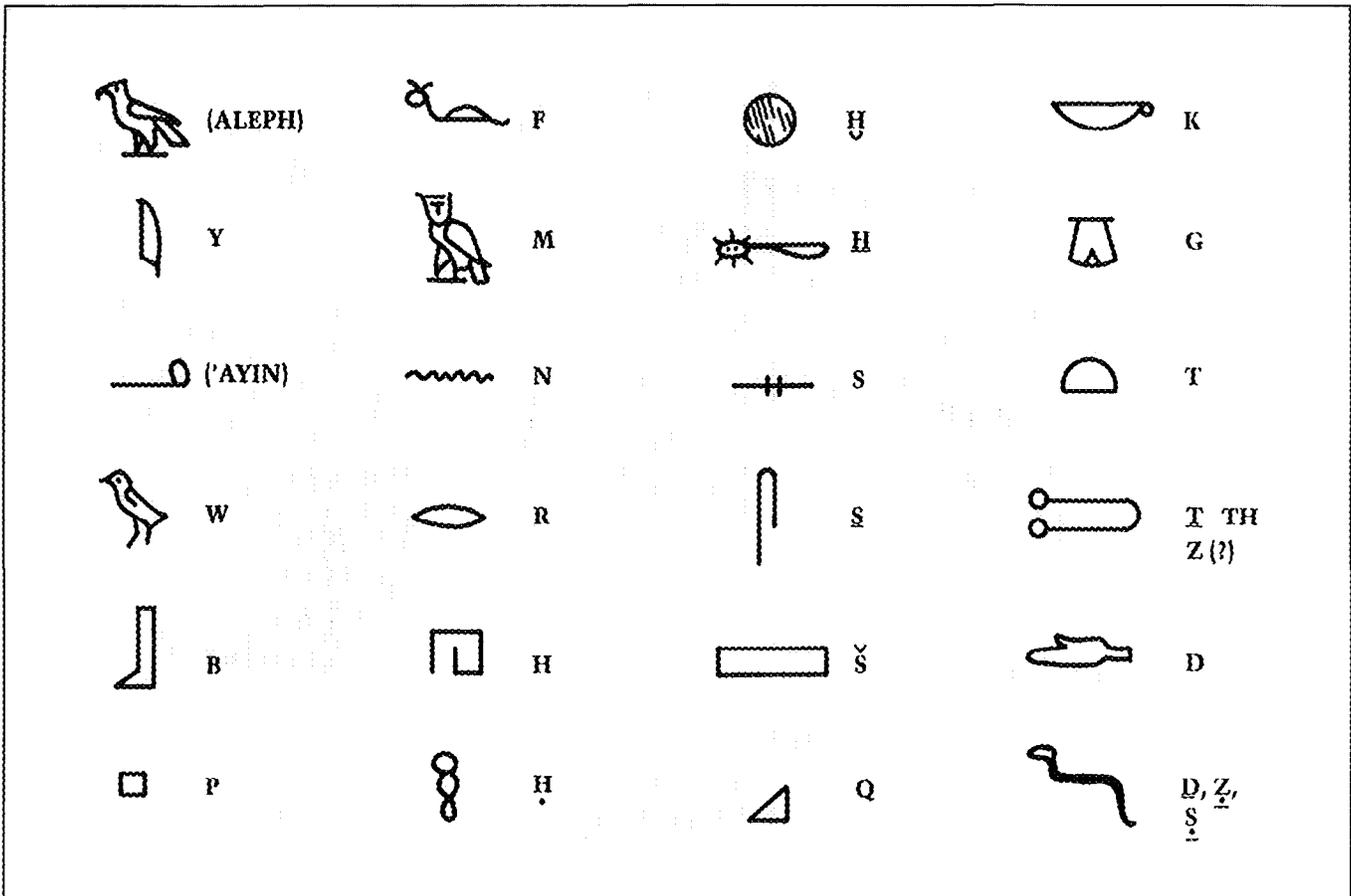
Conventions were used to decide which words could be replaced by signs, and for which words phonetic signs and consonants were required. Ancient writing did not divide words, however, so there was the remaining difficulty of how to know what meaning to attach to the sign sequences. For example, , *mn*, name, could also be read as *ro n*, the mouth of.

The solution to this problem was provided by determinatives—special

signs that were used at the end of various words to indicate the class of ideas to which the words belonged. Thus, there were determinatives for man, a profile of a sitting man; mouth, profile of a sitting man pointing to his mouth; abstract ideas, the first rays of the rising sun, and so on. Sometimes several determinatives were used simultaneously at the end of a word to pin down its meaning.

Determinatives were also used by Akkadians in Babylon, where a sign for GOD accompanies the spelling of a divine name. The closest equivalent for English is the capitalization of the initial letters of proper names.

Surprisingly, Egyptian hieroglyphics have deep structural similarities with modern writing. The use of metaphors is similar to the use of picture signs. Also, the use of metaphors implies a seeming confusion of meaning which can be cleared only by the context.



The Primary Egyptian Alphabet

Furthermore, there are sentences that are syntactically ambiguous but semantically clear.

In other words, the ambiguities inherent in language that arise due to multiple meanings of words, or the multiple ways a sentence can be parsed, are generally resolved by determining the context. To know the context is to be able to resolve the ambiguities and, therefore, we have a problem here that is inherently circular.

This circularity is what makes the learning of a new language, just by the use of grammar and a dictionary, such a difficult task. This is also why computer understanding of natural language is so hard. The difficulty is reduced somewhat by the use of devices such as 'scripts' that point the background context. Clearly, any such approach will work only partially.

I propose, therefore, the use of the idea of determinatives. As text is

recorded in a computer, the analysis of each sentence, or a part thereof, is also simultaneously recorded. This analysis may be done by a human, perhaps by the writer himself. The analysis is coded so as not to be printed into hard copy. This is not likely to completely resolve all ambiguities, but it would lessen them greatly. Further computer processing to understand the text would be considerably easier.

The Egyptian idea of determinatives, when used in the above described way for modern text, turns the problem of computer processing of text inside out, so to speak. This use of determinatives has another important benefit: It allows a distinction to be made between original and copy. The original in the disk could be used for analysis, whereas the bare text in the copy could not. Of course, there is the intriguing optional possibility that the original could be transcribed fully with the code information stored in

microdots that could be read by special optical readers

Subhash C. Kak
 Department of Electrical
 & Computer Engineering,
 Louisiana State University
 Baton Rouge, Louisiana 70803.

Another World View

Editor:

This letter notes some thoughts that occurred to me upon reading Roger Schank's "What Is AI, Anyway?" in the winter issue (*AI Magazine*, volume 9, number 2, Winter 1987).

Early in the article, Mr. Schank states, "The primary goal [of AI] is to build an intelligent machine." I would argue that an intelligent machine is in fact a contradiction in terms. I know that this is heresy in the AI community, but bear with me.

Let us look at just one of Mr.

Schank's "critical features" that an entity should possess to be considered intelligent, namely creativity. It seems to me that it is in the nature of creativity that it is unpredictable. I don't mean that it is unpredictable because we don't know enough in any given instance to be able to predict the outcome—I mean that it is inherently unpredictable. There is a discontinuous break from the past into the future. It's unpredictable because even if you knew everything about the state of things just before the creative act, you still wouldn't have enough information to predict it. Something new is brought forth, something that was not contained in what preceded it, either actually or potentially.

Machines, on the other hand, are inherently predictable. If you know everything about the state of a machine at time t , you can predict, in theory at least, its state at time $t+1$. This is what a machine is. Everything that a machine can be or do in the future is contained in its present state, either actually or potentially.

Let me give an example. It is easy to imagine a machine programmed to write music. There would be some sort of random generator of notes combined with rules regarding preferred intervals, chord progressions, etc. And in fact, such programs have been written, and do produce original compositions, or so I've read at any rate. It is within the realm of the imaginable that such programs could be refined to produce music equal to great human compositions. Don't we then have a machine that does something that we label creative when human beings do it? Perhaps. But it is also true that there are compositions produced by recognized composers that our machine could never produce. For example, John Cage's "Four Minutes and Thirty Three Seconds of Silence," in which the pianist simply sits silently for four minutes and thirty three seconds. The paradigms governing what constitutes music would of necessity be built into the program, and the machine would be unable to make the paradigm shift that lay behind Cage's work. It's conceivable that the program could have been accidentally programmed to produce such a piece, via a program bug, for

example, but that does not alter the argument. There is still nothing the machine can do to produce music that is not inherent in its program.

Now I am aware that this argument, as compelling as it is to me, will be dismissed by the vast majority of AI practitioners. And I think the reason is that there is a fundamental presupposition on the part of most people in the field that the brain is a machine and that human intelligence is produced by the working of this machine. Therefore, my reasoning is clearly flawed, as we can point to the existence of a counterexample to my conclusion. My belief is that this presupposition (which is, really, the fundamental presupposition of AI itself) is in fact incorrect—that intelligence is more than a product of the functioning of the brain. Thus we see that at the heart of AI lie philosophical issues of the most fundamental sort, because if my reasoning is in fact sound, it seems to imply that intelligence exists independently of matter.

Looked at from the perspective of Mr. Schank's primary goal (building an intelligent machine), the history of AI to date, for all its useful and interesting results, has to be regarded as a history of failure. We are farther from the goal than we ever were. From the perspective of his second goal (finding out about the nature of intelligence), however, there has been substantial progress. Ironically, each failure to achieve the first goal deepens our understanding of the second. Just compare Mr. Schank's lists of critical features of intelligence and problems of AI with some of the early ideas of what intelligence was (for example, the ability to solve problems). It is my view that AI will eventually die as a discipline, crushed by the weight of its failure to achieve the impossible, but that in the process it will have made an enormous contribution to the progress of the race: through its failure it will have rendered unavoidable the conclusion that intelligence is not a property of machines, either computers or brains. This will open up whole new realms of possibility for what we are and what we may become. There is a wonderful irony in this: the contribution will have been made by people committed to the

opposite point of view, a contribution that could never have been made by people who held the view all along that we are not machines. Such people are not found in AI labs.

Now I don't expect that your typical reader will immediately abandon his or her world view in favor of mine upon reading this letter. My hope is that I can raise in the minds of at least some of your readers the possibility that what they have been regarding as Truth (that the universe is wholly material) is in fact an assumption, a working hypothesis, if you will, that has not been proved, and about which there is at least as much evidence in contradiction as there is in support

Bruce David
31 Milford St.
Boston Massachusetts 02118

AI Education

Editor:

In a recent issue of *AI Magazine*, an article was published entitled "A Graduate Level Expert Systems Course," by David Brown. Although I found the article very interesting, I noticed the emphasis it put on theoretical AI issues, being directed toward computer science graduate students.

At the Department of Electrical and Computer Engineering, University of Kansas, we also offer a graduate level seminar in knowledge-based systems, but we take a very different approach. Our seminar is intended for engineering students who will either use knowledge-based systems in their future careers or will be called upon to design applications using existing techniques and tools. Because of the very diverse interests of the audience, the seminar covers a variety of AI techniques and systems, thereby exposing students to as many methodologies as possible.

I believe the approach we take better serves the interests and future careers of our engineering students. If you think your readership would be interested in a description of a graduate-level course for knowledge-based systems geared towards engineering students, I would be glad to submit an article discussing our approach



GURU:
When you expect the world from your expert system.

Tired of being limited? It's time to discover GURU—an expert system environment with comprehensive inference engine and rule management controls

Incredible flexibility.

GURU lets you design an expert system that runs exactly the way you want. By using fuzzy variables, certainty factors, reasoning rigor, rule selection order, and numerous environmental and utility variables, you have virtually unlimited control over the consultation environment.

Unprecedented development efficiency.

Using GURU's case saving and replay, you can track the effects that rule changes have on system behavior. You can also use meta rules to examine or alter other rules during a consultation. And using GURU's knowledge tree, you can display the relationships and dependencies between an application's rules, variables, and goals.

Quick and thorough.

By mixing forward and backward chaining, goal search time can be shortened dramatically. And, using GURU's multiple rule firing capability you can re-fire rules as values change. GURU also comes equipped with seamlessly integrated 4th generation decision support capabilities such as data base, spreadsheet, and report generator.

GURU runs on PCs, LANs, and VAXs.

To find out how GURU can exceed your expectations, call 1-800/344-5832 or 317/463-2581.



P.O. Box 248
 Lafayette, IN 47902
 1-800/344-5832
 317/463-2581

Explore the world of expert systems with GURU Tutor!

This full-featured development environment allows you to prototype GURU applications using rule sets, data bases, spreadsheets, and more for only \$75. To order, call us at 1-800/344-5832 or write VISA, MasterCard, and American Express accepted.

GURU is a registered trademark of mdba, Inc. VAX of Digital Equipment Corp.

For free information, circle no. 85

and giving an overview of the course.

Costas Tsatsoulis
 Electrical and Computer Engineering
 The University of Kansas
 Lawrence, Kansas 66045

We welcome letters from our members on this subject. In the meantime, we are reviewing Professor Tsatsoulis's article. —Ed

Publishing Cycles

Editor

I recently received an advertisement (the address label came from the AAAI membership list) for a newsletter called "AI Week." The brochure stated that it was published twice a month. I read no further.

You don't suppose they used AI to determine the publishing cycle?

Fred Kline
 PO Box 9041
 Seattle, Washington 98109

For free information, circle no. 53

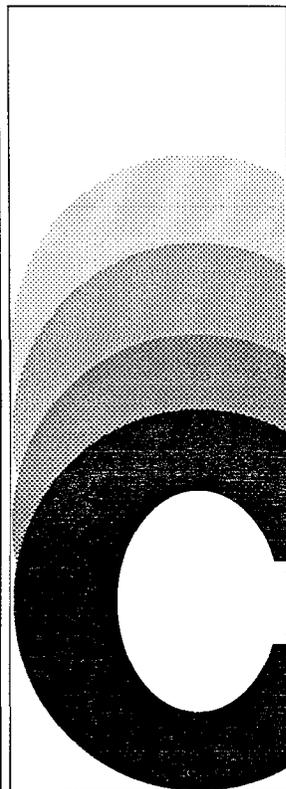
LISP LIBRARY FOR C - \$189

Adds symbolic processing to the C language • Written entirely in C
 Allows the addition of intelligence to existing conventional software
 Good documentation • Full source code • Portable • FAST • SMALL

RULE PROCESSING LIBRARY

\$79

Process LISP rules from C • Add expert systems capabilities to your C programs
 Extensive example



**DRASCH
 COMPUTER
 SOFTWARE**
 187 Slade Road
 Ashford, CT 06278
 (203)429-3817

MasterCard • Visa • COD • Next day shipment