Artificial Intelligence: A Hermeneutic Defense

Thomas F. Gordon *

German National Research Center for Computer Science (GMD)
Institute for Applied Information Technology
Sankt Augustin, Germany

January 2, 1992

1 Introduction

The field of Artificial Intelligence (AI), from its very beginnings in the 1950s, has been criticized for its name, as well as its ambition. Most of the debate concerns the possibility of artificial intelligence, and presumes there is indeed some *thing* which *is* intelligence; the only question has been whether or not artificial systems can be built which exhibit, or have, this thing. That is, the debate has remained for the most part within the *rationalistic* tradition. In this section, I would like to explore two alternative approaches to this issue. The first considers the consequences of viewing artificial intelligence as another *metaphor* for computing. That is, perhaps certain kinds of computing systems can be usefully viewed as being *like* intelligent beings in some significant way. An immediate consequence of this view, of course, would be a lowering of the aims and ambitions of the field. The claim that a system displays something like intelligence is surely much weaker than the claim that it *is* intelligent.

However, I will reject viewing AI as a metaphor for another reason. The notion of metaphor presupposes that words do indeed have a core of certain meaning, their literal meaning. The metaphorical meanings deviate from the literal meaning in various ways, by removing some essential element of the definition of the term. By arguing that AI uses the term "intelligence" metaphorically, one would imply there is another, literal meaning.

Instead, I propose taking the lessons of *hermeneutics* seriously, by accepting that terms have no static, context-independent, literal meaning. No one use of the word "intelligence" can make an exclusive claim to legitimacy. The focus of the debate about AI should not be its *possibility*, but rather the *suitability* of using the expression "artificial intelligence" in particular social or institutional situations for some particular purpose. Thus, whether or not it is helpful to speak

*Many thanks to Joachim Hertzberg, Christoph Lischka, Thomas Christaller and the editors for their helpful comments and suggestions.

of some computer system as being intelligent cannot be decided conclusively in the abstract. The risks and opportunities presented by the contingencies of each *case* need to be considered.

This argument turns Winograd and Flores' use of hermeneutics to criticize Artificial Intelligence on its head. Their argument is unpersuasive, as it attempts to restrict and delineate the meaning of intelligence in a rationalistic way, while at the same time denying the validity of this rationalistic tradition. If we accept the lessons of hermeneutics, then we must also drop a naive correspondence theory of meaning. If intelligence does not denote any particular class of objects, then debates about whether or not certain kinds of machines can be members of such a class become vacuous. However, by abandoning this kind of rationalism, we are left with the task of finding alternative ways to critically examine and evaluate AI systems, while avoiding rationalistic debates about what is or is not intelligence. I discuss three possibilities, based on economics, natural science and law.

Insights from law and jurisprudence will play an important role several times here. For some time now, I have been active in the interdisciplinary field of artificial intelligence and law. In jurisprudence, there is a history of reflection about the nature of reasoning which is comparable in depth and richness to the philosophy of science and mathematics. Within jurisprudence, the limits of a rationalistic perspective have long been appreciated, although the discussion about the consequences of these limits continues. For example, it is recognized that the meaning of concepts is "open-textured" and evolves over time, that normative considerations are of central importance when deciding whether to subsume some event under a general term, and that it is futile to try to construct a "heaven of concepts" for classifying all future events.

Before continuing with a more detailed discussion about these three ways of viewing AI, it may be useful to recall some of the various points of view about the nature of logic. Although the significance of logic for AI is often debated, there does not seem to be much controversy any more about what logic is. This was not always true. I am not well enough acquainted with the history of logic to know why the debate has quieted down; but I doubt it was because the various issues were settled. Has the arena for debating the various issues involved shifted from logic to AI? Stephen Toulmin, in his *The Uses of Argument* [8] listed these positions respecting the subject matter and purpose of logic:

Logic as psychology. Logic is concerned with the "laws of thought", not with pathological, defective thinking processes, but with "proper, rational, normal" thinking, "the working of the intellect of health".

Logic as sociology. Rather than the individual human mind, it is the "habits of inference" which have been "developed in the course of social evolution" that are of interest to logicians.

Logic as technology. Rather than an empirical science about how (healthy) persons actually think, logic is viewed as a craft, a collection of techniques for effective thinking.

Logic as mathematics. Logic is neither science nor art (craft), but a special field of pure mathematics concerned with the properties of an abstract set of objects such as "logical relations".

Logic as "generalized jurisprudence". This is the view developed by Toulmin in *The Uses of Argument*. Logic is concerned with the "soundness of claims", the procedures by which claims are "put forward, disputed and determined". Legal disputes are viewed as just a special case of rational dispute for which the procedures have "hardened into institutions".

One could just substitute "AI" for "logic" in the above list to get a list of some arguable positions about the proper subject matter of AI. The formulation of some of the items would have to be modified somewhat. The technology argument, e.g., would have to distinguish between making tools for assisting effective thinking and machines which themselves think. To my knowledge not all of these positions have been taken; the technology and psychology positions have received the greatest amount of support. Of course, as a lawyer, I would be willing to argue that AI, too, can profitably be viewed as "generalized jurisprudence".

2 The Rationalist Debate About AI

The usual debate about the possibility of artificial intelligence focuses on three issues: 1) What *is* intelligence; 2) Can, as a matter of principle, intelligent machines be constructed; and 3) If the second question is answered affirmatively, how? Taken for granted in these discussions is the adequacy of the scientific method for addressing these questions. Notice also that here the ambition of AI to construct intelligent systems is understood literally. In the next section we will explore the view that AI is not really about intelligence at all.

With respect to the first question, neither the existence of a thing called "intelligence" nor the possibility of delineating the class of intelligent things in terms of necessary and sufficient properties is called into question. The debate centers around which properties are necessary and which are typically associated with intelligence but not strictly required before one is willing to attribute intelligence to some object. The point of John Searle's "Chinese Room" *Gedankenexperiment*, for example, is that *understanding* is, in his view, an essential part of intelligence [7]. Searle claims he would be unmoved by an AI system which could automatically translate one natural language into another. Even if a computer could be programmed to adequately and convincingly perform this task, he would be unwilling the attribute intelligence to the machine as he is convinced his thought experiment shows that mere performance does not imply understanding.

Also along this line are the arguments claiming an intelligent system must be embedded in the "real world". The claim is that a certain kind of robot might be said to be intelligent, but not for example a chess playing program. Chess programs do not have sensors and manipulators. They transform strings of input symbols into strings of output symbols and the locus of intelligence, according to this line of argument, remains with those persons who interpret these strings in order to make their next move in some game of chess. Intelligence here is

viewed as an attribute of autonomous systems struggling to be successful in some environment.

For some, intelligence is a defining characteristic of certain "higher" forms of biological life, humans in particular. From this perspective, no machine can be intelligent simply because a machine is not an animal. That is, the defining characteristics of intelligence are so intimately connected with being an animal that, as a matter of definition, no machine can sensibly be said to be intelligent.

Cognitive science, on the other hand, is based on the premise that it can make sense to talk about intelligence abstracted from biology. (This is what distinguishes cognitive science from cognitive psychology.) Animal and human intelligence are viewed as special cases. This perspective opens the door to defining classes of intelligence where one or more of the defining characteristics of human intelligence are missing. One could postulate a form of intelligence, e.g., where understanding, in Searle's sense, is not required. Perhaps chess machines could be said to display one of these other forms of intelligence. The problem with this strategy is to find a taxonomy of intelligence which is not arbitrary and construed. Why call something "intelligence" if it is not sufficiently related to the common sense meaning of the term?

Whether AI is achievable obviously depends on the particular view of intelligence adopted. In *Understanding Computers and Cognition*, for example, Winograd and Flores deny the possibility of AI by restricting their view of intelligence to human intelligence [9]. The AI projects they criticize do indeed aim to achieve human levels of performance in such domains as natural language understanding.

Following the approach articulated by Rodney Brooks [1], there is a project underway here at the GMD, lead by Christoph Lischka, to construct a small, autonomous, mobile robot displaying the intelligence of a certain kind of lizard. Although the goal is not human levels of intelligence, this project is ambitious enough. Perhaps the robot should be able to catch small insects, for example. Whatever the merits of Winograd and Flores' arguments concerning AI and human intelligence, it remains an open question whether these more modest goals are realizable.

The third issue, how to achieve AI, is too often confused with the second issue, the possibility of AI in principle. Newell and Simon's physical symbol system hypothesis (PSSH) has played an important theoretical role in the history of AI [5]. Its importance is such that the PSSH is often confused with AI itself. Elaine Rich in her textbook on AI, e.g., states that the PSSH lies "at the heart of research in artificial intelligence" [6, p. 3]. Indeed the central role of the PSSH is so great that those committed to other approaches to constructing intelligent systems, such as some connectionists, claim they are no longer doing AI research![1]

[1] According to Rich, Newell and Simon define a physical symbol system as follows: "A physical symbol system consists of a set of entities, called symbols, which are physical patterns that can occur as components of another type of entity called an expression (or symbol structure). ... At any instant of time the system will contain a collection of these symbol structures. Besides these structures, the system also contains a collection of processes that operate on expressions in order to produce other expressions ... A physical symbol system is a machine that produces through time an evolving collection of symbol structures. Such a system exists in a world wider than these

Why should AI be committed to any particular hypothesis concerning the features sufficient or necessary of an intelligent machine? Indeed, why should AI be committed to the digital computer? There should be room within AI for alternative approaches and hypotheses. Physics doesn't stop being physics when a new theory of the universe is proposed. If one views intelligent machines as being the subject matter of AI, then there seems no pressing need to restrict the field to a particular type of machine.

3 AI as Metaphor

It may appear that one way to avoid some of the difficulties of the hard line view that AI is about building intelligent machines is to argue that AI's use of the term "intelligence" is metaphorical: AI systems are not *really* intelligent, they just have some features in common with intelligence. This would allow us to preserve the conventional meaning of intelligence without necessarily limiting AI's ambition regarding the levels of performance to be achieved. The adequacy of the Turing test for testing intelligence would be a non-issue, as real intelligence would not be claimed.

Unfortunately, this argument is not without problems. First, it does not completely avoid the problem of defining intelligence. As AI aspires to be a science, its subject matter needs to be delineated rather more precisely than some unspecified relation to intelligence. The task here is however simpler, as we can be satisfied with a set of features characteristic of intelligence, without being concerned with whether the set is exhaustive or includes all elements necessary for real intelligence. To justify a metaphorical use of intelligence, its literal meaning must be understood to some extent; but it is not necessary to precisely define intelligence in terms of necessary and sufficient conditions. If the claim is made that an AI system *is* intelligent, then the door is open for arguing that some necessary feature of intelligence is missing. The claim that an AI system behaves *as if it were intelligent* is much weaker. The absence of a necessary feature would not rebut the claim.

Secondly, this approach to defining AI is not in the end significantly different than the approach taken in cognitive science. The history of science shows that it is not unusual to apply an everyday word metaphorically to describe a new technical concept. Through use in the scientific discipline the term acquires a new technical meaning. Examples include the terms "field" and "force" in physics.

The use of metaphor has its justification. Languages such as English do not encourage the creation of new terms, and metaphor allows a language to be extended with new senses and shades of meaning for its existing vocabulary. It also allows a new concept to quickly acquire meaning, by inheritance from some existing sense of the term. However, the use of metaphor brings with it the risk of misunderstanding. The complete meaning of the prior sense is not

symbolic expressions themselves." The Physical Symbol System Hypothesis (PSSH) is: "A physical symbol system has the necessary and sufficient means for general intelligent action." [6, pp. 3-4]

carried over into the new context, and it may not be apparent just what the metaphorical use of the term is intended to mean.

The main problem with viewing the "I" in AI metaphorically, is that intelligence is an abstract, *open-textured* concept [3]. If intelligence is to be understood metaphorically within AI then arguably the AI community should make an effort to distinguish between its metaphorical, technical use of the term and its everyday, common sense meaning. However, this is easier said than done. Intelligence has no well-understood literal meaning. Our very understanding of intelligence continues to develop along with our research in AI.

William C. Hill has recently argued that, not only is AI a metaphor, it is a poorly chosen metaphor [4]. Hill claims that most AI research has not been about intelligence at all, not even metaphorically, but about constructing new "computation-based representational media", i.e. new forms of communication. He first claims that AI is a phrase such as "horseless carriage, wireless telegraph, iron horse, glass teletype or artificial writing", a phrase which "describes a new technology wrongly in terms of an old familiar one". But unlike these other terms, "artifical intelligence" distracts attention from the new technology's principal use, the computer's potential for improving communication. Actually and "even worse", Hill goes on to say, AI is not like "wireless telegraph" as it describes new technology not in terms of other familiar technology but in terms of mental phenomena, causing irrelevant arguments about the nature of mind and intelligence.

Hill's arguments have a great deal of merit, at least for those of us who, upon reflection, have become involved in AI precisely because they are interested in new forms of effectively representing and distributing knowledge and ideas. However, each of us needs to decide for ourselves where our interests lie, and a great number of AI scientists are indeed principally interested in pursuing the goal of creating intelligent machines.

4 A Hermeneutic Interpretation of AI

The arguments outlined above about the nature of AI do not give up what Winograd and Flores call the "rationalistic tradition" of Western science. They view this tradition as being based on mistaken premises regarding the nature of understanding and knowledge, drawing principally for support from Heidegger's hermeneutic theory of understanding, Maturana's theory of perception and cognition and Searle's theory of speech acts. Each of these thinkers arrives at much the same epistemological stance, although they start from very different intellectual traditions. Heidegger is an existentialist philosopher, Maturana a biologist and Searle a linguist.

It would be too much to try to replicate Winograd and Flores' arguments here. It is also unnecessary for our present purposes. I do not indeed to challenge or support their position, but to examine some of the consequences for AI of accepting their principal conclusions regarding the nature of knowledge and understanding. So, the next few paragraphs will be limited to a summary of this point of view. Although the arguments are difficult and foreign, at first, the main

insights are not so difficult to grasp when stated informally in everyday terms. Let us start with language.

According to the rationalistic tradition, words have literal meaning. That is, words are thought to correspond to objects in the world. They *denote* things in a context independent way. Searle's speech act theory challenges this notion by arguing that the meaning of a sentence is always dependent on some particular conversational context or situation. The speakers, their goals and intentions need to be considered when trying to get at the meaning of some "illocutionary act".

I am not sure whether Searle dealt with this aspect of meaning in his work, but the context dependence of meaning implies that words acquire new meaning through use. Words have an "open texture". In the philosophy of law, H.L.A. Hart, especially, stressed this quality of legal terms [3]. His position on this subject is a moderate one in that he asserts that terms do have a core of certain meaning. There are cases where a term is clearly applicable. Open-texture is limited to the boundary.

This brief mention of legal reasoning provides a nice opportunity to shift our attention to Heidegger, who adopted and generalized the term "hermeneutics" from its prior context. Prior to Heidegger, hermeneutics had been the art of interpreting legal and religious texts. It had been an approach to trying to understand the intended meanings of a text long after it has been written, by persons divorced in time, space, language and culture from the text's original context.

Heidegger generalized hermeneutics to the problem of an individual trying to understand his world. Not only is the connection between a term in some language to objects in the world tenuous, so too is the connection between an individual's conceptualization of the world and the world itself. The classification of objects, indeed the identification of objects, does not preexist, but occurs during the process of interaction with the environment. The particular division of the world into objects, properties and classes arrived at depends on a person's unique history, goals and perspective. Contrast this view with the rationalist tradition, where there is thought to be an "objective" view of reality, where the goal of science is viewed as arriving at, by application of the scientific method, the one true theory of the world.

As I understand Winograd and Flores, the relevance of Maturana's work here is that he explains in biological terms the dependence of perception on the structure of the perceiving organism. Compare this with behaviorism, where it is assumed that stimuli can be identified, measured and categorized independently of the structure of the particular organism. Winograd and Flores use Maturana's conclusion that there is no objective perception to support their critique of rationalism.

Let me use the word "hermeneutics" as the general term for this alternative orientation towards understanding and knowledge, without necessarily restricting ourselves to Heidegger's particular interpretation. One difficulty with the hermeneutic viewpoint is that any attempt to convey it must use language, and the conventions of our language are so deeply steeped in the rationalistic tradition that the hermeneutic viewpoint appears mysterious, mystic or even

self-contradictory. Any description of the hermeneutic perspective is couched in terms of a theory about perception, knowledge and understanding, which gives the impression that the terms of the theory denote objects in the world and that the theory is subject to verification or, if you prefer, falsification. One plays the game of natural science while refuting the rules of the game. The theory gives rise to the kind of tension experienced with the various instantiations of the Liar's paradox. To understand the hermeneutic perspective, one has to temporarily suspend disbelief as one does when reading a novel or viewing a movie.

In the last section we discussed the possibility of a metaphorical interpretation of AI. The usual interpretation of metaphor does not diverge from the epistemology of the rationalistic tradition. On the contrary, the whole notion of metaphor depends on a distinction between the literal context-free meaning of a term and a novel use which is in conflict in some way with this literal meaning.

Now we are at a point where we can discuss Winograd and Flores' use of the hermeneutic perspective to criticize AI. They argue that AI is deeply embedded in the rationalistic tradition, pointing especially to Newell and Simons' Physical Symbol System Hypothesis. Their interpretation of the PSSH, which I suppose is the usual interpretation, supposes that symbols "can be understood as referring to objects and properties of the world" [9, pg. 74].

There may be a number of ways to rebut Winograd and Flores' treatment of AI. One could argue for another interpretation of the PSSH, for example. Rather than supposing that symbols denote objects in the world, one could hypothesize that symbol processing of a certain kind is a sufficient and necessary condition for higher kinds of intelligent behavior, even though symbols do not denote objects in the world. That is, the PSSH need not imply a naive correspondence theory of meaning.

Let us focus on another kind of rebuttal here, however. This is my main point: If we accept the hermeneutic viewpoint, then we must also accept that the term intelligence does not denote any particular thing, but may vary in meaning depending on its use within some particular context. One can argue, as lawyers do, that a particular interpretation in some particular context (or case) would lead to certain desirable or undesirable consequences, or create certain risks or opportunities, but not generally, abstracted from some concrete context, that only this or that concept of intelligence is legitimate. That is, if we intend to adopt a hermeneutic perspective, we cannot at the same time use rationalistic arguments to deny the possibility of AI. Hermeneutics requires other criteria for evaluating scientific hypotheses, indeed it presupposes an alternative philosophy of science.

Of course I cannot pretend to develop an alternative, hermeneutic philosophy of science here. But it may be that such an alternative science would have some characteristics in common with jurisprudence, that the criteria and methods used to evaluate scientific theories would resemble the methods used to decide legal cases. This idea is explored a bit further in the next section.

5 Evaluating AI Sytems

Rather than trying to define intelligence in the abstract, and then arguing about whether or not artificial intelligence is possible or desirable, it may be more constructive to focus our attention on the problem of evaluating specific AI systems. An understanding of the limits and potential of AI can evolve through the practice of constructing and using specific systems. Concrete systems, used in specific situations, permit the interests of users and system designers to be taken into consideration.

Again, the legal analogy is useful here. Vague terms are often deliberately used in statutes as a way of deferring decision-making from the abstract setting of a parliament or congress to the courts, where the term can be fleshed out during the process of deciding concrete cases. A constitution may refer to "due process" or a statute "reasonable cause", without further defining these terms. This is as it should be. Legislatures lack the vision to foresee all the consequences of a law.

Similarly, the field of AI can be characterized very abstractly, as Elaine Rich, for example, does when she writes AI is "the study of how to make computers do things at which, at the moment, people are better". What is or is not an AI system can then be decided, if not definitely, on a case by case basis.

This approach to delineating the field is not merely an attempt to avoid the difficult issue of defining intelligence. Rather, taking the hermeneutic perspective seriously, it is a recognition that concepts like intelligence cannot be defined.

How then can an AI system be evaluated without returning to the futile problem of defining intelligence? Toulmin claims there can be no general domain-independent method for evaluating arguments. The same claim can be made for evaluating AI systems. Which methods are appropriate depends on the interests and goals of the system designers and users. Without striving for completeness, here are three approaches which come to mind:

The Marketplace Approach. For AI products, an economics perspective may be appropriate. There is no need to decide whether the system really displays intelligence; it is sufficient that some community continues to find the product useful. There is a lot of hype in the AI industry, as in all industries. We in AI should find this no more or less disturbing than in other fields.

The Natural Science Approach. When the purpose of an AI system is to test an hypothesis about some particular cognitive process, then the usual methods of natural science may be appropriate. I hesitate to call this a process of validation. Supporters of Popper's philosophy of science, at least, argue that scientific theories cannot be validated, but at most falsified.

It is not necessary to delve into another discussion about the nature of science here. The only point I would like to make is that the rationalistic tradition has, despite its limits, proven its value. The limits of the rationalistic perspective can be viewed as simplifying assumptions which are appropriate in certain contexts. Apparently, we can ignore the lessons of

hermeneutics for certain tasks, just as we can, to use a tired example, get along well with Newtonian physics in our daily lives without resorting to the complexities of the theories of relativity or quantum mechanics.

The Legal Approach. In the case of knowledge-based or expert systems, especially when such systems are used for making decisions in organizations where the rights and duties of persons in the organization may be affected, it may be appropriate to view the "knowledge base" of the system as a set of laws. The decisions made by such systems must be backed up by cogent argumentation. (This is usually called an "explanation facility".) These decisions should be subject to challenge in some kind of quasi-legal proceeding by the persons affected.

Notice that in such settings, the decisions made by expert systems must be justified by *normative* arguments. Neither logic nor appeals to "cognitive adequacy" are sufficient. Rather, substantial arguments having a legal quality are required. Conflicting interests must be balanced and the appropriateness of subsuming the concrete events of the case under the general terms used in the knowledge base must be addressed.

Formal verification methods can play a role in the above approaches, but cannot themselves offer a complete solution. Formal methods may be used to show, for example, that two forms of representation are equivalent, by constructing a sequence of transformations, known to preserve some relevant property, from one form into the other. Such methods may also be used to derive properties of a knowledge base or program. These formal methods alone are insufficient, as there is no formal way to show that the knowledge structures of an AI system are satisfied by the intended application domain. A stronger statement can be made here. No method, formal or not, can "verify" an AI system generally, abstracted from its application to concrete cases. (Indeed, this is true of all computer systems.) Again, this is an argument from jurisprudence; there is no general method for determining outside the context of particular cases whether the facts of some case are subsumed by the general terms used in the knowledge base. A literal interpretation of a knowledge base will result in unintended or undesirable decisions being made. Knowledge bases, like the law, need to be modified, reinterpreted, and extended as they are applied to particular problems.

The term "verification" is misleading, as its use suggests that there is some way to gain complete confidence in the "correctness" of some AI system. Except perhaps in highly artificial, construed domains, this will never be the case. In practice, there will be arguments for and against the suitability of the system for its intended task, and it will be matter of judgment requiring the exercise of discretion and interpretation to decide the matter. Thus I have chosen to speak of "evaluating" AI systems. Perhaps "judging" would have been still better.

6 Conclusion

AI as a field has always had its antagonists. In some countries, notably Great Britain and West Germany, the field was prevented from advancing as rapidly

as in the U.S. because of negative assessments of AI's legitimacy or potential, or because of conflict between AI and conventional computer science. After more than 30 years of development, however, I think we can say that AI is here to stay, despite ongoing discussion about the nature of intelligence. I am confident of this for a number of reasons. First of all, there is now a thriving AI industry. These commercial interests will not allow AI to die. Second, this industry is a testament to the fact that AI has indeed created useful technology. It is simply not true that AI has produced no or too few tangible results. (It is true that AI has not fulfilled the promises of some of its promoters, but these individuals should be held accountable for their predictions and claims, not the field as a whole.) Third, our understanding of intelligence is evolving at the same time as AI, indeed because of AI. AI is not bound to any particular hypothesis about the nature of intelligence, despite the historical significance of the physical symbol system hypothesis. Indeed, one could argue that critics of the PSSH are not actually critics of AI, but are themselves doing AI by pointing out the limits of a particular hypothesis, at least if they also go to the trouble of proposing a competing hypothesis. The Dreyfus brothers are a good example here [2]. They may not represent the AI mainstream, but they have played the role of ombudsman within the field. They belong to AI. Finally, one cannot use hermeneutics to attack AI. The hermeneutic view implies that there is no single correct definition of intelligence. The notion of intelligence evolving through the practice of AI is just as legitimate, in its special context, as any other.

References

[1] Rodney A. Brooks. Achieving Artificial Intelligence Through Building Robots. A.I. Memo 899, MIT, May 1986.

[2] H. Dreyfus. *What Computers Can't Do: A Critique of Artificial Reason*. Harper and Row, 1972.

[3] H. L. A. Hart. *The Concept of Law*. Oxford University Press, 1961.

[4] William C. Hill. The Mind at AI: Horseless Carriage to Clock. *AI Magazine*, Vol. 10(No. 2), 1989.

[5] A. Newell and H. A. Simon. Computer Science as Empirical Inquiry: Symbols and Search. *Communications of the ACM*, 19(3), March 1976.

[6] Elain Rich. *Artificial Intelligence*. McGraw-Hill, 1983.

[7] John R. Searle. Mind, Brains, and Programs. In Douglas R. Hofstadter and Daniel C. Dennet, editors, *The Mind's I*. Basic Books, 1981.

[8] Stephan Toulmin. *The Uses of Argument*. Cambridge University Press, 1958.

[9] Terry Winograd and Fernando Flores. *Understanding Computers and Cognition*. Ablex, Norwood, New Jersey, 1986.