



Letters

Reactions to Darden

Editor:

We are sympathetic to Lindley Darden's intellectual program. But the various conceptions of abstraction which she discusses are, individually and collectively, inadequate. There are two problems, one concerning the basic nature of abstraction and the other concerning a mechanism by which abstract terms can be related to their definitional base. Darden comes close to the basic nature of abstraction when she asserts that "In difficult cases, forming an abstraction can involve more than merely dropping parts or replacing constants with variables. New, abstract semantic concepts might have to be introduced." This, however, is not quite good enough and, in any case, she doesn't suggest a mechanism by which this can be accomplished. We wish first to discuss the fundamental problem and then to propose a computational mechanism.

Let us begin with Suppes's assertion that a set exists for any property, the so-called "axiom of abstraction." This axiom presupposes a determinable property. We believe that the problem of abstraction arises just because thinkers get intuitions of nondeterminable properties. On one level, the intuitive, the thinker can, for example, distinguish marvelous novels from hack work on a single reading. On another level, the level of conversation, the thinker can only waffle. It is remarkably difficult to explicitly describe the difference between good novels and trash. The reason, we suggest, is that the characterization of a marvelous novel requires an unbounded list of properties—infinite, as it were. This problem confronts physicists and biologists as well as literary critics.

A logician we know, Frederick B. Thompson, has told us about a favorite trick: Asking a biologist to define 'dog.' Thinkers who make progress with this problem seem to create decision procedures, or algorithms, that look at certain properties first, then others, and others again. It is, however, one thing to create algorithms which work for concrete things, such as 'dog' and quite another to create algorithms adequate for abstract things, such as 'marvelous-

ness.' In finite time, these algorithms may fail to decide whether the object in view (the novel) has the abstract property (marvelousness). What matters most is putting the most decisive properties at the head of the list.

When a thinker has an algorithm and wishes to communicate it, he or she must use words; we call the form of words required a 'rationalization.' Good rationalizations use words that name the most decisive properties, and syntax that reveals the working of the algorithm. Some thinkers have poor intuitions; some have good intuitions but make poor algorithms; some have good intuitions and make good algorithms, but utter poor rationalizations. When a thinker has a good intuition, makes a good algorithm, and utters a good rationalization, the world treasures it—sometimes for thousands of years.

As far as we know, the intuitive mechanism by which thinkers arrive at abstract characterizations is pretty much of a mystery, a mystery into which we can claim no particular insight, though we have tried to indicate a neuropsychological context (W L. Benzon and D. G. Hays, *Principles and Development of Natural Intelligence*, *Journal of Social and Biological Structures*, in Press). We do have something more tractable to say about

how rationalizations can be linked to the abstract terms which they help define.

We offer a simple mechanism, that of metalingual definition (D. G. Hays, *The Meaning of a Term is a Function of the Theory in Which It Occurs*, *SIGLASH Newsletter*, 1973, 6, 4:8-11). Consider a simple example: Charity is when someone does something nice for someone else without thought of reward.

"Charity" is the definiendum, "someone ... reward" is the definiens, and "is when" establishes the connection between them. Charity, the definiendum, is defined by a story, which is the definiens. This story is a rationalization, in the sense of the previous discussion. Any story which matches the pattern of the definiens is an instance of charity. The essential point is that charity is not, in any way, a term in the rationalization which defines it. Charity is defined by the patterning of elements in the defining story. While the rationalization is defined in general terms—"someone," "something nice"—the essential definition mechanism is not one of dropping detail, it is one of relating a term to a pattern of terms.

Note that this definitional process has to be recursive. While "someone" is a generalization over concrete things, people, "reward" probably is not. The variety of physical things which can serve as rewards is so great that we cannot meaningfully form a characterization of them; and abstract things, such as fame, can also be rewarding. Hence "reward" must also be abstractly defined. That is, it gets its meaning through a pattern of relationships between various objects and actions. And it is entirely possible that abstractly defined terms would appear in the definition of reward. However, eventually, we will arrive at a layer of terms whose meaning is defined over patterns among physical objects and processes. Thus this simple mechanism serves to relate abstract terms to concrete ones.

We call the mechanism "metalingual definition" out of respect for Roman Jakobson, whose concept of the metalingual function of language (R. Jakobson, *Linguistics and Poetics*,

In *Style in Language*, ed. T. A. Sebeok, Cambridge, Mass.: MIT Press, 1960, 350-377.) inspired the mechanism. In an AI context the mechanism involves the familiar process of defining a concept by using a frame, plan, script, whatever, with the provision that the definiendum be distinctly more abstract than the definiens. For example, "get food" can certainly be defined by an elaborate plan structure which has provisions for everything from getting a burger at MacDonalds, through shopping at the supermarket, to tilling the soil and skinning the deer. But that does not make "get food" abstract, for getting food is not, in our sense, essentially more abstract than each of the objects and processes in the "get food" frame. "Charity," however, is more abstract than the terms which define it. The same holds true, we assert, for such scientific terms as "mass," "phase space," "species," "covalent bond," etc.

If we now return to our discussion of abstraction's basic nature, the definiens in a metalingual definition is only a rationalization. And the rationalization is not, by the terms of that discussion, the source of insight from which the original intuition of an abstract property springs. That insight comes from noting "family resemblances," to use Wittgenstein's term, among various frames, scripts, plans, whatever, that have accumulated in the system. What we find deeply mysterious is the process by which this accumulated information is first sorted out by family resemblance and then rationalizations created.

This concept of abstraction has been developed in a number of papers (W. L. Benzon, *Cognitive Networks and Literary Semantics*, *MLN*, 1976, 91:952-982; W. L. Benzon, *Lust in Action: An Abstraction*, *Language and Style*, 1981, 14:251-270; D. G. Hays, *On 'Alienation': An Essay in the Psycholinguistics of Science*, in *Theories of Alienation*, eds., R. F. Geyer and D. R. Schietzer, Leiden: Martinus Nijhoff, 1976, 169-187; B. Phillips, *A Model for Knowledge and Its Application to Discourse Analysis*, *American Journal of Computational Linguistics*, 1978, Microfiche 82; and M. J. White, *Cognitive Networks and Worldview: The Metaphysical Termi-*

nology of a Millenarian Community, Ph. D. Dissertation, Department of Linguistics, State University of New York at Buffalo, 1975.), though most of these papers do not distinguish between rationalizing an abstraction and intuiting it. Applying this concept to science, however, does involve at least one subtlety about which we can comment. Natural science is devoted to understanding physical objects and processes, not things like charity, which do not appear to be physical at all. But that does not mean scientific concepts must therefore all be concrete concepts.

Consider the concept of carbon. It is a perfectly concrete thing, but we suspect that its characterization is, in fact, essentially abstract. Diamonds are concrete, and so is lampblack. Both consist primarily of carbon. But it is only relatively recently in the history of science that we have known that. Until the chemical revolution and the beginnings of modern atomic theory we had no way of discerning the essential identity between diamond, a hard, transparent, colorless pebble, and lampblack, a fine black powder. Only through the concepts and experimental technique of modern chemistry can we define what carbon is, as opposed to characterizing lampblack and diamond, where basic sensory categories are adequate. A similar point could be made about a shock from an electric eel and lightning. Both are manifestations of a single physical phenomenon, electricity, but their phenomenology is so different that this identity wasn't recognized until relatively recently.

For yet another example, one from biology, consider mankind and *homo sapiens sapiens*. The term "mankind" is embedded in a commonsense discourse in which human beings are essentially different from animals. The term "*homo sapiens sapiens*" is embedded in a scientific discourse in which human beings have evolved from animals and are therefore deeply similar to animals. To use a traditional terminology, the two terms may have the same extension, but their intensions are quite different.

Thus the essential distinction to the concept of abstraction is not so much that between physical objects

and processes and abstract objects and processes as it is a contrast between phenomenologically naive characterizations of the world and, well, abstract characterizations. By phenomenologically naive we mean simply those categories available through the sensorimotor apparatus of the human nervous system. "Black powder" and "hard, transparent, colorless pebble" are phenomenologically naive descriptions while "an atom consisting of 6 protons and 6 electrons" is not phenomenologically naive. These descriptions all characterize physical objects, but we have two different mechanisms linking the characterizations to the objects. One mechanism goes directly through the sensorimotor apparatus while the other mechanism, abstraction metalingual definition, loops through high-level linguistic and inferential structures.

This leads to our final point. Darden finds that AI's pursuit of commonsense knowledge is irrelevant to an understanding of scientific knowledge. We think this position is mistaken. By our account, phenomenologically naive commonsense knowledge is the ground from which all abstraction must start. The recursive nature of the metalingual mechanism makes it possible to move far away from that sensory base, but the base is necessary. Thus we feel that a deep understanding of science is impossible without a deep understanding of the mechanisms of commonsense knowledge. Yes, as Darden asserts, much commonsense knowledge is often badly mistaken. The intellectual mystery lies in understanding how, through abstracting over this mistaken mass of knowledge, science has been able to create ever more effective abstract accounts of the world. We think the mechanism of metalingual definition is one tool in the set of tools needed to convert that obscure mystery into a set of solvable intellectual puzzles.

William Benzon
David Hays
161 Secon Street
Troy, New York 12180

Darden's Response

Editor:

Benzon and Hays raise deep and important issues related to human cognitive endeavors in forming abstract concepts, such as charity and carbon. However, my purpose in the section on abstraction in "Viewing the History of Science as Compiled Hindsight" (*AI Magazine*, Volume 8, Number 2, Summer, 1987) was more modest. My purpose was to do a conceptual analysis of abstraction as it is used in current computational AI work. Whether or not the methods currently used in AI mimic human methods of abstraction is not an issue which I addressed in the article, nor do I know the answer to it. Personally, I am more interested in creative computational methods being developed in AI for representing knowledge and reasoning, quite apart from whether they tell us anything about human cognitive processes. My goal in discussing abstraction, and distinguishing it from simplification and generalization, was to begin a discussion of its nature and methods for doing it. Those methods are of interest to me because they may be useful in representing and using the compiled hindsight that the history of science can provide. I suspect both Darwinian natural selection, which served as a particular instance of a selection type theory, as well as the abstraction for selection that I discussed, would both be abstract concepts in Benzon and Hays's sense of that term.

My purpose was not to discuss how abstract scientific concepts are formed from or justified by observational evidence, although those are topics that I, as a philosopher of science, find quite interesting. My discussion of abstraction was much more limited than the topic of the relation between theoretical terms and observation terms and how the former may be grounded in the latter. That has been a serious concern of philosophers since Locke and Hume formulated their empirical philosophies and a quite specific concern of philosophers of science in this century. (See, for example, Frederick Suppe's *The Structure of Scientific Theories*, 2nd. ed., Illinois University Press, 1977.) Ben-

zon and Hays's concept of "metalingual definition" sounds quite similar to the network model for theories discussed by Mary Hesse (*The Structure of Scientific Inference*, University of California Press, 1974) in which the terms of a theory are seen as part of a network, with some terms grounded in observable predicates, while others receive their meanings through their interconnections with other more theoretical predicates in the network. I am unsure how the abstract formulations of theories that I discussed relate to the issue of definitions of abstract terms. My concern was with methods for forming abstractions of theories and determining conditions for guiding their instantiation, not with how to define the abstract terms embedded in them.

Benzon and Hays say that I provide no computational mechanism for substituting a more abstract concept for a more particular one in forming an abstraction for a theory. I said that such a task might be difficult to implement, but I did make one suggestion. Find an "is a" hierarchy in which the particular concept is located and use the next higher concept. This method will certainly require additional guidance because a concept may be embedded in numerous "is a" hierarchies, none of which might provide an adequate higher level concept for in a particular theoretical context. Thus, Benzon and Hays are correct that more work needs to be done on methods for finding appropriate higher level concepts.

As to commonsense: I hope that science can be understood without a deep understanding of commonsense. Endeavors of philosophers of science for the last fifty years show that it is not easy to determine criteria for demarcating science from nonscience. My guess is that it will be much more difficult, if not impossible, to delineate what counts as commonsense. I decided long ago that I have no idea what commonsense is, since my mother often scolded me with the remark: "Lindley, that just didn't show any commonsense!"

The issue about how humans recognize similarities so as to group things by family resemblance is a deep one. Those of us working on the problem

of analogy are aware of the computational difficulties of detecting similarity. If similarity can be analyzed into identity and difference (perhaps identity at a more abstract level), it is more tractable. (See L. Darden, "Artificial Intelligence and Philosophy of Science: Reasoning by Analogy in Theory Construction," in *PSA* 1982, v. 2, ed. T. Nickles and P. Asquith, Proceedings of the Philosophy of Science Association, pages 147-165 and Dedre Gentner, "Structure Mapping A Theoretical Framework for Analogy," *Cognitive Science* 7 (1983) 155-170.) But Benzon and Hays quite rightly point out that mysteries abound as to how humans group similar things and then form more abstract concepts that capture essential or typical features. AI can usefully explore computational mechanisms connecting methods of inducing concepts from examples, placing them in "is a" hierarchies, and using those hierarchies while forming more structured abstractions, such as one for selection type theories. The deeper concerns raised by Benzon and Hays about humans' abilities to form abstract concepts may provide useful problems and difficult challenges for cognitive science and AI.

Lindley Darden
Committee on the History
and Philosophy of Science
University of Maryland
College Park, Maryland 20742

Reactions to Letovsky Article

Editor:

I have just finished reading the Conference Report from the Fall 1987 issue of *AI Magazine* titled "Ecclesiastes: A Report From the Battlefields of the Mind-Body Problem." Instead of detailing every frustration I found with it, I would simply like to express my disappointment with your magazine for publishing such a poorly written report. I do not read *AI Magazine* to hear about the social activities of a conference attendee and his personal biases toward the participants, but rather to have described, in a neutral manner, the theories and opinions presented. I hope that in the future you will take more care to see that

such sloppy journalism is not included in your magazine.

David K. Young
BBN Laboratories Incorporated
10 Moulton Street
Cambridge, Massachusetts 02238

Editor:

Your Conference Report articles serve a very useful purpose, and their informal style makes them quite interesting and readable. However, the conference report in your Fall 1987 issue, "Ecclesiastes: A Report from the Battlefields of the Mind-Body Problem" by Stan Letovsky, had several major deficiencies. The author violated basic concepts of reporting for a technical magazine: objectivity, completeness, and reasonable lack of intrusion of the reporter himself.

Letovsky was rude to the sponsors of the conference, southern evangelical Christians, who had somehow succeeded in bringing together a rather diverse and interesting group of speakers: "Next morning the proceedings opened with remarks by a man who faintly resembled Jerry Falwell... I occupied myself by passing notes back and forth with the student next to me wherein we communicated wittily on the situation..."

There was a strong "AI groupie" tone to the report: "I ran into Minsky, which is rather like a 10-year-old boy running into his favorite World Wrestling Federation star on the street..." "I sat at a table with Minsky and Moravec and traded favorite scientific and science fiction ideas till late in the evening. It was heaven." The "Dear diary" type of exposition is out of place in a technical magazine.

The author often provides a complete summary when one of his heroes speaks: Minsky, Moravec, and particularly Letovsky. However, the author is remarkably vague about the content of key talks. Sometimes this is due to the author's prejudices in the debate between the dualists (the mind is distinct from the brain) and the monists (the mind can be explained entirely in terms of the brain): "...but I was no objective observer. The dualist papers seemed like pseudoscientific mysticism." Sometimes the vagueness is due to poor note taking: "Arbib pro-

ceeded to the main part of his talk, which I am sorry to say I don't remember at all." Sometimes, it is due to AI-groupiness: "(Lenat) is currently involved in a project to codify over a decade a tremendous body of commonsense knowledge in machine-usable form. As usual, his presentation was polished, his slides cute. The audience was mesmerized."

Conference reports, while giving some local color, should stress content. This report, while capturing the near-orgasmic excitement of the author on being near his heroes, left much to be desired in content, particularly as far as the dualist point of view is concerned, and told me more than I wanted to know about the author's innermost thoughts.

More editorial control will have to be exerted to eliminate sophomoric articles of this kind if *AI Magazine* wants to be considered a professional publication.

Oscar Firschein
SRI AI Laboratory
333 Ravenswood Avenue
Menlo Park, California 94025

Editor:

After seeing Volume 8, No. 3 (Fall 1987) of *AI Magazine*, I can honestly say that the review of the AI and Human Mind conference by Stanley Letovsky was the most unprofessional article I have ever read.

Letovsky started by describing the opening remarks by a "gentleman ... [with a] southern accent," whose name he never bothered to mention, and who introduced Henry Margenau in whose honor the conference was given. Letovsky had "never heard of" Margenau, and "wandered into the hall" before the talk was over. Then, based on second-hand accounts of the talk "at lunch with some other AI graduate students," Letovsky proceeded to characterize Margenau's views as "obvious and obviously wrong," and criticized at length what he assumed Margenau had said.

Letovsky's own views were rather naive. His description of quantum mechanics, for example, had "particles moving along probabilistic paths," a concept that is sufficient for introduc-

tory physics courses but is inadequate for serious discussions in which the more precise concept of state superposition should be used. Although such naivete is itself unobjectionable, it compounds Letovsky's offense of reviewing a talk from second-hand information.

But my primary concern is not Stanley Letovsky and his review.

My concern is that *AI Magazine* would publish this material. It insults the organizers and participants of the conference. It hurts *AI Magazine* by making suspect its editorial judgement; which in turn hurts AAAI and hurts contributors by lowering the prestige of publication in that journal. It also insults readers who, when they read *AI Magazine*, expect to be informed, not offended.

According, I strongly feel that *AI Magazine* should:

1. Publish an apology for this lapse in editorial judgement.
2. Seek a review of the conference by someone with the maturity and knowledge, both in AI and physics, to do a professional job.
3. Establish editorial policy to reject this kind of material in the future.

Besides owing this to the conference, you owe this to people like myself who, by being members of AAAI, have been made party to disseminating the unprofessional, offensive tripe written by Mr. Letovsky.

Leonard R. Kasday
158 Oak Creek Road
East Windsor, New Jersey
08520

Editor's Response

Letovsky's piece was intended to be an entertainment, not a serious conference report. I enjoyed reading it, for its discussion (clearly biased) of the monist-dualist controversy, and for its youthful enthusiasm. It is not representative of a "mainstream" article in the *AI Magazine*, nor will it ever be, as long as I'm the editor. Once in a while, however, I will accept something that's controversial and/or humorous, and perhaps even outrageous.

Bob Engelmorc

Editor

Gee. I really liked the Letovsky essay and was delighted that you published it. But I must admit to being biased because I came out rather well in it.

And that's my point. I think it is fine to say that the essay was controversial and, to many readers, outrageous. However, I don't care for the use of the word "biased". Letovsky certainly takes a position about dualism, and I happen to agree with him; so you might call me biased, too. However, in using the particular word "bias", you, in a subtle way, might appear to move the discussion from the scientific to the religious ground.

What I mean is that if we were not immersed in a culture of people 99% or so of whom come from a dualist (religious) intellectual tradition, Stanley's remarks would be considered less controversial. It was indeed astonishing to find that the conference had fundamentalist undertones. In using the word bias, there is a hint about prejudice. What I mean is that it suggests an orientation that comes from somewhere outside the appropriate context.

Indeed, if we were in a humanistic frame of discussion, in which different positions on religion and such matters were considered to all have equal merit, then I would agree that Letovsky, and myself as well, are clearly biased. However, it seems to me that in the frame of the *AI Magazine*, dualism has no special scientific or technical merit. In that sense, I don't see him as showing any particular "bias". Perhaps the word "disrespectful" is more appropriate because it doesn't imply that he's actually wrong.

What fun!

Marvin Minsky

MIT

Cambridge, Massachusetts

Is AAI A Closed Society?

Editor.

AAAI currently has the appearance of a closed society, with little information flowing from the officers and councillors of the society to the members, and little solicitation of input from the members by the officers and

councillors. This lack of communication is obvious in the official channels that could be used for communication from the officers and councillors to the members.

The first of these channels is the annual business meeting. This meeting could be used to summarize the deliberations that the officers and councillors participated in over the last year, including the deliberations at the executive council meetings, for presenting official statements, such as the treasurers report, and for inviting comment by the members.

Instead the annual business meeting is a joke. There is often no published agenda. Only a token summary of the executive council meeting is made. Only verbal overviews of official statements are given. Few councillors are there to listen to any comments by members. For these and other reasons, few members attend the business meeting.

The second official channel for communication is the *AI Magazine*. The purpose of the *AI Magazine* is even stated in its editorial policy: "The purpose of the *AI Magazine* is ... and to keep its readers posted on AAI-related matters."

This purpose is not very well served. For example, in the current issue of the *AI Magazine* the "News from the American Association for Artificial Intelligence" section consists of one page and is located near the back of the magazine. The last time that the minutes of any AAI meeting were published in the *AI Magazine*, was the publication of the minutes of the Fifth Annual AAI Meeting in the Summer 1985 issue. The last time that any useful information about AAI, besides information on conferences, workshops, and nominations for officers and councillors, was published was also in the Summer 1985 issue.

The third official channel for communication from the officers and councillors of AAI are direct mailings to the members, either postal or electronic. Without timely information, there is no way that members can make informed choices during elections. I cannot remember any direct mailings to members that were not conference or workshop related,

except for the ballot mailings.

In sum, none of the possible official channels of communications to members is being effectively used to transfer information from the officers and councillors of the association.

I strongly feel that this lack of communication is bad for the association. More information flow from the officers and councillors to the members would lead to more understanding of just what the association does. More input from the members would lead to beneficial changes in the goals and procedures of the association. I suggest that the officers and councillors of the AAI take immediate steps to present more information about the activities of the association to the members, and that they solicit more input from the members.

As one step toward this end, the "AAAI News" section of the *AI Magazine* should be expanded to include information about AAI and its current activities. At a minimum, it should include summaries of all deliberations of the council and committees of AAI, the minutes of any formal AAI meetings (including executive council meetings and annual business meetings), full results of elections, and all formal reports by officers of AAI. The section should also be relocated (along with the related "Announcements" section) to just after the table of contents, where it would be more visible. Comments by members on any of the information presented should also be solicited.

The information in this section should be printed as soon as possible, ideally in the first issue following the time that the information is available or the first issue following the activity (if necessary, such as in the case of financial statements, preliminary information should be printed as soon as it is available, and revised or audited information would be printed in a later issue.) The preparation of this information should be a responsibility of the officers and committee chairs of AAI. In this way the information will come directly from those who have the privilege of making the decisions, will have more impact than if simply reported on by someone else, and will contribute more to a feeling of openness in AAI.

This change would go a long way toward giving AAAI the appearance of an open society. However, other changes, such as better attendance by councillors at annual business meetings, should also be made to further foster communication between the officers and councillors of AAAI and its members.

Peter Patel-Schneider
Schlumberger Palo Alto Research
3340 Hillview Avenue
Palo Alto, California 94303

Dr. Patel-Schneider:

Thank you for your letter describing your concerns about the appearance of the AAAI as a closed society because of the apparent lack of communication between the officers and the members

In its short history, the AAAI has developed into a fluid, informal organization. That informality is often the character of a scientific society rather than a professional association. Issues arise in scientific societies which are normally not crucial or contentious, but deal with the organization's objectives which are reflected in its programs and services. Because of this situation, very few people attend the annual business meeting although we do announce it either in the magazine or by mail. The discussions during the business meetings are admittedly less formalized than one would expect, but it simply reflects the basic character of the AAAI.

The members elect officers and councilors to define the association's strategic direction, develop policies to implement that direction and monitor its implementation. The members put a certain amount of trust in the councilors' capabilities to accomplish these objectives. When public discussions between members and officers on every decision or piece of policy occurs, only confusion results. Operations are put into a standstill. A consistent flow of information about the officers' deliberations and new programs is important to all parties, but minute discussions about them only clog the system and diffuse the real issues.

Because the primary objective of our programs and services is to com-

municate scientific and technological information principally through the magazine, the NCAI's technical, exhibits and tutorial programs and workshops, naturally our direct mailings would only be concerned with the content of those services. Communication from the officers about policies are represented in minutes from the committee meetings and executive council meetings, which admittedly haven't been published on a regular basis (1986 Executive Council Meeting Minutes were published in the Fall 1987 issue). Beginning with the Winter 1987 issue, we will have a special editor (and councilor), William J. Clancey, who will produce the AAAI News section every issue. Not only will committee and council deliberations be published, but other association issues will be discussed. The Publication Committee is also considering the establishment of an Opinion Column in which members' concerns and issues may be expressed to the rest of the membership.

Like other societies, it is customary to publish final or audited financial statements when they are available. Since 1985, we have audited our financial statements which have been published (e.g. the 1986 audited financial statement was published in the Fall issue). We will continue to publish the audited financial statement and will try to publish some preliminary financial data as it becomes available. But, I would caution that the releasing of preliminary financial information can be misleading and not reveal the real financial soundness of the corporation.

We appreciate your noteworthy comments and will try to use them as a guide for the expansion of communication within the AAAI.

Claudia C. Mazzetti
Executive Director
AAAI
445 Burgess Drive
Menlo Park, California 94025

Ms. Mazzetti:

I applaud the expanded role of the AAAI News section of the *AI Magazine*.

I feel that this and other moves will expand the communications between

the officers and councilors of the AAAI and its members to a useful level.

Peter F. Patel-Schneider
Schlumberger Palo Alto Research
3340 Hillview Avenue
Palo Alto, California 94303



Editor:

The program works! We found one!
Edward Feigenbaum
Penny Nii
Knowledge Systems Laboratory
Stanford University
701 Welch Road
Palo Alto, California 94304

See the "Blackboard Systems" articles published in the 1986 Summer and Conference issues of *AI Magazine*.
-Ed

Nonmonotonic Reasoning in a Possibilistic Setting

Editor:

In the literature of AI, nonmonotonic reasoning is usually illustrated through the familiar Tweety example or its variants (P. Cohen, and E. Feigenbaum, *Handbook of Artificial Intelligence*, Los Altos, CA, Kaufmann, 1982). A shortcoming of the Twenty example is that the premise *birds can fly* is in reality a proposition in probabilistic logic rather than first-order predicate logic. Thus, what the Twenty example involves is a revision of probabilities rather than truth values.

An example of nonmonotonic reasoning which involves a revision of possibilities and is simpler and more straightforward than the Twenty example is the following:

1. Carol resides in the United States.

2. Boston is in the United States.

From (1) and (2) it follows that

(a) it is possible that Carol resides in Boston.

Now let us add the premises:

4. Carol resides in California.

5. Boston is not in California.

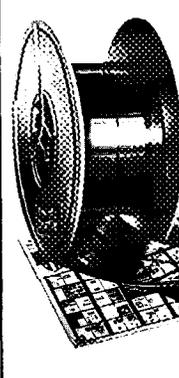
Then, from (1), (2), (3) and (4) it follows that

(b) it is not possible that Carol resides in Boston.

We observe that (b) is the negation of (a). Thus, the reasoning in question involves a reversal in the truth value of a possibility-qualified proposition. Hence, it is an example of nonmonotonic reasoning. It should be noted that the concept of possibility is related to that of nondisprovability in the modal logic models of nonmonotonic reasoning (D. McDermott and J. Doyle, Non-monotonic logic I, *Artificial Intelligence*, 1980, 13, 41-72). However, in the theory of possibility, possibility is a matter of degree and gives rise to a conceptual structure which is different from that of modal logic (L. A. Zadeh, Possibility theory and soft data analysis, in L Cobb. and R. M. Thrall, *Mathematical Frontiers of the Social and Policy Sciences*, 69-129. Boulder, CO: Westview Press, 1980) (D. Dubois and H. Prade,, *Theorie des Possibilites*. Paris: Masson, 1985).

This research supported in part by NSF Grant IST-842416.

L. A. Zadeh
Computer Science Division,
University of California,
Berkeley, California 94720.



This publication is available in microform.

University Microfilms International reproduces this publication in microform: microfiche and 16mm or 35mm film. For information about this publication or any of the more than 13,000 titles we offer, complete and mail the coupon to: University Microfilms International, 300 N. Zeeb Road, Ann Arbor, MI 48106. Call us toll-free for an immediate response: 800-521-3044. Or call collect in Michigan, Alaska and Hawaii: 313-761-4700.

Please send information about these titles:

Name _____

Company/Institution _____

Address _____

City _____

State _____ Zip _____

Phone (____) _____

University Microfilms International

ARTIFICIAL INTELLIGENCE

LEHIGH UNIVERSITY

WIESEMAN CHAIR

Lehigh University is accepting applications for a newly endowed chair, the Wieseman Professorship of Computer Science and Electrical Engineering. We seek a recognized authority in Artificial Intelligence to help shape a growing and dynamic activity into one of international recognition.

The Wieseman Chair will support an intensive research program in the knowledge-based systems fundamental to our cross-disciplinary activities. These include development of expert systems for Lehigh's NSF Engineering Research Center (ATLSS) and the Semiconductor Research Corporation Electronics Packaging Center.

Candidates should send a curriculum vitae and the names of six references to:

Professor L. J. Varnerin
Department of Computer Science & Electrical Engineering
Packard Lab #19
Lehigh University
Bethlehem, PA 18015

Lehigh University is an affirmative action/equal opportunity employer.