

Artificial Intelligence: What it is, and what it should be

Pei Wang

Department of Computer and Information Sciences, Temple University

<http://www.cis.temple.edu/~pwang/>

pei.wang@temple.edu

Abstract

What Artificial Intelligence (AI) is at the current time is not what it should be. The mainstream works in the field are on domain-specific problems and solutions. Instead, AI should focus on general-purpose systems that are adaptive to its environment, and can work with insufficient knowledge and resources. AI in this sense is unified with Cognitive Science at a general level, with the latter focuses on the theoretical aspect of the intelligence-cognition-thinking process in various types of systems (human, machine, etc.).

Introduction

The research of Artificial Intelligence (AI) started half a century ago, as an ambitious attempt to build thinking machines (Turing 1950; Newell & Simon 1963; McCorduck 1979). However, the initial excitement was followed by ungrounded hype and failed predictions. As a consequence, the field gradually turned away from the large topics and to much more special, practical, and manageable problems. This change is taken to be a mistake by some researchers who still believe in the AI dream (Stork 1997), but by many others as a sign of maturity of the field (Hearst & Hirsh 2000). Today, AI has made remarkable progress in its various subfields, though the major challenges remain unanswered (Stone, Hirsh, & others 2005). We have not built a thinking machine yet, and many people have given up this goal as unrealistic, even for the future. The field of AI is not unified by a shared theoretical foundation or a common goal, but by a class of loosely related problems.

Cognitive Science (CogSci) started at the realization that the principles and mechanisms in the brain and in the machine are similar, if not identical. Therefore, the problems studied in several fields, including AI, psychology, linguistics, philosophy, etc., have close relations among them, and should be attacked by an interdisciplinary cooperation. The “cognitive revolution” has more or less changed the involved fields, by encouraging them to borrow concepts and methods from each other, and there have been new ideas and results produced in this process. Even so, since each field has its own goal and interest, it is inevitable that most works labeled

as “CogSci” still bear clear marks of their “home field” (Wilson & Keil 1999). In recent years, the conferences and publications of CogSci have been dominated by cognitive psychology (Goldstone 2005). As a result, most of the works are experimental, targeted on the human mind/brain, without saying too much about the machine (though using the computer as a modeling tool).

Now the 2006 AAAI Spring Symposium on “Between a Rock and a Hard Place: Cognitive Science Principles Meet AI-Hard Problems” shows that there are still people who are concerned about the big picture of the two fields, AI and CogSci, as well as their relationship. To make the discussion more fruitful, the organizers provided a list of focus questions to cover the major aspects of the issue.

From my own experience, the answers to these questions mainly come from the answer to a deeper question, that is, “What is AI?” For the people who are unhappy about the current situation of AI, this question directly leads to another one: “What should AI be?” Of course we can ask the same questions about CogSci, but given the nature of this symposium, that is a secondary issue.

In the following, I will first make a position statement about the deep questions, then, based on it, answer each of the seven focus questions in a consistent manner.

Position statement

The following statement is a brief summary of my ideas on AI as a whole, based on my own research experience in AI and CogSci, especially my practice in the NARS project (Wang 1995; 2004; 2006a; 2006b). Since the statement inevitably touches many issues in the field, to fully discuss its justification and implication is far beyond the capacity of this paper. Instead, I will only present the statement in a general level, and leave the technical discussions to the other publications.

The results of a complete AI research project should consist of:

1. a *theory* on the principles and mechanisms of intelligence,
2. a formal *model* of intelligence based on the above theory,
3. a computer *implementation* of the above model.

The first two items form the *science* aspect of the project, while the last two form the *engineering* aspect of the project.

In the context of AI, the concept “intelligence” should be understood as the general-purpose capability of *adaptation to the environment when working with insufficient knowledge and resources*. Concretely, it means that the system must only assume *finite* time-space supply, always *open* to new tasks, process them in *real time*, and *learn* from its own experience.

CogSci, as the study of the processes like cognition, intelligence, mind, and thinking in general, should be carried out at two levels of abstraction:

1. It should try to identify general principles and mechanisms of these processes in their various forms (in human, machine, animal, group, etc.).
2. It should also study each of these concrete forms of these processes, as special cases of the general principles and mechanisms.

According to this opinion, AI and CogSci should be closely related, even with a large overlap, though the two fields are not identical to each other. Concretely, the theory and model of AI should be based on the result of CogSci at the general level, and the implementation of AI should be one form (though not the only form) of realization of the theory and model.

My own research on intelligent reasoning systems is an attempt to study AI and CogSci according to the above understanding. This research has produced many novel and interesting results, as described in my publications.

The current mainstream AI research is not what it should be, because most of the works are domain-specific, many (though not all) of them ignore the requirement of adaptation, and most of them make various kinds of assumptions on the sufficiency of knowledge and/or resources. Though they are still sound scientific research that produce valuable results, they miss the essence of intelligence that distinguishes the human mind from traditional computer systems. As a result, they fail to find the identity of AI that makes it fundamentally different from conventional computer science and application, fail to produce novel concepts and methods that work in environments where the traditional ones do not work, and fail to contribute to the ultimate goal of AI, that is, to build thinking machines.

The current mainstream CogSci research is not what it should be, either, because most of the works ignore the difference between “general cognition” and “human cognition”. As a result, though the research results are valuable for the study of the human mind-brain complex, they do not provide enough help for the related fields, especially AI. Even for the small number of AI works that are affected by CogSci research, the influence is often misleading, because they tend to take “faithfully duplicating human cognition/intelligence” as the preferred way to achieve artificial cognition/intelligence, which, though a possibility, is still far from justified as the best way to achieve real AI.

Questions and Answers

In the following, each of the focus questions is answered by the implications of the above opinion on AI and CogSci in general.

1. Is cognitive science relevant to AI problems?

Yes, CogSci and AI should have a close relationship. The science aspect of AI and the general level of CogSci are about the same issue, to a large extent.

The current mainstream AI ignores the results of CogSci, because it has been focused on domain-specific problems. The solutions obtained in these situations usually cannot be generalized, and they are often very different from how the same problems are solved by the human mind.

For example, in many AI approaches the notion of a “concept” is still used in the “classical” sense, where it is defined by a sufficient and necessary criterion, and whether something is an instance of the concept is a matter of yes/no. On the contrary, in CogSci the classical theory has been abandoned long ago as a description of human categorization, except as a rough approximation in certain special situations. Many new models of categorization have been proposed and studied in CogSci, including prototype theory, exemplar theory, and theory theory. Though each of them still has its own problem, in general they all fit better with the reality of human cognition than the classical theory does (Laurence & Margolis 1999).

Many AI researchers may be aware of this situation, but take it as irrelevant — if they do not attempt to simulate the human mind, why does it matter if their categorization model is not the one used in the mind?

It matters because the classical theory assumes a static and simple domain. For such a domain, a theory equipped with classical concepts can provide a complete and accurate description. The AI systems developed in this way work well in these domains, but cannot handle uncertainties and surprises not anticipated by the designer — the systems are “brittle” (Holland 1986). To escape from this brittleness, one thing that must be done is to switch from the classical theory to a new theory that captures the plasticity and fluidity of concepts (Hofstadter & FARG 1995).

This example tells us, if certain AI research continues to ignore CogSci, it can still go on as before, but it will not be easy to solve the really hard problems. The mechanisms of the human mind are not formed arbitrarily, but are adapted through a long evolution process in the attempt to serve certain purposes. If AI also attempts to achieve similar goals, it should consider these mechanisms carefully, though they are not the only possibilities. Failing to do so will leave AI to its old path, where successes can be achieved, but they look more like the successes of computer science, while the key features of intelligence, like flexibility, creativity, and efficiency, will still be missing.

At the same time, this is also a challenge to CogSci. To make its results applicable to AI, CogSci should try to generalize its conclusions to such a level that it is not bounded to the details of the *human* mind/brain. This is possible, because many cognitive mechanisms not only have a neurological explanation, but also have a functional explanation, as a way to achieve certain goals under certain conditions. As soon as the (functional) preconditions and effects of a cognitive mechanism are clarified, its relevance and applicability to AI become clear.

2. Are “good enough” solutions valuable?

Yes, in a sense this is what AI is all about — to find “good enough” solutions when the optimal ones cannot be obtained, due to the restriction of knowledge and resources.

Of course, whenever optimal solutions can be obtained, they are preferred over “good enough” solutions for all practical purposes, except that the research then does not belong to AI anymore.

One major problem in mainstream AI research is that people only pay attention to the *quality* of solutions a system can produce, but not to the *condition* under which the solutions are produced. It can be argued that what we call “intelligence” is not a capability that always finds optimal solutions for all problems (at least the human mind does not have such a capability), but a capability that usually finds good enough solutions under the restriction of available knowledge and resources (Medin & Ross 1992; Wang 1995; 2006b).

When we say that the AI problems are “hard”, we surely do not mean that the problems in other fields of computer science are easy. The difficulty of the AI problems comes from the fact that they cannot be properly handled by the traditional theories. The theoretical foundation of computer science consists of mathematical logic, formal semantics, discrete mathematics, theory of computability and computational complexity, etc., which all come from the study of mathematics. In mathematics, we are usually only interested in optimal solutions. On the contrary, AI problems are hard, often because the system has to answer questions for which it has insufficient knowledge and resources, and the traditional theories tell us little about what to do in this situation. Otherwise, the problem will be “easy”, because “sufficient knowledge” means that the system has an algorithm for the optimal solution, and “sufficient resources” means that the system can afford the time-space expense of the algorithm. In this situation the problem is considered as solved in computer science.

Therefore, AI can be seen as a field where the problems are “hard” in the sense that they are unsolvable or intractable in computer science. By definition, the solutions found in AI are not optimal (otherwise the problems are not AI problems at all), but still “good enough”, especially compared with random guesses. Contrary to a common misconception, a solution found by an AI system does not have higher quality than a solution found by a traditional computing system. Actually it is the opposite: the latter is always better than the former, and the former is valuable only when the latter cannot be obtained. AI should not pursue optimal solutions, because that domain has been covered by traditional theories quite well, and the real challenge is outside that coverage.

This situation will not be changed by the progress of hardware. It is true that with the coming of larger, faster, and cheaper computers, many previously intractable problems have become tractable, but even so, there are still, and will always be, many problems for which a system has insufficient knowledge and resources, and it cannot wait for the collection of more information or the invention of new machines.

3. Are multilevel heterogeneous approaches beneficial?

“Multilevel heterogeneous approaches” may be beneficial for certain purpose, but it is not the best way to achieve AI.

Since “intelligence” in essence is a unified mechanism, at the core an AI system should depend on a single approach to achieve coherence, efficiency, and flexibility. However, this core should provide a platform on which various other heterogeneous approaches can be plugged in as parts and applied whenever suitable.

It is always true that for a given problem, different approaches usually have different advantages, and no single approach is the best in all aspects. However, it does not mean that the best solution can be provided by bundling these heterogeneous parts together to get a hybrid system. Though it is not hard technically to pass data among these parts, it is difficult to do so in a consistent and efficient way, because these approaches are usually based on different assumptions.

To use heterogeneous approaches in the same system, it is important to integrate them around an intelligent core (Wang 2004; Nilsson 2005). The core does not need to directly provide best solutions for all problems, but it should be general, compact, consistent, flexible, and efficient. It should be able to represent various other parts as (hardware or software) tools, with their preconditions and expected effects. The core will then decide which tool to apply to a given problem, invoke the tool, monitor the process, and check the result. When there is no proper tool available, the core should be able to directly deal with the problem, and to get a good enough result.

The relation between such a core and the heterogeneous parts is similar to the relation between a pair of hands and various tools and machines. For any given problem, it is always possible to design and build a tool or a machine that solves the problem better than bare hands, but it does not mean that the hands can be replaced by bundling the tools and machines together.

The request for a single core does not contradict with the acceptance of heterogeneous tools, because they are not at the same level. The situation is also similar to the relation between an operating system and various application programs. Though different operating systems obviously have their advantages and disadvantages, to run more than one of them in the same machine at the same time can be a disaster. Intelligence is mainly displayed at the “system software” level, where general consistency is more important than performance on a concrete problem.

Though it also supports heterogeneous tools, this “core-based” approach is fundamentally different from what people usually mean by “multilevel heterogeneous approaches”. It does not try to directly integrate the existing AI techniques. Instead, it tries to build an intelligent core which is different from any existing AI technique, because of the condition under which it works. It is only when this core is mostly finished, the other approaches should be taken into account, and they should be handled in a “plug-and-play” manner, that is, not integrated into the core as necessary components for the system to work.

4. Is adaptiveness an essential component of intelligence?

Yes, adaptiveness is one (though not the only one) defining feature of intelligence. As Piaget put it, intelligence is a “highly developed form of mental adaptation” (Piaget 1960).

As argued previously, AI problems are hard because the system has insufficient knowledge and resources for them. However, just to acknowledge such insufficiency does not make the system intelligent. The system should attempt to solve these problems as well as possible, with the available knowledge and resources. “Insufficient knowledge” means the system does not know the optimal solution for a problem. For example, if the future cannot be assumed to be identical to the past, the previous knowledge cannot guarantee a correct prediction. This is Hume’s problem of induction (Hume 1748). In this situation, how can one solution be judged as better than another one?

In an adaptive system, though no prediction can be proved to be infallible, the system prefers conclusions that are better supported by its experience. With the coming of new evidence from time to time, the system adapts to its environment by learning new regularities in its observations, adjusting its beliefs according to these regularities, and predicting the future according to these beliefs. In this way, the validity of induction is justified as *adapting to experience*, not as *preserving truth* (as in deductive logic).

However, it does not mean that adaptiveness guarantees the best solution for a given problem. Actually it is usually the contrary, that is, the best solution is obtained through careful analysis of the problem, given sufficient knowledge and resources. But that is not AI anymore. Adaptation, or learning, is valuable only when the system lacks certain knowledge. We learn something only when we did not know it previously, and we do so only when we have to.

It is important to distinguish what a system *can do* from what a system *can learn to do*. These two capabilities are not always correlated in a system. To judge how intelligent a system is, it is the latter, not the former, that matters. We have the feeling that conventional computer systems are not intelligent, not because it cannot do anything (on the contrary, they can do many things perfectly), but because they cannot learn at all, or cannot learn much.

An implication of the above conclusion is that to measure the level of intelligence of a computer system, it is not a good idea to use a fixed set of concrete problems, such as certain “IQ test”. Such tests are already controversial when used among human beings, and the results will be even less meaningful when used in AI. Since human infants have similar innate capabilities, it makes some sense to identify their learning ability with their problem-solving ability at a certain age. For AI systems, however, these two abilities have no necessary correlation at all. As soon as a set of testing problem is determined, it is possible to design a computer system that works well on these problems, but cannot do anything else. To really measure intelligence, we should not check a system’s ability of problem-solving at a given moment, but how much this ability can change over time, that is, the scope and speed of adaptation.

5. Are the most efficient solutions problem-specific?

Though the most efficient solutions are almost always problem-specific, to find such solutions should not be taken as the goal of AI.

Problem-specific solutions, though efficient for the target problems, usually cannot be used in other situations. Therefore, a system equipped only with problem-specific solutions cannot handle novel problems, even when they are similar to the old problems. Since a central issue of AI is to deal with novel situations, to use problem-specific solutions will not work.

Some people believe that if AI has covered enough problem-specific solutions, general-purpose solutions can be derived from them. This is not the case, because problem-specific solutions are usually formulated in such a way that the general principles and the domain-specific tricks are tangled together, and to separate them is very difficult, if not impossible. On the other hand, general-purpose approaches start at a set of assumptions that rule out the possibility for domain-specific tricks to be smuggled in, and the design will therefore be fundamentally different from problem-specific ones.

By “general-purpose solutions”, I do not mean ideas like “General Problem Solver” (Newell & Simon 1963), that attempt to directly treat all problems in a uniform manner. Instead, I mean a system that is designed to be general purpose, without any assumption about the problem domain, and can learn whatever knowledge embedded in its experience. Since the system is adaptive, it will be inevitable for it to become domain-specific to a certain extent. If its experience comes from a certain domain, then its beliefs and goals will adapt to the environment of that domain.

This situation is similar to the situation of the human mind. Nobody is born to be an expert of a certain domain. Expertise comes from the learning process after birth. In this way, a general-purpose system can also obtain efficient problem-specific solutions. Differently from the “hard-wired” problem-specific solutions, this “learned expertise” allows more flexibility of being revised and applied to novel-but-similar situations.

In this way, the general-purpose knowledge and problem-specific knowledge co-exist in different levels in a system. Take the intelligent reasoning system I designed and developed as an example (Wang 1995; 2006b). In this system, the “innate” components, including the knowledge representation language, the inference rules, and the control mechanism, are all domain-independent. On the other hand, the experience of the system, which provides all the tasks and beliefs, is domain-specific. The system can learn problem-specific solutions, though they are not coded in the system design.

In summary, what I suggest is to build a (general-purpose) child, and let it grow into a (special-purpose) adult, which is an idea that has been proposed many times before (Turing 1950; Nilsson 2005), but has not been sufficiently explored in engineering practice.

6. Are specialized, modular components a reasonable approach to general intelligence?

Yes, modular components are needed for AI, but only in a certain sense, and to a small extent, compared to the current common belief in the field.

At the boundary of a general-purpose intelligence system, it is reasonable to identify certain modular components, each of which is developed to carry out a special function, such as sensorimotor, language interface, etc. At the core of the system, however, all the major processes should be carried out by a unified mechanism.

Many people believe that AI should be studied in a “divide and conquer” fashion, and they see intelligence as consisting of modular components and functions, like searching, recognizing, categorizing, reasoning, planning, decision making, problem solving, learning, and so on.

Though this opinion is intuitively attractive, it has certain weaknesses.

One problem is, when specified in isolation, a formalized function is often quite different from its “natural form” in the human mind. For example, to study analogy without perception leads to distorted cognitive models (Chalmers, French, & Hofstadter 1992).

Furthermore, even if we can successfully produce all the desired functions, it does not mean that we can easily integrate them into one consistent system, because different functions may be developed under fundamentally different assumptions. According to the past experience in building integrated AI systems, “Component development is crucial; connecting the components is more crucial” (Roland & Shiman 2002).

A common counterargument runs something like this: “Intelligence is very complex, so we have to start from a single function to make the study tractable.” For many systems with weak internal connections, this is often a good choice, but for a system like the mind, the situation may be just the opposite.

When the so-called “functions” are actually phenomena produced by a complex-but-unified mechanism, reproducing all of them together (by duplicating the mechanism) may be actually simpler than reproducing only one of them in isolation. For example, we can grow a tree, but we still cannot generate a leaf *alone*, although a leaf is just a small part of a tree.

There are reasons to believe that intelligence is such a complex-but-unified phenomenon. As Piaget said: “Intelligence in action is, in effect, irreducible to everything that is not itself and, moreover, it appears as a total system of which one cannot conceive one part without bringing in all of it.” (Piaget 1963)

My AI system is developed according to the belief that intelligence is basically a single mechanism, and the various “functions” are different phenomena when the observers focus on different aspects of the mechanism. By setting up a proper foundation, the system has obtained unified functions like reasoning, learning, categorizing, planning, decision making, and problem solving. For detailed descriptions, see (Wang 2004; 2006b).

7. Can artificial intelligence contribute to our understanding of human cognition?

Yes. At a general level, AI should follow the same principles as human cognition. It can contribute to CogSci in several ways.

First, AI provides a testing ground for various CogSci theories. Many theories on cognition, intelligence, and thinking have been proposed in philosophy and psychology, and it is not always easy to check them against the reality of the human mind. Because an AI model eventually needs to be implemented in a computer system, it provides an accurate way to test a theory. We will learn if a theory has enough details to be valuable, and if it indeed reproduces the process it describes.

For example, previously I mentioned the various theories on categorization that originated in psychology. If they are implemented in separate AI systems, we can observe the behavior of the systems, so as to evaluate the strength and weakness of each theory.

Also, AI can help CogSci to distinguish human specific features from common features of all cognitive systems (human, machine, animal, and so on). In the current CogSci research, people tend to ignore this distinction, and as a result, observations in the human mind often fail to be properly generalized.

For example, we all agree that (some kind of) memory is necessary for cognition, so the system can remember things. However, whether “forgetting” is also necessary for cognition is still an open problem. We often hear people say that “never forgetting things” is an advantage computers have over the human mind, but if a computer must work with insufficient storage space and processing time, it must also delete things from time to time to make space for new information, as well as to make certain things more accessible (that means to make other things less accessible). Therefore, forgetting becomes necessary. Of course, the concrete forgetting rate of an AI system does not have to be identical to that of an average human — at least we have not seen a strong reason for that.

When AI conclusions are applied to CogSci, we do need to keep in mind about the differences among the fields. For example, AI is mainly interested in *normative* models (How should a computer behave?), while psychology is mainly interested in *descriptive* models (How does the human mind actually behave?). On the other hand, the two still have close relationship, because the basic assumptions behind the normative models should come from general descriptions of the mind.

For example, in the current study of “human errors”, psychologists often take certain normative theory, such as first-order predicate calculus or probability theory, as the criterion to judge the correctness of human behavior. From the viewpoint of AI, we know that every normative theory has its assumptions and applicable conditions, and beyond them the theory is no longer “normative” at all. In this way, we can find problems in psychological research, even though we cannot use AI research to directly replace it (Wang 1995; 2006b).

Conclusions

The seven focus questions can be divided into three groups:

- The identity of AI (Question 2 and 4). Here my opinion is to take AI to mean “an adaptive system that provides good enough solutions”.
- The relation between AI and CogSci (Question 1 and 7). Here my opinion is to let AI move closer to CogSci, and away from Computer Science.
- The path toward AI (Question 3, 5, and 6). Here my opinion is to start at a unified, general-purpose core system, which can learn domain-specific knowledge and use problem-specific tools.

Though most of the current AI works qualify as sound scientific research and valuable engineering practice, they fail to address the key issues in intelligence and cognition, which distinguish the human mind from traditional computer systems.

To solve this problem, it is necessary to clarify the goal of AI, and to realize that it is not to find perfect solutions to various domain-specific problems, but to develop a general-purpose mechanism to deal with problems beyond the current knowledge/resources scope of the system by learning from experience.

In the future, we do hope to see a closer cooperation of AI and CogSci that eventually leads to a unified theory of intelligence and cognition, which is a direction that has been explored only by a few researchers (Newell 1990; Anderson & Lebiere 1998).

In recent years, there is a noticeable rebound of hope and interest in general-purpose AI systems, under the names of “Artificial General Intelligence” (Goertzel & Pennachin 2006) or “Human-level Artificial Intelligence” (Nilsson 2005). This time, as intelligent beings, we should be able to learn from the past experience, and do a better job (within our knowledge-resources restriction, of course).

References

- Anderson, J., and Lebiere, C. 1998. *The Atomic Components of Thought*. Mahwah, New Jersey: Lawrence Erlbaum Associates.
- Chalmers, D.; French, R.; and Hofstadter, D. 1992. High-level perception, representation, and analogy: a critique of artificial intelligence methodology. *Journal of Experimental and Theoretical Artificial Intelligence* 4:185–211.
- Goertzel, B., and Pennachin, C., eds. 2006. *Artificial General Intelligence*. New York: Springer.
- Goldstone, R. 2005. Returning to a new home. *Cognitive Science* 29:1–4.
- Hearst, M., and Hirsh, H. 2000. AI’s greatest trends and controversies. *IEEE Intelligent Systems* 8–17.
- Hofstadter, D., and FARG. 1995. *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. New York: Basic Books.
- Holland, J. 1986. Escaping brittleness: the possibilities of general purpose learning algorithms applied to parallel rule-based systems. In Michalski, R.; Carbonell, J.; and Mitchell, T., eds., *Machine Learning: an artificial intelligence approach*, volume II. Los Altos, California: Morgan Kaufmann. chapter 20, 593–624.
- Hume, D. 1748. *An Enquiry Concerning Human Understanding*. London.
- Laurence, S., and Margolis, E. 1999. Concepts and cognitive science. In Margolis, E., and Laurence, S., eds., *Concepts: Core Readings*. Cambridge, Massachusetts: MIT Press. 3–81.
- McCorduck, P. 1979. *Machines Who Think*. New York: W. H. Freeman & Co.
- Medin, D., and Ross, B. 1992. *Cognitive Psychology*. Fort Worth: Harcourt Brace Jovanovich.
- Newell, A., and Simon, H. 1963. GPS, a program that simulates human thought. In Feigenbaum, E., and Feldman, J., eds., *Computers and Thought*. McGraw-Hill, New York. 279–293.
- Newell, A. 1990. *Unified Theories of Cognition*. Cambridge, Massachusetts: Harvard University Press.
- Nilsson, N. 2005. Human-level artificial intelligence? Be serious! *AI Magazine* 26(4):68–75.
- Piaget, J. 1960. *The Psychology of Intelligence*. Paterson, New Jersey: Littlefield, Adams & Co.
- Piaget, J. 1963. *The Origins of Intelligence in Children*. New York: W.W. Norton & Company, Inc. Translated by M. Cook.
- Roland, A., and Shiman, P. 2002. *Strategic computing : DARPA and the quest for machine intelligence, 1983-1993*. Cambridge, Massachusetts: MIT Press.
- Stone, M.; Hirsh, H.; and others. 2005. Artificial intelligence: The next twenty-five years. *AI Magazine* 26(4):85–97.
- Stork, D. 1997. Scientist on the set: An interview with Marvin Minsky. In Stork, D., ed., *HAL’s Legacy: 2001’s Computer as Dream and Reality*. Cambridge, Massachusetts: MIT Press. 15–30.
- Turing, A. 1950. Computing machinery and intelligence. *Mind* LIX:433–460.
- Wang, P. 1995. *Non-Axiomatic Reasoning System: Exploring the Essence of Intelligence*. Ph.D. Dissertation, Indiana University.
- Wang, P. 2004. Toward a unified artificial intelligence. In *Papers from the 2004 AAAI Fall Symposium on Achieving Human-Level Intelligence through Integrated Research and Systems*, 83–90.
- Wang, P. 2006a. The logic of intelligence. In Goertzel, B., and Pennachin, C., eds., *Artificial General Intelligence*. New York: Springer.
- Wang, P. 2006b. Rigid Flexibility — The Logic of Intelligence. Manuscript under review.
- Wilson, R., and Keil, F., eds. 1999. *The MIT Encyclopedia of the Cognitive Sciences*. Cambridge, Massachusetts: MIT Press.