

ARTIFICIAL INTELLIGENCE: ENGINEERING, SCIENCE, OR SLOGAN?

Nils J. Nilsson
Artificial Intelligence Center
SRI International
Menlo Park, California 94025

ABSTRACT

This paper presents the view that artificial intelligence (AI) is primarily concerned with propositional languages for representing knowledge and with techniques for manipulating these representations. In this respect, AI is analogous to applied mathematics; its representations and techniques can be applied in a variety of other subject areas. Typically, AI research (or should be) more concerned with the general form and properties of representational languages and methods than it is with the content being described by these languages. Notable exceptions involve "commonsense" knowledge about the everyday world (no other specialty claims this subject area as its own), and metaknowledge (or knowledge about the properties and uses of knowledge itself). In these areas AI is concerned with content as well as form. We also observe that the technology that seems to underly peripheral sensory and motor activities (analogous to low-level animal or human vision and muscle control) seems to be quite different from the technology that seems to underly cognitive reasoning and problem solving. Some definitions of AI would include peripheral as well as cognitive processes; here we argue against including the peripheral processes.

I. INTRODUCTION

Artificial intelligence (AI) is a large and growing field. Graduate students study and perform doctoral research in AI at many universities throughout the world, and scientists and engineers at academic and other research centers contribute to AI's body of concepts and techniques. Industrial organizations are not only beginning to apply AI ideas to manufacturing technology, but are using them in an increasing number of new products. There are AI organizations and societies, AI textbooks, AI journals and magazines, and AI meetings. What is the nature of all of this activity? Is AI a coherent subject area? If so, is it a part of computer science, of psychology, or of philosophy; or is it actually an amalgam of parts of these subjects and perhaps others? Is AI

primarily a scientific discipline whose goal is to gather and analyze knowledge about intelligent behavior, or is it an engineering enterprise whose goal is to synthesize intelligent artifacts? Is AI getting anywhere?

There are a number of reasons it is important for those of us who are involved in AI to ask and answer questions like these. First, several people outside AI, whose opinions command well-deserved respect, are inclined to provide answers that challenge our views. Some think that AI research is rather vacuous -- based more on slogans and showmanship than on achievement. AI researchers ought not to forfeit to others the task of defining our goals and prospects and describing our accomplishments. Moreover, progress in a field and the view of that field held by its practitioners are highly interrelated. There are undoubtedly some views of AI that are more fruitful than others. We ought to be guided by the most productive paradigms. Finally, as a field mature, it becomes possible to teach its accumulated knowledge as a set of goals, methodologies, and principles. We need a coherent picture of AI if we are to teach it to students. For these reasons, at least, it is important to keep asking questions about the nature of AI.

II. A BROAD AND A NARROW VIEW

Our first task is to delineate the subject matter of artificial intelligence. There are a number of reasonable alternatives in characterizing the subject matter of any field, and I am sure that boundaries are never set very exactly or with universal agreement. We will not be concerned here with precision, but I think there are two rather different choices

First, we could take AI to encompass *all* those processes that account for intelligence or adaptive behavior in humans and other animals. Among the many activities that underlie such behavior, it seems reasonable to distinguish between peripheral and central processes, in which the *peripheral* ones are those that are quite close to

the boundary between the environment and the animal or machine that inhabits it. Peripheral perceptual processes, for example, might include optical or acoustic transduction, as well as the *first* stages of image or auditory-signal processing. Peripheral output processes might include motor routines and the feedback loops that contribute to short-term stability. If AI were defined to subsume peripheral processes, it would include the work of Marr [1] and of Barrow and Tenenbaum [2], for example. By the same criteria, it should also then include much of the work done in acoustic processing of speech sounds [3].

With regard to humans, I am inclined to consider as *central* those cognitive processes that are involved in reasoning and planning. Work on automatic methods of deduction, commonsense reasoning, plan synthesis, and natural-language understanding and generation are examples of AI research on central processes.

Some would define AI broadly to include both peripheral and central processes. I admit that it is often difficult to separate them. Psychological research has shown that higher processing often controls the biases and modes of peripheral processing. Besides, AI research itself has shown that it is usually not productive to arrange processing in isolated levels. Nevertheless, it seems to me that the kinds of techniques and representations employed in the early stages of image processing are vastly different from those employed, for example, in commonsense reasoning. Put simply, I fear that the variety of processes underlying adaptive behavior is just too large to constitute a coherent, single field of study. For this reason, I think that AI will ultimately fracture along a cleavage line somewhere between the most peripheral and the most central processes. Consequently, I am inclined to a somewhat narrower definition of AI -- one that encompasses the central processes only and leaves the study of peripheral processes to other disciplines.

I shall probably not be able to give an entirely satisfactory definition of the "central" processes that I think ought to be included in AI, but I hope that by the end of this paper the reader will have a "rough feeling" for them. I certainly mean to include those reasoning processes of which humans are "conscious". By doing so, I include, for example, the techniques used in computer programs for diagnosing diseases [4], evaluating ore deposits [5], and determining chemical structure [6]. These kinds of programs are commonly called *expert systems*. Yet there are many (presumably central) processes that apparently require reasoning abilities of which humans do not seem to be conscious. Such processes as are involved in understanding and generating natural-language sentences are examples. I assume that "unconscious" efforts such as these are among those to which Knuth [7] refers when he says: "I'm intrigued that AI has by now succeeded in doing essentially everything that requires 'thinking' but has failed to do most of what

people and animals do 'without thinking' -- that, somehow, is much harder!"

Perhaps as important as the processes themselves is the "knowledge" they manipulate. In fact, the subject of knowledge representation formalisms is a good starting point for a more detailed explanation of just what I think AI is.

III. BUILDING ON THE NARROW VIEW

Whatever else we decide to include under AI, I would like to join those who claim for it the task of developing, maintaining, and using computer-interpretable formalisms for representing knowledge. In this view, AI includes what several have called *applied epistemology*. Stated thus, the claim may seem rather audacious. Perhaps a somewhat less contentious way of putting it is to say that AI research indicates that it is now possible to have a unified field of study tantamount to applied epistemology. I am saying here that we may as well call this field *artificial intelligence*, since it comprises so much existing AI research and so few persons outside AI are working on epistemological problems with the same precision and scope.

Before we talk about the kinds of knowledge representation formalisms being studied in AI, it may be useful to mention two rather well-known, conventional media for representing knowledge, namely, natural language (such as English) and mathematical notation. The first of these is known to be difficult to use with precision and is not yet sufficiently computer-interpretable. Mathematical notation can, of course, be interpreted by computers; despite its precision, however, it lacks the power to say much that can be said in English. Because I want to assign to AI, among other things, the role of "keeper" of advanced knowledge representation formalisms, we can obtain some insight into the nature of AI by investigating the roles played by the "keepers" of these more conventional languages.

Let us take natural language, for example. We can express almost any kind of knowledge in English. A high school or college student takes many different courses, say, physics, biology, chemistry, and history. A large part of what he learns about these subjects is expressed in English. The student might also take a course in English. This subject, English, is about a language for expressing knowledge about all of these other subjects (and about English itself). In English class, a student learns about syntax, style, exposition, and so on. These topics are studied more or less independently of the subject matter of the sentences used to illustrate the topics. There is an interesting interaction between a subject such as physics and English. As new physical phenomena are discovered, new English words and phrases must be invented (or old

ones, such as "charm", given additional, appropriate meanings). Conversely, the more enriched English becomes with words and phrases denoting a wealth of concepts, the better able it is to describe physical phenomena. Interconnected though they might be, we usually have no trouble determining whether we are studying physics or English.

There is a large amount of knowledge about physics, chemistry, and many other subjects that is difficult or cumbersome to express in English, yet can easily be expressed using mathematical notation. Here too, we have a special field of study based on this notation, namely mathematics. Applied mathematicians concern themselves with representational problems in a discipline such as physics, and develop mathematical forms and techniques for resolving such problems -- thus contributing both to mathematics and to physics. Admittedly there is often a fuzzy boundary between applied mathematics and its object of treatment, such as physics, biology, or chemistry. Just as mathematical techniques are often illustrated by physical examples, so do physicists need to be intimately familiar with mathematical techniques to make progress in their field. The same person may sometimes pursue the goal of an applied mathematician (namely, to use and enrich mathematical formalisms and techniques) and, at other times, pursue the goal of a physicist (namely, to know and understand the physical universe). Nevertheless, these goals are different, and it remains reasonable to treat mathematics as a subject that is separate from those to which the language of mathematics can be applied.

One might view some of the representational formalism being studied in AI as attempts to create languages that possess the precision and computer-interpretable properties of a mathematical notation, but to do so for a much wider range of concepts than those dealt with by classical mathematics. Mathematical notation is useful for denoting numbers and certain other kinds of abstract structures that have a quantitative character. Yet, there are many other kinds of knowledge, which English seems to be able to represent (albeit somewhat imprecisely), that are not representable in conventional mathematical notation. AI research seeks representations that can be given a declarative interpretation, like those of mathematics, for expressing nonmathematical knowledge. Many potentially useful formalisms exist, including logical formulas, rewriting rules, semantic networks, production rules and other declarative notations. To give them a name, let us call them *propositional* representations. I predict that one of the chief concerns of AI will be the use of these representations (and such extensions as may prove necessary and beneficial) to represent and reason about -- well, about anything at all.

In addition to its preoccupation with logical and other propositional formalism, AI is vitally interested in

ways of manipulating these formalisms and in computational techniques for doing so. This latter concern gives AI a definite engineering character. It is the creator not only of powerful knowledge representation languages, but also of techniques and systems that manipulate knowledge to produce useful results. In this respect, AI is analogous to a combination of applied mathematics, numerical analysis, and those portions of computer systems technology that are concerned with the algorithmic languages used for solving mathematical problems. We have something pretty close to my perception of AI if we form a similar combination based on propositional formalisms, rather than on conventional mathematics.

Thus, AI seeks to express knowledge by using propositional formalisms, representing them as a computer data structures that can be manipulated flexibly and efficiently. To accomplish this objective, AI must be concerned with the relevant computer languages, systems, and environments.

Let us try to describe some of the kinds of knowledge AI researchers are attempting to represent. First, we have the kind of knowledge that, in English, might be represented by sentences stating particular facts about some situation. An example from economic geology (the field concerned with locating and evaluating mineral deposits) is "The geologic prospect is cut by a thoroughgoing fault system". Sentences of this kind can be viewed as expressing relations among individuals in some domain and, perhaps, functions of individuals. Next we have more general knowledge that refers either to indefinite individuals or, universally, to all the individuals in some set. Examples are "all granitic rocks are igneous" and "some sedimentary strata are oil-bearing". AI turns to a formal logical language, such as first-order logic, in which to express sentences like these.

Whether expressed in English or in logic, the knowledge about a subject typically requires a very large number of sentences to express it. AI is also quite concerned, therefore, with devising efficient methods for storing and retrieving knowledge expressions. As knowledge about a field grows, additional sentences are of course required. More importantly, the knowledge representation language must be expandable to include expressions for new individuals, functions, and relations. These can either be precisely *defined* in terms of existing more primitive concepts, or they may simply derive their meanings from the set of expressions in which they participate.

Hierarchical structuring of knowledge serves a key role in its efficient use and retrieval. Experts in particular fields are able to recognize as single concepts what novices in those fields would consider rather complex patterns of primitive concepts. Part of the growth of a specialized discipline involves the invention by experts of

these higher-level ideas. For example, Rich and his colleagues [8] have been exploring the use of hierarchical structures to represent the knowledge utilized by skilled computer programmers in designing programs. A skilled programmer can look at programs and understand them more thoroughly and quickly than would a novice, because the expert sees them as embodying concepts he knows about -- such as tail recursion or initialization, for example. In some fields, many of the constructs that contribute to expertise are not easy to express in English and are not learned explicitly by students. Instead they are absorbed more or less haphazardly "along the way". One of the very important applications of AI is in helping specialists to make explicit and then express this kind of expert knowledge in logical languages.

Besides expressing the technical knowledge of physics, chemistry, computer programming, civil engineering, or other such fields, AI is concerned with elucidating and representing what might be called *commonsense knowledge*. The commonsense knowledge possessed by all humans involves, perhaps, hundreds of thousands of facts like "objects fall unless they are supported", "material objects do not suddenly disappear", and "one can get wet in the rain". Hayes [9] has been attempting to find representations for some of the commonsense knowledge that we all have about physics, especially about liquids. He calls this body of knowledge *naive physics* to differentiate it from more technical and mathematical physics. We also need to be able to represent everyday knowledge about time, the history of events, and alternative courses of events. McDermott has begun work on formalizing some of these ideas [10]. To build systems that can mimic some of the reasoning abilities of humans, we shall also need naive psychology, naive biology, and other bodies of knowledge that all humans need and use.

It is in the area of commonsense knowledge that AI has a close contact with philosophy. McCarthy, in particular, has been concerned about studying (from an AI viewpoint) some long-standing philosophical problems concerning causality, counterfactuals, knowledge, and belief [11]. (Many of these topics might fall under the heading of naive psychology.) A key problem is to illuminate what it means for humans (or other computers) to know or believe things, to have goals, wants, plans, or fears. In the philosophical literature, knowledge, belief, want, and the like are called *propositional attitudes*. A propositional attitude is a relation between an agent and a proposition. For example, to say that agent A1 *believes* that neutrinos have mass is to state a relationship (or attitude) of *belief* between agent A1 and proposition that neutrinos have mass. Representing knowledge about propositional attitudes and methods for reasoning about them are currently important research topics in AI [12].

AI's involvement with commonsense knowledge is

a bit different from its involvement with more technical disciplines. In the latter, there are specialists whose job it is to expand knowledge. But there is no recognized discipline whose specialty is commonsense knowledge, so the research task of making it explicit has fallen, by default, to AI researchers.

AI is also concerned with knowledge about how knowledge itself should be structured and about how to use it most efficiently. For this *metaknowledge*, just as for commonsense knowledge, the AI researchers themselves are responsible for content as well as form. Even though AI can be considered to be a part of computer science, it is not a very well-behaved part; it can stretch beyond its boundaries to make statements about other subjects, about computer science, and, particularly, about itself. Just as we can have books written in English about English, just as we can have metamathematics, so we can also have knowledge about artificial intelligence expressed in AI knowledge representation languages.

An important idea in the efficient use of knowledge involves "procedural" representations. AI researchers discovered that there are circumstances in which it is more efficient to represent knowledge in a computer program or procedure than it is to represent it declaratively in a propositional formalism. On the other hand, it is harder to reason about the consequences of procedurally represented knowledge than about the consequences of propositionally represented knowledge. Weyhrauch's research in FOL [13] has produced a nice synthesis of these two approaches in which procedures can be used as the referents of expressions in a logical language. Such procedure constitute a portion of a *partial model* for the language. One then can choose between reasoning "syntactically" by manipulating expressions in the language or "semantically" by executing procedures (if available). Such a design displays a clean interface between what might be called the AI components of a system (those handling logical reasoning) and the more conventional or peripheral parts (those doing arithmetic, for example).

Of all the subject outside computer science, psychology has a particularly close relationship with AI. Psychology claims as its subject matter many of the natural phenomena (behavior, cognition, learning) that AI is attempting to understand and replicate in computer systems. One might expect, therefore, that AI has much to learn from psychology.

The relationship between psychology and AI is analogous to that between certain subdivisions of physiology (those dealing with spinal reflex arcs, nerve transmission, heart beat oscillators, etc.) and electrical engineering. Although several people in a field called *bionics* thought engineers could learn much from biological systems, it turned out that knowledge seemed to flow more from engineering to biology than in the other direction. Physiologists now understand some of

their subject matter in terms of constructs invented by engineers -- constructs such as feedback loops, stability, flip-flop circuits, and so on. I suspect that AI has similarly informed and will continue to inform psychology. Before scientists can make sense of natural phenomena, they need appropriate concepts and vocabulary. The concepts invented by AI researchers in the process of building intelligent machines will allow psychologists to construct more powerful models with which to explain human or animal intelligence.

It oversimplified, of course, to concentrate on the flow of information from AI to psychology. There have probably been many instances in which AI research has been illuminated by the work of psychologists. At the very least, the phenomena studied by psychologists provide AI research with some of its goals. But, on the whole I agree with Newell when he says: "... AI (and computer science) can live and prosper without psychology, but psychology cannot prosper without AI" [14].

Some AI researchers have suggested that AI's proper horizon includes *all* intelligent behavior, whether performed by animals or by machine. In such a view, psychology becomes a branch of AI, or at least a branch of some expanded subject, whatever it is called. (Some have suggested the name *cognitive science*.) I doubt, however, that a field that embraces the study of intelligent artifacts and intelligent animals could maintain the coherence needed to keep it together. A similar attempt at combining the study of natural and artificial mechanisms of control and communication -- *cybernetics* -- seems not to have succeeded.

IV. WHAT MIGHT BE LEFT OUT OF AI

So far we have asserted that AI concentrates on languages for representing certain kinds of knowledge, as well as on the mechanisms for processing those representations. We have also given some examples of the kinds of knowledge for which AI seeks representations and processing methods. But nothing we have said necessarily excludes conventional mathematical notations and techniques. After all, we might include logic as part of mathematics (which is sometimes done), and be left with just a single field that deals with all kinds of computer-interpretable knowledge representations (numerical, algebraic, logical) as well as with the techniques for manipulating them. Such a solution is perfectly conceivable, of course, but it still seems preferable to divide such a large field into components -- one of which deals with the more conventional mathematical constructs and methods, while the other is concerned essentially with logical apparatus for representing and manipulating that knowledge best expressed by natural-language sentences. AI concedes the

first of these components to the mathematicians, the applied mathematicians, the numerical analysts, and the algebraic computer language designers, i.e., to those who, after all, "got there" first.

The boundary between AI and certain other parts of computer science -- such as operating systems, compilers, parsers, and database systems -- seems less well-defined. One could argue that the representations and methods used in building compilers, for example, are not standard mathematical ones and are not that much different from some of the representations and methods used in AI. Nevertheless, on historical and practical grounds AI is well advised to leave the main custodianship of these other parts of computer science to those who are apparently managing very well without any direct help or interference from AI.

I have already mentioned that AI probably ought not to attempt to include the kinds of processing that occur at the sensory and effector peripheries of the nervous systems of animals. In my opinion, there has been some very good work indeed on these subjects, but the representations and methods used seem to be standard mathematical ones. We should not necessarily expect that all the different kinds of neural processing performed by animals can be neatly explained by the same theoretical constructs. Let us note, however, that excluding peripheral processing does not imply that the study of *higher*-level perceptual and motor reasoning is not properly a part of AI. The DARPA speech-understanding systems, for example, integrated low-level acoustic processing with representations and processes that I would definitely want to include in the AI repertoire [15].

V. APPLICATIONS OF AI

There are two kinds of applications of AI. In one, AI formalisms are used merely to *represent* knowledge about a certain subject, say, economic geology. Exploiting the extra power that AI formalisms bring to this task, experts in subjects like economic geology find that they can communicate ideas to their colleagues and students that had been difficult to state clearly and precisely in English. The language and terminology of several fields have already been enriched by AI formalisms. Knuth supports a similar claim when he observed "I believe the knowledge gained while building AI programs is more important than the use of the programs, but I realize that most people won't see this" [7]. Thus, progress in AI leads to progress in other subjects, and sometimes it is difficult to separate AI research from research in economic geology, for example. Usually, however, the AI researcher is interested in general methods that will be useful in many domains, while the economic geologist is interested primarily in his own field.

I anticipated that the language provided by AI will benefit a large number of disciplines. The knowledge obtained during a lifetime of experience by skilled medical practitioners will be available for use by others because it will be written down with a precision that English simply cannot capture. Complex legal arguments, which torture English beyond recognition, will be expressed in new formalisms more suited to the task. Although this prediction may sound extreme, I believe that AI formalisms (based primarily on logic and its extensions) will augment their more conventional mathematical counterparts to supplant English and other natural languages as the best medium for representing scientific, commercial, legal, and much commonsense knowledge.

The other major application of AI techniques, of course, is in the construction of systems that have access to knowledge of the kind we have just been discussing, and that perform in a manner similar to a skilled human who has such access. This application is what usually comes to mind when we think of AI -- namely, surrogates for humans in various intellectual or perceptual capacities. I agree with Knuth that, as surrogates, these systems have not yet had a significant effect. They are often brittle and do not yet cover a sufficiently wide range of situations to be truly useful. Systems of this kind have primarily been experiments conducted by AI researchers to test various representational and processing strategies. But we observe a steady increase from year to year in the power and range of performance of even the "breadboard" systems; soon some of these systems will have practical applications.

VI. IMPLICATIONS FOR AI TEACHING AND RESEARCH

The main reasons for attempting to gain a perspective regarding AI and how it relates to other disciplines is so that its subject matter can better be organized and taught and so that AI research can be pursued more productively. Let us consider educational strategies first. If propositional formalisms (especially logic) and their use in knowledge representation and manipulation are as important as I believe they are, then we ought to stress these topics in training AI researchers. Progress in AI research is slowed by the fact that many AI researchers do not know much about logic, even argue against it, and at the same time propose *ad hoc* substitute languages whose inadequacies, compared with logic, should have been obvious.

Other important components of an AI "curriculum" are topics found in AI textbooks: heuristic search and the important role it plays in the efficient manipulation of knowledge structures; deduction and planning processes; efficient indexing methods for storing and retrieving expressions; AI programming techniques using LISP.

Some acquaintance with suitable AI programs, presented as case studies, is also important. Familiarity with the latest research results in areas related to commonsense knowledge and reasoning would then round out the core of the program.

Let us turn next to AI research. An important concern for AI as a field is how basic AI research can keep in sufficiently close contact with its wide spectrum of applications. If basic research is to be productive and relevant, it must be continually stimulated by such contact. Similarly, work on applications needs to be informed by the results of basic research. It is very important, therefore, to maintain close, reciprocal contact between basic research and the various applications of AI.

One model that has proved useful for relating AI basic research and applications might be called the "onion model". At the core of the onion is basic AI research on such topics as knowledge representation languages, commonsense reasoning, deduction, planning, and heuristic search. One layer out in the onion is a shell that consists of the major research subdivisions in AI: natural-language processing, vision, expert systems, and problem-solving. These subdivisions, and others like them, are often used to divide AI research laboratories into subgroups and AI conferences into sessions.

At the next outer layer are what I call the "first-level applications" of AI. These applications of AI ideas are implemented by AI researchers for the purpose of advancing basic AI knowledge about the core topics or about AI subdivisions -- not necessarily to achieve anything useful as an application per se. Examples of such first-level applications are MYCIN [16] (a medical diagnosis program), PROSPECTOR [5] (a geological consulting program), and the DARPA speech-understanding systems [15]. These are systems that can be constructed within an AI laboratory with the help of consultants who have an expert's knowledge about the domain of the applications. Obviously, an AI laboratory must choose its first-level applications prudently. It cannot possibly work in all applications areas, yet it needs to work in enough different ones to ensure the generality of the methods being developed.

Continuing outward to subsequent layers, we first encounter applications of AI ideas that are done with the intention of achieving useful results in specific domains. These are what I call the "second-level applications". Here we find the development of robots that might be prototypes for those that will actually perform in factories, of medical systems to diagnose diseases in an actual clinical setting, of program verification systems that have demonstrable utility on large-scale computer programs, and the literally hundreds of other useful systems that might embody AI ideas. Each of these applications requires a substantial contribution from experts in the area of application. Because of this requirement, these applications are seldom carried to their successful

conclusions in AI laboratories -- although, to provide additional contact with the real world, a large AI laboratory may want to be involved in one or two of these areas.

Thus, a typical AI laboratory pursues work near the center of the onion. Depending on the size of the laboratory, a certain number of first-level applications and, perhaps, a second-level application or two may be included. Additional consulting on other second-level applications contributes to technology transfer and keeps the research community informed about real problems.

Just as applied mathematicians sometimes change hats to become electrical engineers, physicists, or other specialists, so do AI researchers sometimes become absorbed by the subject matter of an application. Without a clear idea of just what constitutes AI research, its goals, subject matter, and techniques, it is easy even for AI researchers not to be aware of having left AI. Some of these researchers are doing very good work in such specialities as VLSI design, computer program verification, or chemistry. I do not regard this work as AI research, however, unless it contributes generally to AI methodology. Although it is undoubtedly salutary for some AI people to devote themselves to particular applications, it is important that the core of AI retain its sense of cohesion and original purpose.

VII. CONCLUSIONS

AI, like certain other subdivisions of computer science, is concerned with representational formalisms, techniques for manipulating them, and implementations of these formalisms and techniques as computer programs. In my opinion, AI's special niche involves propositional representations. (Among such we include those formalisms useful for expressing and manipulating knowledge about knowledge representation formalisms themselves -- even the mathematical ones. Logic has traditionally played a key role in such metarepresentations.) The more conventional mathematical formalisms and manipulations are best left to applied mathematics. Adopting this point of view implies that certain topics, traditionally thought to be a part of AI, ought to be conceded to other disciplines.

AI is like mathematics in the sense that each can be

applied in a wide range of other subjects. There should usually be little ambiguity in deciding whether a piece of work ought to be regarded as an achievement in AI (or applied mathematics) -- or in a field that uses AI or mathematics as a tool. For example, when a physicist uses the diffusion equation to express the properties of heat conduction, he sees himself as contributing to thermodynamics, not to applied mathematics. Similarly when someone builds a useful organic-chemistry synthesis system employing established AI ideas, I would regard it primarily as progress in, say, computational chemistry, not in AI. (If the work also resulted in new advances (of general utility) in AI representational systems or in their manipulation, it might in that case also properly be regarded as a contribution to AI.) By this criterion, many AI researchers are not really doing AI research, but are doing work in chemistry, geology, VLSI design or what-have-you, using AI methods. AI, as a field, should be less concerned about whether the development of mathematical theorem provers, formula manipulators, program verifiers and synthesizers, robot control systems, and other applications is thought to *be* AI. It should be enough that identifiable AI methods are used. As AI ideas come to be employed routinely in many disciplines, we shall begin to see AI journals and conferences concentrate less on straightforward applications, no matter how successful, and more on innovative developments in the general methodology of AI (perhaps illustrated by exemplary applications.)

In this context it is easier for us to tolerate the annoying slogan, "if it's successful, it isn't AI". In a similar sense, if it's successful it isn't applied mathematics -- it's physics, chemistry, or some other subject. Yet applied mathematics prospers and can point to its own successes. So can artificial intelligence.

ACKNOWLEDGMENTS

I want to thank Barbara Grosz, John McCarthy, Bob Moore, Stan Rosenschein and Dave Wilkins for their helpful comments. This paper is based on a talk given at the Computer Science Department of Carnegie-Mellon University on April 22, 1981 in its Distinguished Lecturer Series. Discussions following that talk, as well as similar talks at Stanford and MIT, influenced the final form of the paper ■

REFERENCES

- [1] Marr, D., "Early Processing of Visual Information", *Phil. Trans. Royal Society (Series B)*, Vol. 275, pp. 483-524 (1976).
- [2] Barrow, H.G., and J.M. Tenenbaum, "Recovering Intrinsic Scene Characteristics from Images," in *Computer Vision Systems*, A. Hansen and E. Riseman, (eds.), pp. 3-26 (Academic Press, New York, New York, 1978).
- [3] Zue, V.W., and R.M. Schwartz, "Acoustic Processing and Phonetic Analysis," in *Trends in Speech Recognition*, W.A. Lea (ed.), pp. 101-124 (Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1980).
- [4] Weiss, S.M. et al., "A Model-based Method for Computer-aided Medical Decision-making," *Artificial Intelligence*, Vol. 11, No. 1 and 2, pp. 145-172 (1978).
- [5] Duda, R., J. Gaschnig and P. Hart, "Model Design in the Prospector Consultant System for Mineral Exploration," in *Expert Systems in the Microelectronic Age*, D Michie (ed.), pp. 153-167, (Edinburgh University Press, Edinburgh, Scotland, 1979).
- [6] Buchanan, B., and E. Feigenbaum, "Dendral and Meta-Dendral: Their Applications Dimension," *Artificial Intelligence*, Vol. 11, No. 1 and 2, pp. 5-24, 1978).
- [7] Private communication, 1981.
- [8] Rich, C., *A Representation for Programming Knowledge and Applications to Recognition, Generation, and Cataloging of Programs*, Ph.D. Dissertation, Massachusetts Institute of Technology, Cambridge, Massachusetts (1980).
- [9] Hayes, P.J., "The Naive Physics Manifesto," in *Expert Systems in The Microelectronic Age*, D. Michie (ed.), pp. 242-270, Edinburgh University press, Edinburgh, Scotland (1979).
- [10] McDermott, D., "A Temporal Logic for Reasoning About Processes and Plans," Department of Computer Science Research Report No. 196, Yale University, New Haven, Connecticut (March, 1981).
- [11] McCarthy, J., "Epistemological Problems of Artificial Intelligence," in *Proc. 5th Intl. Joint Conf. on Artificial Intelligence*, pp. 1038-1044 (1977).
- [12] Moore, R.C., *Reasoning About Knowledge and Action*, Artificial Intelligence Center Technical Note 191, SRI International, Menlo Park, California (October 1980).
- [13] Weyhrauch, R., "Prolegomena to a Theory of Mechanized Formal Reasoning," *Artificial Intelligence*, Vol. 13, pp. 133-170 (1980).
- [14] Newell, A., "A Textbook That Points the Way," *Contemporary Psychology*, Vol. 26, No. 1, pp. 50-51 (1981).
- [15] Medress, M.F., et al., "Speech Understanding Systems: Report of a Steering Committee," *Artificial Intelligence*, Vol. 9, No. 3, pp. 307-316 (1977).
- [16] Shortliffe, E., "Consultation Systems for Physicians," in *Proc. Canadian Soc. for Computational Studies of Intelligence (CSCSI)*, University of Victoria, Victoria, British Columbia (1980). ■

ANNOUNCEMENT

RECEIVING THE AI MAGAZINE

Recently, a mailing was made to the entire AAAI membership roster. If you are a member of AAAI and did not receive our recent letter, or if you are not receiving the AI Magazine please notify us. If you received our mailing please take the time to fill out the form we enclosed. A question on it pertains to receipt of the AI Magazine.

To date, there have been three prior issues of the AI Magazine:

Vol. 1, No. 1, Spring, 1980
Vol. 2, No. 1, Winter, 1981
Vol. 2, No. 2, Summer, 1981

If you wish to order back issues of the magazine. Notify.

AAAI
445 Burgess Drive
Menlo Park, CA 94025
(415) 328-3123
ArpaNet Address: AAAI-OFFICE@SUMEX-AIM