Differing Methodological Perspectives in Artificial Intelligence Research

Rogers P. Hall and Dennis F. Kibler

Irvine Computational Intelligence Project, Department of Information and Computer Science, University of California, Irvine, California 92717

Abstract

A variety of proposals for preferred methodological approaches has been advanced in the recent artificial intelligence (AI) literature Rather than advocating a particular approach, this article attempts to explain the apparent confusion of efforts in the field in terms of differences among underlying methodological perspectives held by practicing researchers The article presents a review of such perspectives discussed in the existing literature and then considers a descriptive and relatively specific typology of these differing research perspectives. It is argued that researchers should make their methodological orientations explicit when communicating research results, to increase both the quality of research reports and their comprehensibility for other participants in the field. For a reader of the AI literature, an understanding of the various methodological perspectives will be of immediate benefit, giving a framework for understanding and evaluating research reports In addition, explicit attention to methodological commitments might be a step towards providing a coherent intellectual structure that can be more easily assimilated by newcomers to the field.

More than a quarter century after its beginnings, AI has yet to produce a commonly accepted statement of purpose or description of conventional research practices. Studies are reported in a wide range of publications. While some focus on the field (*e.g.*, *Artificial Intelligence*), others are concerned with different research areas (*e.g.*, *Behavioral and Brain Sciences*). What results is a profusion of literature that is difficult to encompass for stu-

dents and practitioners alike. If the study of AI is to be considered, and conducted as a scientific endeavor rather than an amorphous enterprise whose subject matter is constantly shifting (or even disappearing as results are incorporated into other fields), one might profitably ask if distinct methodological perspectives can be identified by which to organize some of the current confusion of efforts. Perhaps, as others have pointed out, "there are undoubtedly some views of AI that are more fruitful than others ... We ought to be guided by the most productive paradigms" (Nilsson, 1982).

Concern for methodological issues in AI research is on the upswing (Ohlsson, 1983; Bundy, 1983a; Cercone and McCalla, 1983). However, this interest appears to be prescriptive focusing on what AI researchers *should be* rather than what they actually *are* doing. For example, Cercone and McCalla relegate the multitude of differing AI approaches to a rather constraining spatial metaphor, a "pie" composed of problem areas like vision, expert systems, or learning. They then specify "design objectives that any *ideal* AI system *should* meet" (p. 4, italics added). These objectives include development of a working system, external validation of system capabilities, and identification of generalized results. While these may well be desirable characteristics for some idealized view of AI research, they fall short of specifying the underlying assumptions or basic objectives of practicing researchers. This despite remarks that

"AI is a turbulent, exciting, audacious research area with a multitude of different approaches and influences which should continue to gain in credibility and importance in the years to come (p. 21)."

We argue that prescriptive methodological analysis of AI simply adds to the prevailing sense of confusion. We propose instead a descriptive analysis of what participants in the field actually do, as evident in exemplary published

Address all correspondence to first author

studies. We will present a variety of perspectives for viewing AI research, which have been previously reported in the literature, and we will attempt to condense those perspectives into a descriptive epistemological framework or typology. No attempt will be made to evaluate their relative efficacy of differing perspectives. However, we will argue that researchers in AI should make their methodological perspectives explicit when publishing research results so that both practitioners holding alternative perspectives and newcomers to the field can understand them better. Even without explicit methodological commitments on the part of writers, a reader sensitive to methodological issues should be better able to evaluate published research reports.

Are Research Perspectives Important in AI?

In a relatively new field of intellectual exploration, adherence to (or even identification of) methodological perspectives might be considered somewhat counterproductive. The flexibility with which investigators could approach a tremendous variety of potential research questions might be unnecessarily constrained if the means of investigation were rigorously defined. We believe, however, that identifiable methodological approaches have evolved and have come to play an increasingly important, though implicit, role in shaping public discussions of the "state of AI." Clearly, researchers are doing something, and this activity should be amenable to descriptive analysis.

As a preface to this analysis, some description of what organizing research perspectives provide for scientific activity seems in order. "Perspective" in this article refers to a disciplinary matrix or context within which researchers practice their trade. This concept of a matrix is similar to the most general notion of a "paradigm" as advanced by Kuhn (1970) and includes sets of shared values within a common conceptual vocabulary (Weimer, 1979). Although we will not address the developmental maturity of AI as a scientific endeavor, we will assume that multiple, competing paradigmatic orientations can exist simultaneously in a single field (Masterman, 1970). Hence proponents of differing perspectives often talk past each other when discussing aspects of their work in terms of those different perspectives.

From a functional point of view, a shared methodological perspective provides important constraints on exploratory activity for a particular community of researchers. First, the perspective guides selection of appropriate research problems or phenomena to be examined much as a map defines boundaries in a geographic territory being explored. Once a problem is selected for study, the shared perspective specifies how acceptable research should be conducted and indirectly provides a medium of communication (*i.e.*, journals and conferences) through which research findings can be examined by the community as a whole. This critical evaluation of findings proceeds in accordance with a shared set of criteria for gauging the quality of research efforts. Finally, the research perspective provides a framework within which new participants in the research community can assimilate the skills necessary for conducting and evaluating research.

Paradoxically, there are indications that the field of AI lacks the kind of direction a clear explication of such methodological constraints might provide There is considerable disagreement over what constitutes important research problems, as is evident, for example, in the divergence of views collected in a recent survey of research on knowledge representation (Brachman and Smith, 1980). The survey indicated no clear consensus on what "knowledge" was to be represented or what "representation" entailed. This sort of confusion is also evident in the selection of appropriate research methodologies, leading to arguments like the persistent methodological squabble between "scruffy" and "neat" views of AI research (Bundy, 1982).

Perhaps as a result of such problems, published studies turn up in a dizzying variety of journals, conference proceedings, and books, which are hard to follow even for the experienced AI researcher, much less neophytes. Assuming that readers can find published studies that relate to their interests, there seems to be no consensus as to what constitutes "good work." In the opinion of one leading AI researcher, this confusion over appropriate problems and methods has reached a point where most recent work submitted to conferences has been rejected as "junk" (Schank, 1983). Such confusion must also have consequences for the training of new AI researchers, forcing them into what amounts to isolated apprenticeships with curricula that have been described as "scandalous" (Nilsson, 1980a) in their lack of formal methodological training.

This serious state of disarray might suggest that organizing research perspectives currently do not exist in the field. We contend, however, that several perspectives do exist, that they are seldom explicitly stated, and that in some respects they are in opposition to each other. Without an explicit examination of these perspectives, critical evaluation of AI with respect to problem choice, methods, and quality of results suffers, and the study of AI is difficult to present as a coherent academic discipline.

Perspectives in Print

This section contains a roughly chronological selection of previously published methodological analyses of AI research. Each of these studies seeks to partition research in AI into several distinct categories, depending on the author's underlying assumptions. Obviously, this sort of partitioning can be achieved in a variety of ways:

- By existing or desired scientific qualities (*e g*, Ohlsson's call for principles, 1983).

- By position on central problematic issues (*e.g.*, Schank's endorsement of cognitive modeling, 1983)

- By task domain (*e.g*, Cercone and McCalla's AI pie, 1983)

- By examination of the intent of the researcher and the methods applied in pursuing those intentions.

Our review will not be exhaustive, as we have chosen to exclude analyses which are overtly prescriptive in nature (*e.g.*, Ohlsson, 1983 or Schank, 1983) or which subdivide AI research topically (*e.g.*, Cercone and McCalla, 1983). The former approach speaks very little to what practicing AI researchers actually do in their work. The latter blurs important methodological distinctions in trying to make content-level comparisons more accessible to the reader. Instead, we will focus on analyses that attempt to describe what researchers intend to do in their work and what methods they employ towards that end. *i.e.*, descriptions of existing methodological perspectives AI research. To the extent possible, problem choice, preferred methods, and criteria for successful research will be discussed for each methodological approach. However, as will soon become clear, existing descriptive analyses of methodological issues in AI often do not achieve a level of detail that allows sharp distinctions with respect to these issues.

Newell's Analysis

One of the earliest descriptive analyses of research strategies comes from Allen Newell (1973). Research orientations are broken into three classes, and exemplary studies are presented for each. First, Newell describes a class in which the *exploration of intelligent functions* provides the major research focus. In this approach, a problem task is chosen that is assumed to require intelligent (typically human) behavior, and a computational mechanism is proposed that is sufficient to support the accomplishment of this task. As an exemplar, Newell cites Green's (1969) work in automatic theorem proving.

The second class characterizes AI as a *science of weak methods* where "weak" is intended to designate general usefulness of a particular technique across a variety of problem domains, despite low information content with respect to any particular domain. Thus, the methodological approach is to describe a general technique or method, demonstrating its effectiveness across problem domains in the hope of eventually establishing a collection of general methods that are useful for constructing "intelligent" systems—much as numerical methods are applicable across a wide range of problems. An exemplar would be Newell's (1969) examination of ill-structured problems.

Newell's final class views AI as *theoretical psychology*, in the sense of viewing human cognition as the performance of an information-processing system. In this class, a computer model of some cognitive process is proposed and

then validated by comparison with features of the human behavior being modeled. An exemplar would be Newell and Simon's (1972) analysis of cryptarithmetic problem solving in which program control structure is compared with human strategies evident in problem-solving protocols.

Weizenbaum's Analysis

An incompletely specified but highly visible account of AI research practices can be found in Weizenbaum (1976). According to Weizenbaum, AI researchers proceed in one of three modes. First, in a *performance* mode, researchers are purely concerned with building practical software systems that satisfy a need for some artifact capable of impressive levels of performance. As an example, Weizenbaum makes a general reference to robotics research but cites no particular study. In contrast, AI researchers working in *theory* mode strive to uncover general principles of intelligent behavior without explicit regard to issues of implementation, much as turn of the century aerodynamicists studied principles of flight. No exemplars are given for this view of AI. Weizenbaum's third mode, *simulation*, involves the construction of computer models of human cognition which can be compared with actual human behavior as a means of validation. As an exemplar, Weizenbaum suggests Newell and Simon's (1963) work on GPS.

Feigenbaum's Analysis

In the context of describing an emerging discipline of "knowledge engineering," Feigenbaum (1977) argues for the continuing applicability of a view of AI research proposed much earlier (Feigenbaum and Feldman, 1963). In the first of two representative classes, researchers attempt to build useful intelligent systems and to develop a methodology that supports such construction. This view is particularily appealing to Feigenbaum, and his description of it is replete with terminology suggestive of an engineering discipline (*e.g.*, "workbench," "knowledge engineers," and "toolkit"). Artifacts as "intelligent agents" are characterized by their use of heuristic search guided by a considerable amount of domain specific knowledge. Numerous exemplars are given, with the DENDRAL project (Feigenbaum *et al.*, 1971) being distinguished by both its longevity as a research program and its popularity among researchers in chemistry. Feigenbaum claims the second class of AI research closely follows Newell's (1973) view of AI as theoretical psychology, but offers no exemplars, or even any discussion.

Lenat's Analysis

In describing his research on mathematical discovery, Lenat (1978) proclaims a single paradigm for AI research, in which intelligent behavior is viewed as the output of a symbol-processing system. Having selected some human

cognitive activity, the AI researcher proposes a theory of information-processing to support that activity, in the form of a computer program. The behavior of the running program is then examined to determine the locus of "intelligent" behavior with the hope of uncovering a unified theory of intelligence. As an exemplar Lenat discusses his own work (1977) on automatic theory formation in mathematics, stressing a view of intelligence as heuristic search. Unfortunately, while this description is compelling, alternate methodological viewpoints are not discussed, making his analysis descriptive in only a narrow sense.

Hayes's Analysis

Hayes (1978) contrasts AI research methodology with that of general systems theory, arguing that AI defers generality in favor of working programs. According to Hayes, *applied AI* focuses on creating useful artifacts in highly circumscribed task domains. As an exemplar, Hayes suggests the work of Waltz (1975) which supports natural language access to a large database of aircraft maintenance and flight information. Hayes describes a second form, *scientific AI*, which concentrates on the construction of working programs as experimental evidence for theoretical explanations of intelligent behavior. Schank and Abelson's (1977) description of conceptual dependency theory in natural language understanding is given as an exemplar. Thus, both forms of AI research rely on the construction of working programs. For applied AI, the programs are the end product while for scientific AI they serve as confirming experiments. Confirmation, as described, is simply a demonstration that a particular behavior could be achieved with a particular computational artifact.

Ringle's Analysis

Finally, Ringle (1979) proposes a relatively detailed research taxonomy consisting of four classes. The first class, *AI technology*, is similar to Newell's exploration, Weizenbaum's performance, and Hayes's applied perspectives. The approach taken is to construct reliable, cost-effective artifacts that demonstrate intelligent functioning without regard for human behavior or processes. As an example, Ringle cites the work of Buchanan *et al.,* (1969) on the analysis of mass spectrogram data in the DENDRAL project.

The second perspective, AI *simulation*, is concerned with overt human behavior but breaks into two approaches that differ in the extent to which internal human cognitive processes are considered. In the first approach, which Ringle terms "demonstrative simulation," computer programs are constructed to produce human-like behavior without regard for internal cognitive processing. As exemplars, Ringle suggests many of the early game-playing systems, which sought a human level of performance without concern for emulating human processing. According to Ringle (1983b), demonstrative simulations are no longer

undertaken in "mainstream AI." The second approach, "investigative simulation," involves computer demonstration of human-like behavior followed by hypotheses about similarities between machine processing and human cognitive activities. Newell and Simon's (1963) use of protocol analyses in validating the computational mechanisms of the GPS system (specifically means-ends analysis) is given as an example.

The third perspective in Ringle's taxonomy is termed *AI modeling* and involves the construction of computer programs that are intended specifically as models of internal human cognitive representation and processing. Ringle cites the work of Hunt (1973) on human memory as an exemplar. This modeling perspective is distinct from investigative simulation in that research moves from theory to object rather than constructing working programs and then advancing hypotheses of similarity between program and cognitive structures. Ringle (1983b) later suggests that simulation and modeling perspectives might be combined under the rubric of "cognitive simulation," with exploratory activity moving both ways between theory and object. The relationship between cognitive simulation in AI and research in experimental psychology is discussed at length (Ringle, 1983a).

Finally, Ringle describes *AI theory* in which general principles of intelligence are advanced without regard for particular implementations or human cognition. Ringle's claim is that this perspective amounts to a form of applied epistemology that, "when mature, will *subsume* the theory of human intelligence" (p. 12). As an exemplar, Minsky's (1975) theory of frames is given.

Perspectives in Perspective

If the reader at this point feels somewhat ambivalent towards the research perspectives discussed above, this paper will be in some measure a success. Perspectives overlap across the differing frameworks, and none clearly specifies a "modus operandi" for AI researchers that one might expect from an epistemological description of a disciplinary matrix within which research occurs. Specifically, none of the studies reviewed clearly describes issues of problem choice, practical methodology, or the critical evaluation of research reports.

The reasons for the sparseness of the foregoing descriptions are diverse and merit some consideration at this point. Of central importance, there is no consensus from within the field or among observers of the field as to what the term "artificial intelligence" means. While many might agree that *artificial* (particularily man-made) systems are worthy of study (Simon, 1969), there is little agreement on the extent to which intelligent systems should be artificial, in the sense that the artifact differs significantly from the original, natural object.

More importantly, however, the term *intelligence* has no agreed upon meaning. Before attributing this confusion

to AI specifically, we should note that the meaning of intelligence was a hotly contested issue many years before the emergence of AI, and remains unresolved unless elusive operational definitions like "intelligence is what intelligence tests measure" are acceptable. Hence, we when researchers speak of choosing a problem task that is assumed to require "intelligent" (perhaps human-like) behavior, the space of potential problems is hardly well-defined.

In fact, the definition of intelligent behavior (whether human or otherwise) appears to reflect the goals of the researcher more than any attributes of the behavior itself. For example, in Ringle's taxonomy it is unclear whether research problems (in the form of some task domain) that interest AI "technologists" would interest followers of any of the other perspectives. As Ringle (1983a) suggests, the method and intent of the researcher appear to overshadow choice of a domain. Hence, researchers might approach a particular task domain (e.g., vision) from differing methodological perspectives and achieve strikingly different results.
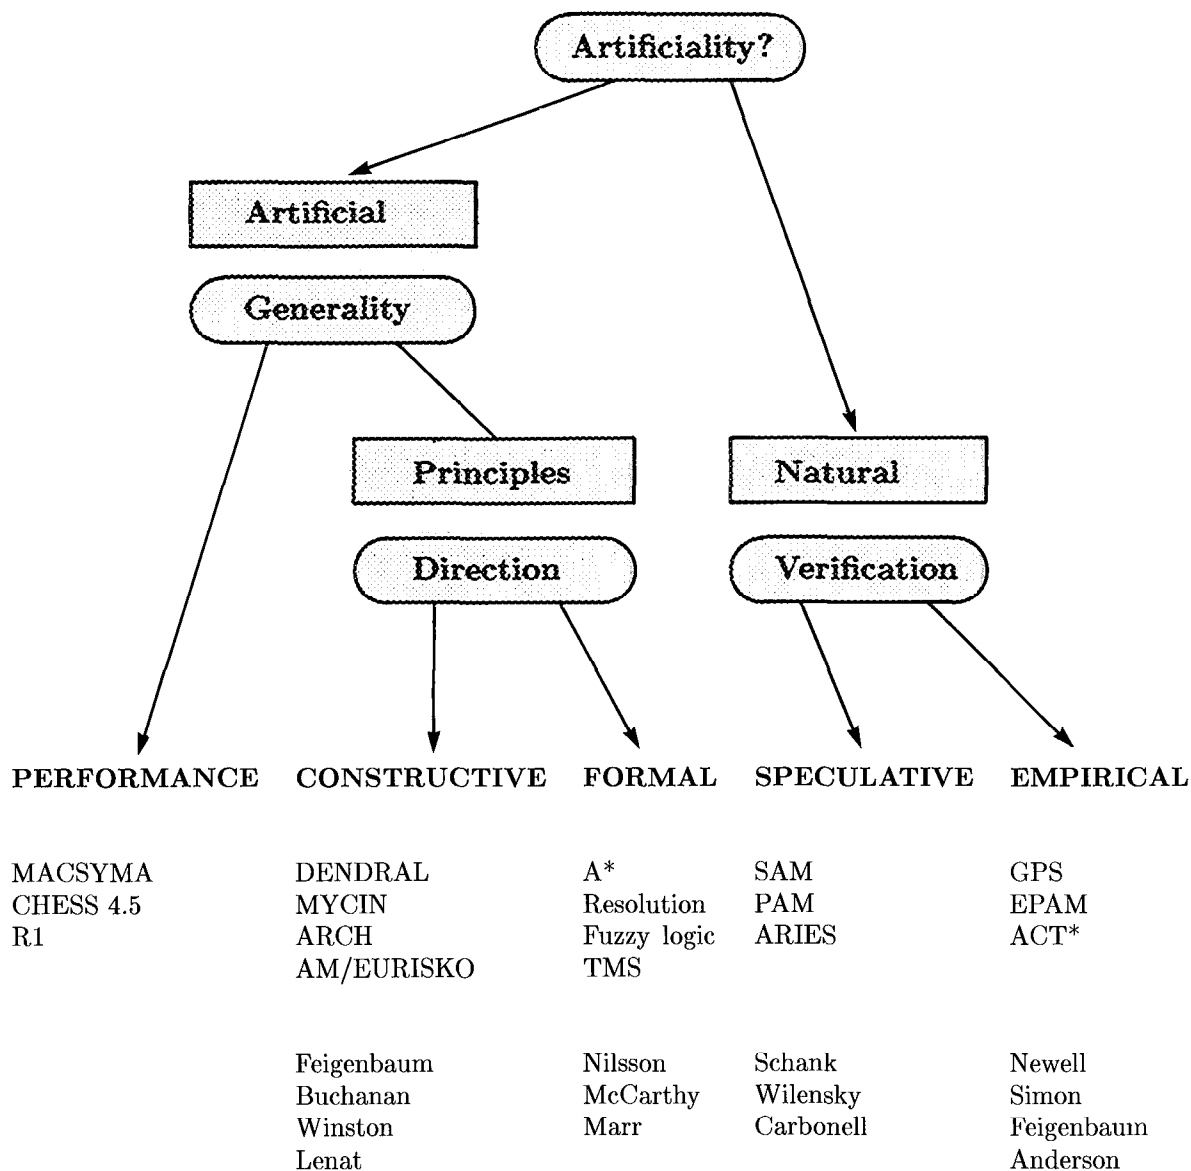
Considering method and intent more directly, we can to identify interesting common points among the analyses reviewed. With respect to intent, two issues seem paramount and reflect the definitional difficulties described above. First, research perspectives appear to diverge over the issue of the *artificiality* of the computational mechanisms under study. Of the six analyses reviewed, all but Lenat (1978) and Hayes (1978) find an explicit distinction between research undertaken with the intent of exploring human cognitive phenomena and studies of "intelligent" functioning by any means possible (*i.e.*, computational mechanisms).

Second, perspectives can be differentiated in terms of the *generality* of reasoning methods sought. Newell's (1973) distinction between the exploration of intelligent functions and the study of weak methods clearly exploits opposite ends of this generality continuum. In fact, all but Feigenbaum (1977) and Lenat (1978) make a similar distinction between computational techniques developed for a specific applications area and techniques explored with the intent of discovering general reasoning mechanisms that will be effective across varied task domains. In the extreme, endorsement of generality leads to research perspectives that purport to encompass both computational and human reasoning capabilities. An example would be Ringle's (1979) AI theory, which, as described, may eventually "subsume" theories of human intelligence.

In sum, dimensions of artificiality and generality play an important role in differentiating perspectives. Table 1 shows how the analyses reviewed have fit these dimen-

	Artificial	Natural
Specific	Newell: exploration Weizenbaum: performance Hayes: applied Ringle: technology	
General	Newell: weak methods Weizenbaum: theory Feigenbaum: engineering Lenat: unified theory Hayes: scientific Ringle: theory	Newell: theoretical psychology Weizenbaum: simulation Feigenbaum: theoretical psychology Ringle: simulation

Partitioning of reviewed perspectives by dimensions of artificiality and generality.

Table 1.

The figure is a tree diagram with nodes: **Artificiality?** branching to **Artificial** and **Natural**. **Artificial** leads to **Generality**, which branches to **Principles**. **Principles** leads to **Direction**. **Natural** leads to **Verification**.

PERFORMANCE	CONSTRUCTIVE	FORMAL	SPECULATIVE	EMPIRICAL
MACSYMA	DENDRAL	A*	SAM	GPS
CHESS 4.5	MYCIN	Resolution	PAM	EPAM
R1	ARCH	Fuzzy logic	ARIES	ACT*
	AM/EURISKO	TMS		
	Feigenbaum	Nilsson	Schank	Newell
	Buchanan	McCarthy	Wilensky	Simon
	Winston	Marr	Carbonell	Feigenbaum
	Lenat			Anderson

A typology formed from differing intentions and methods of AI researchers. Exemplary research projects are listed for each perspective along with well–known participants.

Figure 1.

sions. It is interesting that none appears to fit within the cell formed by interest in specific reasoning mechanisms for human intellectual capabilities. In fact, it is difficult to find AI studies undertaken with these goals. This is not the case in other research approaches to intelligence, however. In experimental psychology, for example, differential theories of human intelligence provide an interesting history of just such a division between specific and general

intellectual abilities (Sternberg, 1977). The issue of generality in computational studies of human reasoning does not appear to have become so well defined.

Not only are cells of Table 1 incompletely filled, but the two factors of artificiality and generality do not account completely for the diversity of research activity existing within particular cells. For example, it is not clear that Feigenbaum's (1977) engineering perspective overlaps very cleanly with Newell's (1973) description of weak methods. Additional factors concerning methodological approach must also be considered. In the next section, we present a five-part typology that both covers existing descriptive frameworks and speaks directly to issues of problem choice, methods, and criteria for assessing quality.

An Epistemological Framework for AI Research

In attempting to delineate a set of AI research perspectives with respect to problem choice, methodological approach, and criteria for evaluation, we do not claim that the resultant typology is wholly accurate. However, we offer empirical evidence for the fit of these perspectives in the form of an examination of selected studies. In addition, we will attempt to demonstrate the diversity of results that have been achieved within each perspective and suggest that each perspective is productive within the context of its underlying methodological assumptions and that results are often shared across perspectives.

Artificial Versus Natural

As noted earlier, most of the analyses reviewed in this paper have much in common. We propose a unifying typology of research perspectives, described by the tree structure shown in Figure 1. In particular, there seem to be two primary classes of research activity separated by the level of concern for explaining natural (*i.e.*, human or animal) intelligence. This reflects the dimension of artificiality discussed previously. Within these broad classes, termed *artificial* and *natural* for reference, further subdivisions follow on the basis of relatively coarse grounds of methodology and intent.

Performance Versus Principle

For researchers not concerned with exclusively natural modes of intelligent functioning, an intentional distinction is possible between those interested solely in program *performance* and those interested in uncovering more general *principles* of intelligence. This reflects the dimension of generality discussed previously, and Ohlsson's advocacy of investigating principled mechanisms for information processing (1983) exemplifies prescriptive concern for this issue. The MACSYMA system (1975) for symbolic manipulation of algebraic formulas would be an example of the performance perspective.

Constructive Versus Form

For researchers interested in more general principles, a methodological division similar to Ringle's distinction between simulation and modeling appears useful. One approach to such general principles is a bottom-up construction of functional AI systems in "real-world" (*i.e.*, difficult) task domains in the hope that general principles useful for intelligent systems will emerge as computational requirements are accommodated. The work of Buchanan *et al.*, (1969) on DENDRAL provides an exemplar of this *constructive* approach in which the researcher moves from object (program) to theory (general principles).

In contrast, many AI researchers interested in general principles of intelligence use a top-down approach in which computational structures and mechanisms are described formally (usually independent of any particular task) before being demonstrated in domains that are often highly constrained. In Ringle's terminology, researchers proceed from theory (general principles) to object (programs). Hart, Nilsson, and Raphael's (1968) work on heuristic search in graphs is an exemplar of this *formal* approach. We should not exaggerate this directionality between theory and object, however, both constructive and formal approaches evidence an interplay of inductive (bottom-up) and deductive (top-down) activity common in other fields of scientific inquiry.

Empirical Versus Speculative

For AI researchers primarily concerned with the character of naturally occurring intelligent systems, we can also define a subdivision along methodological lines between researchers who present explicit demonstrations of correspondence between behavior of artificial and natural systems and researchers who profess an interest in natural intelligence, but make less rigorous efforts to demonstrate empirical evidence for correspondence between artificial and natural systems. Newell and Simon's (1972) work on human problem solving can be considered an exemplary of the former, *empirical* perspective, while Schank's (1977) description of episodic memory organization might serve as an exemplar for the latter, *speculative* perspective. We should note that with this last methodological distinction, we do not intend to contrast "careful" and "reckless" computational approaches to natural intelligence. Rather, empirical and speculative perspectives represent alternatives in terms of problem choice, preferred methods, and criteria for successful results. These distinctions will be discussed more fully in following sections.

In addition to the studies cited above, Figure 1 lists other studies in order to give some feeling for the diversity of work within each particular perspective. These studies were selected because they are well-known; they are listed in roughly chronological order beneath the exemplars for each perspective. Beneath these studies, well known proponents of each perspective are listed. As we will discuss

shortly, consistent proponents of the performance perspective are not evident in the AI literature, although instances of such work continue to appear.

Shifting and Sharing

To suggest that these five perspectives describe completely disjoint communities of practicing researchers would be misleading. In fact, even a casual inspection of the literature suggests that people shift between perspectives over time and that issues are shared across perspectives with differing research goals. Feigenbaum provides a clear example of a researcher shifting between perspectives. His thesis research with the EPAM simulation of human verbal learning (Simon and Feigenbaum, 1964) falls within the empirical tradition, while much of his subsequent work in knowledge engineering (Feigenbaum, 1977) clearly represents the constructive perspective. Interestingly, a recent paper by Feigenbaum and Simon (1984) gives a defense of major findings with the earlier empirical work. Thus AI research cannot be partitioned solely on the basis of major participants (a possibility not discussed in earlier sections). Feigenbaum's activities in multiple perspectives are not aberrant or unusual; it is easy to imagine other situations in which individuals might shift in a similar fashion. For example, a researcher doing formal work on knowledge representation techniques might collaborate on the construction of some practical artifact, with reports of both activities reaching the literature.

Alternately, it is possible for research undertaken in one perspective to initiate complementary activity in a different perspective, resulting in a sharing of ideas across multiple perspectives. For example, the many contexts in which search proved an important component of constructive and performance studies undoubtedly provided much of the initial impetus for more formal approaches to heuristic search. Consider also the widespread use and endorsement of semantic nets. Initiated in the empirical tradition with the work of Quillian (1968) on modeling human semantic memory, this network-like formalism for representing information has since been so widely adopted as to be found in published reports from each of the other methodological perspectives.

While it might be argued that semantic nets are example of a shared knowledge representation tool that is independent of the user's methodology, this network formalism serves clearly different roles in studies conducted within different methodological perspectives. Lenat's (1977) use of a network formalism for representing mathematical concepts, an example of the constructive perspective in our typology, is clearly different from Etherington and Reiter's (1983) work on inheritance hierarchies with exceptions, a strong exemplar of our formal perspective. The former uses semantic nets precisely in the sense of a tool, while the latter study takes network formalisms as the central object of study. (Brachman (1979) discusses the varied history of semantic networks in detail.) In summary, there is ample evidence that the same idea can be shared across perspectives, but that this sharing is tempered by dramatically different focus.

Because of this tendency on the part of researchers and issues to shift quite freely between perspectives, the typology we are presenting might better be considered as a description of *ideal types* from which variance is to be expected. Nonetheless, we contend that published research reports can generally be identified as subscribing to one of the five perspectives, and that such a classification can be useful for readers wishing to evaluate the quality of these reports and their relation to other work. To justify this claim, we must sharpen the methodological and intentional divisions depicted earlier by giving more detailed descriptions of the three components a useful research perspective should provide: guidance in problem choice, a characteristic methodology, and a set of criteria by which good work can be identified.

Perspectives in Particular

In this section, we will describe the three components mentioned above for each of the five proposed research perspectives and we will attempt to assign studies to each of the perspectives. For each perspective, these exemplars are examined first, followed by a more general summary of perspective components.

Performance AI

As an exemplar of AI research oriented entirely towards impressive levels of performance, we will look at the widely used MACSYMA system for online algebraic manipulation (Moses, 1971; MACSYMA, 1975; Barr and Feigenbaum, 1982). Unlike much previous work on symbolic integration which was concerned with general AI techniques, the MACSYMA project displays a clear goal of integration, performance at or beyond levels of human performance. Extensive domain-specific expertise (often in the form of highly specific mathematical algorithms) is used to generate solutions without regard for human approaches to similar problems or the relationship of computational mechanisms to established AI techniques.

As a problem choice, symbolic integration represents a constrained, nontrivial task that is often difficult for humans. In addition, substantial levels of performance in this problem domain promise to be of considerable use to individuals who regularly face difficult integration problems in their work (*e.g.*, researchers in plasma physics).

The methodology of the MACSYMA project is to develop and use whatever integration techniques seem promising for the solution of particular classes of symbolic problems. Although representational and processing issues of general interest to AI do emerge (*e.g.*, inheritance hierarchies to guide inferences over symbol types), there is little

a priori interest in such issues. Rather, the emphasis is on what techniques might be applied so as to minimize the time/space complexity of generating solutions within the problem domain or to make the system more accessible to users.

As a criterion for success, the usefulness of the MAC-SYMA system for online users looms large. In addition, comparisons are drawn between successively more sophisticated system capabilities in terms of the classes of problems that can be solved and the time/space complexity of solutions to those problems. In the case of MACSYMA, program performance actually exceeds that of all but a few human experts.

To generalize, the performance AI perspective selects problems often undertaken with some difficulty by humans. Hence, moving toward solution of these problems would typically be of some practical significance. Methods consist of developing and applying processing and representational techniques that approach solutions with minimal computational requirements. As evidence of success, practical usefulness and improved performance over previous computational approaches weigh heavily.

Constructive AI

As an example of the constructive (bottom-up) approach to general principles of intelligence in AI, we will examine the heuristic DENDRAL project (Buchanan *et al.*, 1969, 1970; Feigenbaum *et al.*, 1971). Taking the study of scientific hypothesis formation as a general goal, the DENDRAL project can be seen as a study of the merits of generality versus expertise with respect to the performance power of AI techniques. This effort is not to be confused with the performance perspective, and interestingly enough, the investigators on the DENDRAL project are careful to point out that "attention given to the program as an application of artificial intelligence research has tended to obscure the more general concerns of the project investigators" (Feigenbaum, *et al.*, 1971, p. 166).

As a problem choice, molecular structure elucidation in organic chemistry serves as a complex, real-world problem which is "complex enough and rich enough in internal structure and theory to provide many firm foundation points on which to erect a meta-level for the study of theory formation processes" (Feigenbaum, *et al.*, 1971, p. 187). The problem is chosen to provide a "forcing function" in which domain requirements will guide system design and illuminate representational and processing issues of more general significance in AI.

Methodologically, the DENDRAL project proceeds as an iterative interplay between program design/construction and performance/experimentation, with general issues emerging in the process of accommodating demands in the task domain. For example, in attacking the amine family of chemical structures, the need for strong heuristic constraint on the space of possible structures leads to

the construction of a powerful planning mechanism. This development, combined with the incorporation of simple hypotheses concerning likely fragmentation patterns, introduces problems of consistency among multiple sources of knowledge, which in turn leads to an appreciation for the desirability of separating knowledge representation from processing details. Reflecting on this chain of events initiated by an attempt to constrain search, the investigators report "there are a number of ways to do this, some of which were tried with success, some with failure. The failures were at least as illuminating as the successes" (Feigenbaum, *et al.*, 1971, p. 171). As might be expected, rewriting substantial sections of code is described as a common activity in the project (Buchanan, *et al.*, 1970). Hence, demands originally quite specific to the task at hand can be seen to force consideration of issues of general significance in AI.

A variety of evidence is given for the success of the DENDRAL project. Comparisons of the complete space of structural candidates with generated and suggested candidates are given as evidence of efficient and correct structure determination for selected molecular classes (Feigenbaum, *et al.*, 1971). In addition, published reports of candidate structure spaces appearing in widely read chemistry journals are given as evidence that chemists found DENDRAL's performance interesting. Favorable comparisons with structure identification by human experts (graduate students and a post-doctoral fellow) for selected molecular classes are also given (Buchanan, *et al.*, 1970). Finally, arguments are made that the DENDRAL programs can be extended to accommodate new molecular classes and rapidly accumulating theoretical knowledge of mass spectrometry.

In summary, the DENDRAL project provides considerable insight into the manner in which constructive AI is routinely done. Complex, real world problems are chosen to create an experimental design atmosphere in which issues of general interest to AI research are regularly forced into active consideration. Solutions to these design problems, as incorporated into functioning software systems, are evaluated in terms of efficiency, credibility with human experts in the domain area, and demonstrated or promised extensibility.

Formal AI

The original work of Hart, Nilsson, and Raphael (1968) and subsequent descriptions (Nilsson, 1971; 1980b) of heuristic search serve as widely read exemplars of the formal perspective in AI research. Apart from any concern for naturally intelligent behavior, these reports provide an abstract framework for using domain-specific information in determining minimum cost solutions for a large class of specific problems, expressed more generally as graph search problems.

As a problem choice, heuristic graph search represents

an abstraction of problems encountered in many applications areas (*e.g.,* navigational routing, circuit design, or problem solving). It is the intent of the authors to give a general theory of heuristic search that encompasses a variety of techniques previously reported in the AI literature.

Methodologically, these reports give a formal problem definition of finding minimum cost paths for a restricted class of graphs that serve as a general representational medium for a variety of search problems. A generalized algorithm is developed, which, using suitably restricted evaluation functions to determine which node to consider next, can be shown always to yield a minimum cost solution path between start and goal nodes, providing such a path exists. Claims for the correctness of this general algorithm are proven, and in later reports (Nilsson, 1971; 1980b), performance comparisons are made between particular algorithms using differently informed evaluation functions.

Criteria for success in these reports include acceptable proofs of algorithm correctness and performance increases for more "informed" versions of the algorithm. More generally, it is shown that the proposed formalism does indeed cover a wide class of search techniques, from blind search to heuristic search, in which the chosen evaluation function provides a relatively tight lower bound on actual minimum cost solution paths.

To generalize, problem choice in the formal AI perspective appears to focus on recurring problems across multiple domains, which identify the need for general techniques. Hence, general techniques should not be considered to arise in a vacuum, rather they emerge as a result of a perceived need for an encompassing formal framework for some related classes of existing problems. Methodologically, work proceeds by giving a formal problem specification, detailing some computational mechanism (*e.g.,* an algorithm) for solving the problem, and then giving some justification for the appropriateness of that mechanism. Although not appearing in the study described above, a demonstration of the proposed general technique in a particular (typically constrained) task environment is commonly used. Successful work in this perspective requires unambiguous and adequately descriptive specifications of problem and solution mechanism for an "important" class of related specific problems. A problem class can be considered important if particular manifestations of the problem recur in the literature and these various manifestations can be meaningfully viewed as members of a more general class of problems which have yet to see an encompassing solution. Demonstration of computational sufficiency must be convincing (*e.g.,* proofs or assumptions should be believable), and comparisons with alternate methods of solution should be favorable.

Speculative AI

Schank and Abelson's (1977) treatise on natural language understanding stands as a clear exemplar of AI research in naturally occurring intelligent behavior without the onerous task of empirical verification that is characteristic of empirical AI. For Schank and Abelson, the focus of research is squarely on proposing a theory that can account for human abilities in understanding and generating routine connected discourse in a natural language. Although discussion periodically turns to a more general theory of "knowledge systems" that might encompass both human and machine performance, concern for human functioning is clearly emphasized.

As a problem choice, the authors constrain the immense domain of natural language use by focussing on what they term the "naive psychology" and "naive physics" of everyday human discourse (Schank and Abelson, 1977, p. 4). Rather than strictly defining a task environment, these constraints are taken to provide a starting point with relatively simple (*i.e.,* common-sense) forms of knowledge. Problem choice within this perspective is discussed more directly by Lehnert (1984) in terms of local versus global task orientations. Rather than choosing highly constrained (local) problems typical of traditional psychological investigation (*e.g.,* free word association), research within the speculative perspective considers self-contained tasks of broad (global) scope which occur as part of the normal human behavioral repertoire (*e.g.,* question answering).

Methodologically, the approach taken by Schank and Abelson marks a sharp divergence from traditional psychological or linguistic approaches to language use, which the authors find unnecessarily restrictive. In their words, "we are willing to theorize far in advance of the usual kind of experimental validation because we need a large theory, whereas experimental validation comes by tiny bits and pieces" (Schank and Abelson, 1977, p. 7).

More rigorous demonstrations of empirical fidelity characteristic of the empirical perspective are attacked on the the basis of ecological validity. In short, Schank and Abelson question the utility of studying narrowly constrained task domains (a local task orientation in Lehnert's terminology) with experimental materials stripped of external validity (*e.g.,* memorization of nonsense syllables). Impatience with and distrust of empirical demonstrations of theoretical validity are offset by what is described as a painstaking process of implementing theorized mechanisms of natural language understanding in clearly specified computer programs. Difficulties with program construction are taken as theoretical inadequacies that must be remedied before implementation will be successful. The source of theoretical speculations appears to be primarily introspection guided by "intuitive necessity" and "internal consistency" (Schank and Abelson, 1977, p. 21). Occasionally, observations of the behavior of others (*e.g.,* the daughter of one of the authors) are used in discussing theoretical propositions, but these observations appear to be used more as anecdotes to motivate theoretical speculation

than as empirical demonstrations of theoretical validity. Lastly, separate programming efforts (*e.g.*, the script and plan applying systems, SAM and PAM) are described as working implementations of theoretical components.

Criteria for success amount to the reader's sense of psychological plausibility bolstered by evidence in the form of working programs. For example, the relatively appealing proposition of script-based episodic memory for stereotypical human experiences is supported with output of the SAM (script applier mechanism) program, that includes descriptions of events which are never mentioned as part of the input story, underscoring the importance of expectations in understanding natural language.

In summary, the speculative AI approach to studying naturally occurring intelligent behavior selects problems that are common representatives of the system under study, unlike the rather restricted problems chosen in the empirical perspective discussed next. Introspection guides the formation of theoretical propositions concerning intelligent behavior, which are then tested by attempting to embody those propositions in clearly specifiable computer programs. To the extent that implementational difficulties arise, theoretical propositions are reconsidered. Successful work within this perspective consists of working programs that are taken as support for the sufficiency of psychologically plausible theories.

Empirical AI

The widely known GPS (General Problem Solver) project of Newell and Simon (1963, 1972) serves as an exemplar of the empirical approach to modeling naturally intelligent systems—in this case, the behavior of humans in well-defined problem-solving tasks. There is little doubt that these authors intend a model of human cognition (*e.g.*, "GPS, a program that simulates human thought," Newell and Simon, 1963), and their published reports are quite strongly connected to the psychological literature.

As a problem choice, Newell and Simon suggest human tasks for which AI can provide potential representational and computational strategies (*i.e.*, data structures and means of accessing them that might have some psychological validity). Cryptarithmetic puzzles, theorem proving in logic, and chess are chosen for detailed explication in the GPS project. Extension of a similar experimental approach to a wider range of human functioning is suggested (Newell and Simon, 1972), but only after a thorough understanding of human behavior in less complex domains.

The primary methodology evident in the GPS project is a sustained interplay between program construction and a comparison of program/human performance in the current domain. This comparison proceeds as a detailed ideographic analysis of verbal problem-solving protocols. Program revisions are proposed to accommodate discrepancies between program traces and human protocols. The level of detail at which this comparison is done varies, with con-

siderable interest in the identification of subprocesses in human performance that can be shown to correspond to the actions of subcomponents within the entire program For example, segments of the human protocol such as, "I'm looking for a way, now, to get rid of that horseshoe. Ah...here it is, R6" (Newell and Simon, 1972, p. 461) are taken to correspond to a search through the set of operators (rules of inference in this case) during a subject's first exposure to the GPS experimental task, before a "table of connections" (used in the GPS program to index rules by difference reductions) has been acquired.

As evidence for success of the GPS project, the investigators give detailed descriptions of correspondence between program functioning and human behavior. When discrepancies do arise, to the extent that the program can be modified without major reorganization, the program is considered a valid simulation of human problem solving. Interestingly, more serious discrepancies (*e.g.*, the tendency of some subjects to backtrack and correct previous rule invocations) are viewed positively as uncovering additions which "could significantly increase the total capabilities of the program" (Newell and Simon, 1972, p. 472) rather than as failures. Finally, the viability of the theory (*i.e.*, humans and programs as information processing systems) is demonstrated over multiple task domains as evidence of general applicability.

More abstractly, the empirical AI perspective can be seen to choose well-defined problems that natural systems perform well. Problems chosen within the empirical perspective appear to be more narrowly constrained (of a more local orientation, to use Lehnert's terminology, 1983) than those selected within the speculative perspective. Beyond the exemplar discussed here, this tendency might be seen to extend into choosing tasks for which some empirical database exists with respect to behavior of natural systems. The methodology is essentially experimental, incrementally modifying the artificial system as a result of careful comparison with behavior of the natural system. Newell and Simon's rather strict reliance on ideographic analysis should not be generalized to the perspective as a whole. In general, the primary methodological point is that of detailed, empirical comparison between the program as model and the naturally occurring system. Criteria for success include: empirically demonstrable correspondence between program performance and natural behavior, robustness of program design with respect to incremental changes (taken as an indicator of the fidelity of the model), and (perhaps less crucial than the previous criteria) the extensibility of model concepts to varied task domains.

Conclusion

The review of descriptive methodological analyses and the five-fold synthesized typology we have presented make four central contributions to AI.

First, we provide a description of what is actually done in AI research, using five of the most widely known published studies in the field as examples. This is in contrast to the more common prescriptive analyses of methodological issues in AI which, when they mention existing work in the field at all, tend to focus on selected issues from a single perspective. Bundy (1983a), for example, argues strongly for investigation of computational techniques from what appears to be a constructive perspective.

Second, we provide a typology of divergent perspectives that helps to explain persistent, polemic disagreements between major players in the field of AI. For example, the much heralded "Great Debate" between Schank and McCarthy over methodologies for AI research at AAAI-83 can be understood as a public confrontation between major proponents of speculative and formal perspectives, respectively. While we will not argue that the typology presented in this paper is necessarily wholly accurate, we will claim that this methodological description makes the pluralistic nature of AI research more explicit than have previous reviews.

Third, this paper gives a clearer explication of the assumptions underlying research perspectives than can be found in previous descriptive analyses. Although we will not claim to have described such assumptions completely, we have begun a process of description that contributes to an understanding of AI as a field with multiple, competing research communities.

Finally, we hope to have presented a methodological framework which readers might use to make published reports in this field more accessible and more easily evaluated. Even if explicit attention to methodological assumptions is not included in such reports, an awareness of these issues on the part of the reader is desirable.

We should in fairness also be explicit about what this paper does not do. First, despite our enthusiasm for descriptive methodological analyses, this paper does not completely describe the field. Guidelines for problem choice, preferred methodological approach, and criteria for success are not developed at a level of detail that would allow unambiguous classification of an arbitrary study. Perhaps this is not a reasonable or even desirable goal, however, since a perspective describes the shared aspirations of a community of researchers rather than a particular study. Classification of a particular study is further complicated by the fact that individuals and issues appear to shift between perspectives over time. We do hope, however, that the arguments we have presented will serve as an impetus for a wider consideration of differing methodologies in AI.

Second, we do not feel that detailed examination of individual studies or our own classification of other studies gives sufficient empirical support for the methodological typology as described. Other avenues of investigation, such as an analysis of citation activity or funding histories, might provide more convincing evidence for particular perspectives. One corroborating source is an intruiging historical survey of "intellectual issues" in the field of AI by Newell (1983). Although his discussion is more wide ranging, Newell identifies several issues which have divided the field of AI at various times. These issues closely parallel the intentional and methodological divisions used in our five-fold typology, and, in Newell's opinion, they have yet to be resolved. Again, rather than claiming to have completely described this situation, we hope only to have sparked a wider range of descriptive activity.

Third, and most emphatically, we have not suggested which methodological perspectives are "most promising" and thus what sorts of work future AI research "should" attempt. In fact, we have presented an explicit argument that judgments of promise must be evaluated relative to goals which differ across perspectives. Hence, without a reasonable understanding of the diversity of actual practice in AI, prescriptive judgments are at best premature and at worst parochial.

Contributions and shortcomings of this article aside, we must return to the practical issue of making sense out of the bewildering array of activity in AI. Given an arbitrary research report from the AI literature, an attempt to classify it according to the current typology would require careful consideration of a variety of issues. Unfortunately, methodological orientations are seldom presented with any clarity in published research reports, making the identification of what was actually done or intended as part of a research project quite difficult to determine. As was pointed out quite candidly after a recent conference (Ohlsson, 1983), researchers in AI have a proclivity for writing and talking about what they would like for their research to demonstrate, as opposed to what is actually being demonstrated. If the accepted vocabulary of AI is, indeed, "about as precise as that of poetry and about as substantive as that of advertising copy" (Doyle, 1983, p. 53), some changes in the manner in which research reports are presented would seem desirable.

We have argued not only for a typology of methods in AI research, but also for the importance of making one's methodological orientation explicit when communicating results with the rest of the AI community. We hope that this will mark a beginning in what should be an ongoing public discussion aimed at describing the current state of AI research. The benefits of this sort of discussion, we hope, are clear. Not only might adherents of divergent approaches begin to appreciate, or at least to understand, the motivations of other researchers, but the field might become more solidly accessible to newcomers.

References

Anderson, J R (1983) *The architecture of cognition* Cambridge, MA: Harvard University Press

Barr, A., & E Feigenbaum (1982) *The handbook of artificial intelligence Vol 2* Los Altos, CA: WIlliam Kaufmann, Inc.

Brachman, R J (1979) On the epistemological status of semantic networks In N V Findler (Ed), *Associative networks: Representation and use of knowledge by computers* New York: Academic Press

Brachman, R J , & B C Smith (1980) *SIGART Newsletter 70*, special issue on knowledge representation.

Buchanan, B., Sutherland, G , & E A Feigenbaum (1969) Heuristic DENDRAL: a program for generating explanatory hypotheses in organic chemistry. *Machine Intelligence* 4: 209-254

Buchanan, B , Sutherland, G. L , & E A Feigenbaum (1970) Rediscovering some problems of artificial intelligence in the context of organic chemistry *Machine Intelligence* 5: 253-280

Bundy, A (1982) What is the well-dressed AI educator wearing now? *AI Magazine* 3, (1): 13-14.

Doyle, J (1979) A truth maintenance system *Artificial Intelligence* 12(3), 231-272

Doyle, J (1983) What is rational psychology? Toward a modern mental philosophy *AI Magazine*, (4): 3:50-53

Etherington, D. W., & R Reiter (1983) *On inheritance hierarchies with exceptions* AAAI-83, 104-108

Feigenbaum, E A , & Feldman J A (1963) *Computers and thought* New York: McGraw-Hill

Feigenbaum, E A. (1977) *The art of artificial intelligence—themes and case studies of knowledge engineering* IJCAI-5, 1014-1029

Feigenbaum, E A , & H A Simon (1984) EPAM-like models of recognition and learning *Cognitive Science* 8: 305-336.

Green, C (1969) *Application of theorem proving to question-answering systems* IJCAI-1, 219-237

Hart, P E , Nilsson, N., & B Raphael (1968) A formal basis for the heuristic determination of minimum cost paths *IEEE Transactions on Systems Science and Cybernetics* v SSC-4(2):100-107

Hayes, P (1978) Trends in artificial intelligence *International Journal of Man-Machine Studies* 10:295-299

Hunt, E (1973) The memory we must have. In Schank & Colby (Eds) *Computer models of thought and language* San Francisco: W H Freeman

Kuhn, T S. (1970) *The structure of scientific revolutions, 2nd ed* Chicago: University of Chicago Press

Lehnert, W G (1984) Paradigmatic issues in cognitive science. In W. Kintsch, J.R. Miller & P.G. Polson (Eds) *Methods and tactics in cognitive science* Hillsdale, NJ: Lawrence Erlbaum

Lenat, D.B. (1977) *Automated theory formation in mathematics* IJCAI-5, 833-842

Lenat, D B (1978) The ubiquity of discovery *Artificial Intelligence* 9:257-285.

Lenat, D B (1982) The nature of heuristics *Artificial Intelligence* 19:189-249

MACSYMA (1975) *MACSYMA reference manual*, The Mathlab Group, Project MAC, MIT

Marr, D. (1977) Artificial intelligence -- a personal view *Artificial Intelligence* 9:37-48

Marr, D (1982) *Vision* San Francisco: W H Freeman

Masterman, M. (1970) The nature of a paradigm. In Lakatos & Musgrave (Eds) *Criticism and the growth of knowledge* Cambridge: Cambridge University Press

McCarthy, J , & P.J Hayes (1969) Some philosophical problems from the standpoint of artificial intelligence In D Michie & B Meltzer (Eds.) *Machine Intelligence* 4:463-502.

McDermott, J. (1982) R1: a rule-based configurer of computer systems *Artificial Intelligence* 19:39-88

Minsky, M (1975) A framework for representing knowledge. In P. Winston (Ed) *The psychology of computer vision* New York: McGraw Hill

Moses, J. (1971) Symbolic integration: the stormy decade *Communications of the ACM* 14(8):548-560

Newell, A., & H.A Simon (1963) GPS, a program that simulates human thought In Feigenbaum & Feldman (Eds) *Computers and thought* New York: McGraw Hill

Newell A (1969) Heuristic programming: ill-structured problems. In Aronofsky (Ed.) *Progress in Operations Research 3*, 360-414

Newell A , & Simon H A (1972) *Human problem solving* Englewood Cliffs, NJ: Prentice-Hall.

Newell A. (1973) Artificial intelligence and the concept of mind. In Schank and Colby (Eds.) *Computer models of thought and language* San Francisco: W H Freeman

Newell A. (1983) Intellectual issues in the history of artificial intelligence In F Machlup and U Mansfield (Eds) *The study of information: interdisciplinary messages* New York: John Wiley.

Nilsson, N.J. (1971) *Problem-solving methods in artificial intelligence* New York: McGraw Hill

Nilsson, N J (1980a) The interplay between experimental and theoretical methods in artificial intelligence. *Cognition and Brain Theory* 4 (1), 69-74.

Nilsson, N J (1980b) *Principles of artificial intelligence* New York: McGraw Hill

Nilsson, N J (1982) Artificial intelligence: engineering, science or slogan? *AI Magazine* 3, (1):2-9.

Ohlsson S (1983) Tell me your problems: a psychologist visits AAAI-82. *AISBQ 46*, 27-30

Quillian M R (1968) Semantic memory. In M. Minsky (Ed) *Semantic information processing* Cambridge: MIT Press, 227-270

Ringle M (1979) Philosophy and artificial intelligence In M. Ringle (Ed) *Philosophical perspectives in artificial intelligence*. Atlantic Heights: Humanities Press

Ringle M (1983a) Psychological studies and artificial intelligence *AI Magazine* 4, (1):37-43

Ringle M (1983b) Personal communication

Robinson J.A. (1965) A machine-oriented logic based on the resolution principle *Journal of the ACM 12*, 23-41

Schank R.C and Abelson R.P. (1977) *Scripts, plans, goals, and understanding* Hillsdale: Lawrence Erlbaum

Schank R C (1983) The current state of AI: one man's opinion *AI Magazine* 4, (1):3-8.

Shortliffe E H (1976) *Computer-based medical consultations: MYCIN* New York: North-Holland.

Simon H A (1969) *The sciences of the artificial* Cambridge: MIT Press

Simon H A & E A Feigenbaum (1964) An information-processing theory of some effects of similarity, familiarization, and meaningfulness in verbal learning. *Journal of Verbal Learning and Verbal Behavior* 3, 385-396.

Slate D J & L R Atkin (1977) CHESS 4 5—the Northwestern University chess program. In P W. Frey (Ed) *Chess skill in man and machine*. New York: Springer-Verlag

Sternberg R J (1977) *Intelligence, information processing, and analogical reasoning: the componential analysis of human abilities* Hillsdale, NJ: Lawrence Erlbaum.

Waltz D (1975) *Natural language access to a large data base: an engineering approach* IJCAI-4, 868-872

Weimer W B (1979) *Notes on the methodology of scientific research* New York: Lawrence Erlbaum

Weizenbaum J (1976) *Computer power and human reason: from judgment to calculation* San Francisco: W H Freeman

Wilensky R (1983) *Planning and understanding: a computational approach to human reasoning* Reading, Massachusetts: Addison-Wesley

Winston P H (1975) *The psychology of computer vision* New York: McGraw-Hill

Zadeh L A (1975) Fuzzy logic and approximate reasoning *Synthese* 30.