

Insufficient Knowledge and Resources — A Biological Constraint and Its Functional Implications

Pei Wang

Temple University, Philadelphia, USA

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

Abstract

Insufficient knowledge and resources is not only a biological constraint on human and animal intelligence, but also has important functional implications for artificial intelligence (AI) systems. Traditional theories dominating AI research typically assume some kind of sufficiency of knowledge and resources, so cannot solve many problems in the field. AI needs new theories obeying this constraint, which cannot be obtained by minor revisions or extensions of the traditional theories. The practice of NARS, an AI project, shows that such new theories are feasible and promising in providing a new theoretical foundation for AI.

AI and Biological Constraints

Since its inception, Artificial Intelligence (AI) has been driven by two contrary intuitions, both hinted in the name of the field.

On one hand, the only form of “intelligence” we know for sure comes from biological systems. Therefore, it seems natural for AI research to be biologically inspired, and for AI systems to be developed under biological constraints. After all, there is existing proof that this approach works. Though many researchers prefer the more abstract languages provided by neuroscience or psychology over the language of biology, their work can still be considered as “biologically inspired”, in a broad sense. To some researchers, “This is what the brain does” is a good justification for a design decision in an AI system (Reeke and Edelman 1988).

On the other hand, the “artificial” in the name refers to computers, which are not biological in hardware, and so do not necessarily follow biological principles. Therefore, it also seems natural for AI to be considered as a branch of computer science, and for the systems to be designed under computational considerations (which usually ignore biological constraints). After all, the biological form of intelligence should only be one of the possible forms, otherwise AI is doomed to fail. Some researchers are afraid that to follow biology (or cognitive science) too closely may unnecessarily limit our imagination on how an intelligent system can be built, and the “biological way” is often not the best way for computers to solve problems (Korf 2006).

Copyright © 2009, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

As usual, the result of such an intuition conflict is a compromise between the two. Many AI researchers get inspirations and constraints from biology, neuroscience, and psychology, on how intelligence is produced in biological systems. Even so, these ideas are introduced into AI systems only when they have clear *functional implications* for AI. Therefore, being “biologically inspired/constrained” is not an ends by itself, but a means to an ends, which is non-biological by nature. In AI research, a biological inspiration should not be directly used as a justification for a design decision. After all, there are many biological principles that seem to be irrelevant to intelligence.

So, the real issue here is not at the extreme positions — nobody has seriously argued that AI should not be biologically *inspired*, nor that *all* biological constraints must be obeyed in AI. Instead, it is on how to posit AI with respect to biology (and neuroscience, psychology, etc), and concretely, on which biologically inspirations are really fruitful in AI research and which biological constraints are functionally necessary in AI design.

To solve the above problem means that, for any concrete biological inspiration or constraint, it is not enough to show that it is indeed present in the human brain or other biological systems. For it to be relevant to AI, its functional necessity should also be established, that is, not only that it has certain desired consequences, but also that the same effect cannot be better achieved in an AI system in another way.

AIKR as a constraint

To make the discussion more concrete, assume we are interested in *problem-solving systems*, either biological or not. Such a system obtains knowledge and problems from its environment, and solves the problems according to the knowledge. The problem-solving process costs certain resources, mainly processing time and storage space.

Not all such systems are considered as “intelligent”. The main conclusion to be advocated in this paper is the **Assumption of Insufficient Knowledge and Resources (AIKR)**: *An intelligent system should be able to solve problems with insufficient knowledge and resources* (Wang 1993; 2006).

In this context, “insufficient knowledge and resources” is further specified by the following three properties:

Finite: The system's hardware includes a constant number of processors (each with a fixed maximum processing speed), and a constant amount of memory space.

Real-time: New knowledge and problems may come to the system at any moment, and a problem usually has time requirement for its solution.

Open: The system is open to knowledge and problem of any content, as far as they can be represented in the format acceptable by the system.

For human beings and animals showing intelligence, "problem solving" is an abstract way to describe how such a system interacts with its environment to maintain or to achieve its goals of survival and reproduction, as well as the other goals derived from them. The system's *knowledge* is its internal connections, either innate or acquired from its experience, that link its internal drives and external sensations to its actions and reactions. For such a system, the insufficiency of knowledge and resources is a biological constraint by nature. The system has *finite* information-processing capability, though has to *open* to novel situations in its environment (which is not under its control), and to respond to them in *real time*. Otherwise it cannot survive.

When we say such a system "is able to solve problems with insufficient knowledge and resources", it does not mean that it can *always* find the *best* solution for *every* problem, nor even that it can find any solution for each problem. Instead, we mean such a system can work in such a situation and solve some of the problems to its satisfaction, even though it may fail to solve some other problems, and some of its "solutions" may turn out to be wrong. Actually, if the system can solve all problems perfectly, by definition it is not working under AIKR. Therefore, AIKR not only specifies the working environment of a system, but also excludes the possibility of certain expectations for a system working in such an environment.

Why AIKR is usually disobeyed in AI

Though AIKR has a biological origin, it does not mean that it cannot be applied to artificial systems — there is nothing fundamentally "biological" in the above three requirements. Any concrete computer system has finite processing and storage capacity, and is often desired to be open and to work in real-time. Even so, few AI systems has been designed to fully obey AIKR.

Though every concrete computer system is finite, this constraint is ignored by many theoretical models of computation, such as a Turing Machine (Hopcroft and Ullman 1979). For instance, many systems are designed to assume that additional memory is always available whenever needed, and to leave the problem to the human administrator when this is not the case.

Most computer systems do not work in real time. Usually, a problem does not show up with a time requirement associated. Instead, the processing time of a problem is determined by the algorithm of the system by which the problem is solved, plus the hardware/software platform on which the algorithm is implemented. For the same problem instance, the system spends the same amount of

(time and space) resources on it, independent of the context in which the processing happens. This is true even for the so-called "real-time systems" (Laffey et al. 1988; Strosnider and Paul 1994), which are *designed* to satisfy the time requirement of a practical application, though when the system actually interacts with its environment, it often does not respond to the (change of) time requirement in the current context. Furthermore, these systems usually only handle time requirements in the form of deadlines, though in reality a time requirement may take another form, such as "as soon as possible", with various degrees of urgency.

Most computer systems only accept problems with content anticipated by the system designer, and novel problems beyond that will be simply rejected, or mistreated — this is why these systems are considered "brittle" (Holland 1986). There are also various restrictions on the content of knowledge the system can accept. For example, most reasoning systems cannot tolerate any contradiction in the given knowledge, and systems based on probability theory assume the given data correspond to a consistent probability distribution function. When such an assumption is violated, the system cannot function properly.

In summary, even though a computer system always has *limited* knowledge and resources, they can still be *sufficient* with respect to the problems the system attempts to solve, since the system can simply ignore the problems beyond its limitation. Intuitively, this is why people consider conventional computer systems unintelligent — though such a system can be very capable when solving certain problems, it is rigid and brittle when facing unanticipated situations.

This issue is not unnoticed. On the contrary, in a sense most AI research projects aim at problems for which there are insufficient knowledge and resources. Fields like Machine Learning and Reasoning under Uncertainty explicitly focus on the problems caused by imperfect knowledge and inexhaustible possibilities, which are also the causes of many difficulties in fields like vision, robotics, and natural language processing. Even so, there are few systems developed under AIKR. Instead, an AI system typically acknowledges AIKR on certain aspects, while on the others still assume the sufficiency of knowledge and resources.

Such an attitude usually comes from two considerations: *idealization* and *simplification*.

To some researchers working on normative models of intelligence, since AIKR is a *constraint*, it limits the model's capability and is therefore undesired, so it should not be allowed whenever possible. To them, in a theory or a formal model of intelligence, *idealization* should be used to exclude the messy details of the reality. After all, Turing Machine is highly unrealistic, but still plays a key role in theoretical computer science. In this way, at least an idealized model of intelligence can be specified, so as to establish a goal for the concrete systems to *approach* (if not to *reach*). An example of this opinion can be found in (Legg and Hutter 2007): "We consider the addition of resource limitations to the definition of intelligence to be either superfluous, or wrong. . . . Normally we do not judge the intelligence of something relative to the resources it uses." Based on such a belief on intelligence, Hutter proposes the AIXI model, which pro-

vides optimal solutions to a certain type of problem, while requiring unlimited amount of resources (Hutter 2005).

To many other researchers who are building concrete intelligent systems, AIKR is a reality they have to face. However, they consider it as a collection of problems, which is better to be handled in a divide-and-conquer manner. This consideration of *simplification* lead them to focus on one (or a few) of the issues, while to leave the the other issues to a future time or to someone else. They hope in this way the problems will become easier to solve, both in theory and in practice. We can find this attitude in most AI projects.

In the following, we will argue against both above attitudes toward AIKR.

Destructive implications of AIKR

Normative models of intelligence study what an intelligent (or rational) system *should do*. In current AI research, these models are usually based on the following traditional theories:

First-order predicate calculus, which specifies valid inference rules on binary propositions;

Probability theory, which specifies valid inference rules on random events and beliefs;

Model-theoretic semantics, which defines “meaning” and “truth” by mapping a language into a model;

Computability and complexity theory, which analyzes problem-solving processes that follow algorithms.

None of these theories, in its standard form, obeys AIKR. On the contrary, they all assume the sufficiency of knowledge and resources in certain form. If we analyze the difficulties they run into in AI, we will see that most of the issues come from here.

Though first-order predicate calculus successfully plays a fundamental role in mathematics and computer science (Halpern et al. 2001), the “logician AI” school (McCarthy 1988; Nilsson 1991) has been widely criticized for its rigidity (McDermott 1987; Birnbaum 1991). In (Wang 2004a), the problems in the logical approach towards AI is clustered into three groups:

- the uncertainties in knowledge and inference process,
- the justification of non-deductive inference rules,
- the counterintuitive conclusions produced by logic.

It is argued in (Wang 2004a) that all these problems come from the usage of a mathematical logic (which does not obey AIKR) in a non-mathematical domain (where AIKR becomes necessary).

Probability theory allows uncertainty in predictions and beliefs, but its application in AI typically assumes that all kinds of uncertainty can be captured by a consistent probability distribution defined over a propositional space (Pearl 1988). Consequently, it lacks a general mechanism to handle belief revisions caused by changes in background knowledge, which is not in the predetermined propositional space (Wang 2004b).

Model-theoretic semantics treats *meaning* as denoted object in the model, and *truth-value* as distance between a

statement and a fact in the model (Tarski 1944; Barwise and Etchemendy 1989), and therefore they are independent of the system’s experience and activities. Consequently, such a semantics cannot capture the experience-dependency and context-sensitivity in meaning and truth-value, which are needed for an adaptive system (Lakoff 1988; Wang 2005).

Algorithmic problem-solving has the advantage of predictability and repeatability, but lacks the flexibility and scalability required by intelligence. Again, here the problem comes from the assumption of sufficient knowledge (i.e., there is a known algorithm) and resources (i.e., the system can afford the resources required by the algorithm) (Wang 2009).

Therefore, to design a system under AIKR, new theories are necessary, which are fundamentally different from the traditional theories, though may still be similar to them here or there.

The above conclusion does not mean that the traditional theories are “wrong”. They are still normative theories about a certain type of rationality. It is just that under AIKR, a different type of rationality is needed, which in turn requests different normative theories.

For a system based on AIKR, many requirements of classic rationality cannot be achieved anymore. The system can no longer know “objective truth”, or always make correct predictions. Nor can it have guaranteed completeness, consistency, predictability, repeatability, etc. On the other hand, it does not mean everything goes. Under AIKR, a *rational* system is one that attempts to *adapt* to its environment, though its attempts may fail. While the future is unpredictable, it is better to behave according to the past; while the resource supply cannot meet the demand, it is better to use them in the most efficient manner, according to the system’s estimation. Consequently, a system designed under AIKR can still be “rational”, though not in the same sense as specified by classic “rationality”.

Similar opinions have been expressed previously, with notions like “bounded rationality” (Simon 1957), “Type II rationality” (Good 1983), “minimal rationality” (Cherniak 1986), “general principle of rationality” (Anderson 1990), “limited rationality” (Russell and Wefald 1991), and so on. Each of these notions assumes a less ideal and more realistic working environment, then specifies the principles that lead to the mostly desired results. AIKR is similar to the above notions in spirit, though it is more radical than all of them — as explained previously, “insufficient” does not merely mean “bounded” or “limited”.

Therefore, a main implication of AIKR is a proposed paradigm shift in AI research, by suggesting the systems to be designed according to different assumptions on the environment, different working principles, and different evaluation criteria. According to this opinion, the traditional theories are not applicable to AI systems in general (though still useful here or there), since their fundamental assumptions are no longer satisfied in this domain.

Now let us revisit the two objections to AIKR mentioned previously.

The objection based on idealization is wrong, because this opinion, as explicitly expressed in (Legg and Hutter

2007), fails to see that there are different types of idealization. AIKR is also an idealization, except that it is closer to the reality of intelligent systems. Rationality (or intelligence) can be defined with either finite resources or infinite resources. However, the two definitions lead to very different design decisions, so the “intelligence” defined in (Legg and Hutter 2007) does not provide proper guidance to the design of concrete AI systems which must work under AIKR. Here the situation is different from that of Turing Machine, because resource restriction is a fundamental constraint for *intelligent* systems, though not for *computational* systems. AIKR is a major reason for the cognitive facilities to be what they are, as observed by psychologists: “Much of intelligent behavior can be understood in terms of strategies for coping with too little information and too many possibilities” (Medin and Ross 1992). No wonder AIXI is so different from the human mind, either in structure or in process (Hutter 2005).

The objection based on simplification is wrong, because this opinion leads to a mixture of different notions of rationality — in certain parts of a system, a new notion of rationality is introduced, while in the others, the classic rationality is still followed. Though such a system can be used for certain limited purposes, they have serious theoretical defects when the research goal is to develop a theory and model of intelligence that is general-purpose and comparable to human intelligence. The divide-and-conquer strategy does not work well here, because a change on the assumption on knowledge and resources usually completely changes the nature of the problem. For example, an open system needs to be able to revise its beliefs according to new evidence. It is an important aspect of common-sense reasoning that non-monotonic logic attempts to capture, by allowing the system to derive revisable conclusions using “default rules” (McCarthy 1988; Reiter 1987). However, the default rule, though having empirical contents, cannot be revised by new experience, largely because non-monotonic logic is still binary, so cannot measure the *degree* of evidential support a default rule gets. Consequently, the system is not fully adaptive, and has no general way to handle the inconsistent conclusions derived from different default rules (Wang 1995b). If we try to extend the logic from binary to multi-valued, we will see that the situation is changed so much that the previous results in binary logic become hardly relevant. Similarly, it is almost impossible to revise AIXI into a system working under AIKR, because the resulting model (if it can be done) will bear little resemblance to the original AIXI.

Constructive implications of AIKR

AIKR is not only a constraint that *should* be obeyed in AI system, but also one that *can* be obeyed. It is possible to design an AI system which is finite, real-time, and open.

The above claim is supported by the practice of the NARS (Non-Axiomatic Reasoning System) project, which is an AI system designed in the framework of reasoning system, and with prototype systems implemented (Wang 1993; 1995a; 2006).

NARS is designed to fully obey AIKR. It is “non-axiomatic” in the sense that its domain knowledge all comes

from its experience (i.e., input knowledge), and are revisable by new experience. The problems the system faces are often beyond its current knowledge scope, so no absolutely correct or optimal conclusion can be guaranteed. Instead, as an adaptive system, NARS uses its past experience to predict the future, and uses its finite computational resource supply to satisfy its excessive resource demands. In this situation, a “rational” solution to a problem is not one with a guaranteed quality, but one that is mostly consistent with the system’s experience and can be obtained with the available resources supply.

Under AIKR, model-theoretic semantics does not work, since meaning of terms and truth-value of beliefs in NARS cannot be taken to be constants corresponding to the objects and facts in an outside world. Instead, truth-value has to mean the extent to which the system takes a belief to be true, according to the evidence in the system’s experience; meaning has to indicate what a term expresses to the system, that is, the role it plays in the system’s experience. Such an *experience-grounded semantics* is fundamentally different from model-theoretic semantics, and cannot be obtained by extending or partially revising the latter.

When the beliefs and problems of NARS are represented in a formal language, what is represented is not a “model of world” that is independent of the system’s experience, but a “summary of experience” that is dependent of the system’s body, environment, and the history of interaction between the system and the environment. This summary must be expendable to cover novel situations, which is usually not identical to any past situation known to the system. For this purpose, a conceptual vocabulary is needed, in which each concept indicates a certain stable pattern in the experience, and a belief specifies the substitutability between the concepts, that is, whether (or to what extent) one concept can be used as another. To represent and derive this kind of beliefs in NARS, a *term logic* is used, which is fundamentally different from first-order predicate logic. In representation language, a predicate logic uses a *predicate-arguments* format, while a term logic uses a *subject-copula-predicate* format. In inference rules, a predicate logic uses *truth-functional* rules, and reasons purely according to the truth-value relationship between the premises and the conclusion, while a term logic uses *sylogistic* rules, and reasons mainly according to the transitivity of the copula.

Since NARS is open to experience of any content (as long as it can be represented in the system’s representation language), the truth-value of a belief is uncertain, since it may have both positive and negative evidence, and the impact of future evidence will need to be considered, too. Obviously, a binary truth-value is no longer suitable, but even replacing it with a probability value is not enough, because under AIKR, the system cannot have a consistent probability distribution function on its belief space to summarize all of its experience. Instead, in NARS each belief has two numbers attached to indicate its evidential support: a *frequency* value (representing the proportion of positive evidence among all evidence) and a *confidence* value (representing the proportion of current evidence among evidence at a near future, after the coming of new evidence of a certain amount). To-

gether, these two numbers provide information similar to a probability value, though they are only based on past experience, and can be revised by new evidence. Furthermore, it is not assumed that all the truth-values in the system at the same time are consistent. Whenever an inconsistency happens, the system resolves it, either by merging the involved evidence pools (if they are distinct), or by picking the conclusion with more evidence behind it (if the evidence cannot be simply merged). In each step of inference, the rule treats the premises as the only evidence to be considered, and assign a truth-value to the conclusion accordingly. Some of the inference rules have certain similarity with the rules in traditional logic or probability theory, but in general the inference rules of NARS cannot be fully derived from any traditional theory.

Since NARS is open to novel problems, and must process them in real time, with various time requirements and changing context, it cannot follow predetermined algorithm for each problem. Instead, the system usually solves the problems in a case-by-case manner. For each concrete occurrence of a problem, the system processes it according to available knowledge and using available resources. Under AIKR, the system cannot assume that it knows everything about the problem, nor that it can take all available knowledge into consideration. NARS maintains priority distributions among its problems and beliefs, and in each step, the items with higher priority will have a higher chance to be accessed. The priority distributions are constantly adjusted by the system itself, according to the change in the environment and the feedback collected for its operations, so as to achieve the highest overall (expected) resource-efficiency. Consequently, the system's problem-solving processes become flexible, creative, context-sensitive, and non-algorithmic. These processes cannot be analyzed according to the theory of computability and complexity, because even for the same problem instance, the system's processing path and expense may change from situation to situation.

NARS fully obeys AIKR. The system is *finite*, not only because it is implemented in a computer system with fixed processing capacity, but also because it manages its own resources without outside interference — whenever some processes are started or speed up, usually there are others stopped or slowed down; whenever new items are added into memory, usually there are old items removed from it. The system is *open*, because it can be given any beliefs and problems, as far as they can be expressed in the representation language of the system, even when the beliefs conflict with each other or the problems go beyond the system's current capability. The system works in *real-time*, because new input can show up at any moment, and the problems can have various time requirements attached to them — some may need immediate attention, though have little long-term importance, while some others may be the opposite.

NARS is built in stages. As described in (Wang 2006), the representation language and inference rules are extended incrementally, with more expressive and inferential power added in each stage. Even so, AIKR is established from the very beginning and followed all the way. This type of divide-and-conquer is different from the common type criti-

cized previously. According to common practice in AI, even if AIKR were accepted as necessary, people would build a finite system first, then make it open, and finally make it to work in real time. However, in that way very likely the design of an earlier stage will be mostly useless for a later stage, because the new constraint has completely changed the problem.

In summary, AIKR not only implies that the traditional theories cannot be applied to AI systems, but also suggests new theories for each aspect of the system. Limited by paper length, here we cannot introduce the technical details of NARS, but for the current purpose, it is enough to list the basic ideas, as well as why they are different from the traditional theories. Technical descriptions of NARS can be found in (Wang 2006), as well as in the other publications online at the author's website.

Conclusion

Though the Assumption of Insufficient Knowledge and Resources (AIKR) originates as a constraint on biological intelligent systems, it should be recognized as a fundamental assumption for all forms of intelligence in general. Many cognitive facilities observed in human intelligence can be explained as strategies or mechanisms to work under this constraint.

The traditional theories dominating AI research do not accept AIKR, and minor changes in them will not fix the problem. This mismatch between the theories AI needs and the theories AI uses also suggests a reason for why the mainstream AI has not progressed as fast as expected.

AI research needs a paradigm shift. Here the key issue is not technical, but conceptual. Many people are more comfortable to stay with the traditional theories than to explore new alternatives, because of the well-established position of the former. However, these people fail to realize that the successes of the traditional theories are mostly achieved in other domains, which the research goal and the environmental constraints are fundamentally different from those in the field of AI.

What AI needs are new theories that fully obey AIKR. To establish a new theory is difficult, but still more promising than the other option, that is, try to force the traditional theories to work in a domain where their fundamental assumptions are not satisfied.

The practice of NARS shows that it is technically feasible to build an AI system while fully obeying AIKR, and such a system displays many properties that are similar to the human intelligence. Furthermore, such a system can solve many traditional AI problems in a consistent manner. This evidence supports the conclusion that AIKR should be taken as a cornerstone in the theoretical foundation of AI.

Acknowledgments

Thanks to Stephen Harris for helpful comments and English corrections.

References

- Anderson, J. R. 1990. *The Adaptive Character of Thought*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Barwise, J., and Etchemendy, J. 1989. Model-theoretic semantics. In Posner, M. I., ed., *Foundations of Cognitive Science*. Cambridge, Massachusetts: MIT Press. 207–243.
- Birnbaum, L. 1991. Rigor mortis: a response to Nilsson's "Logic and artificial intelligence". *Artificial Intelligence* 47:57–77.
- Cherniak, C. 1986. *Minimal Rationality*. Cambridge, Massachusetts: MIT Press.
- Good, I. J. 1983. *Good Thinking: The Foundations of Probability and Its Applications*. Minneapolis: University of Minnesota Press.
- Halpern, J. Y.; Harper, R.; Immerman, N.; Kolaitis, P. G.; Vardi, M. Y.; and Vianu, V. 2001. On the unusual effectiveness of logic in computer science. *The Bulletin of Symbolic Logic* 7(2):213–236.
- Holland, J. H. 1986. Escaping brittleness: the possibilities of general purpose learning algorithms applied to parallel rule-based systems. In Michalski, R. S.; Carbonell, J. G.; and Mitchell, T. M., eds., *Machine Learning: an artificial intelligence approach*, volume II. Los Altos, California: Morgan Kaufmann. 593–624.
- Hopcroft, J. E., and Ullman, J. D. 1979. *Introduction to Automata Theory, Language, and Computation*. Reading, Massachusetts: Addison-Wesley.
- Hutter, M. 2005. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Berlin: Springer.
- Korf, R. 2006. Why AI and cognitive science may have little in common. Invited speech at AAAI 2006 Spring Symposium on Cognitive Science Principles Meet AI-Hard Problems.
- Laffey, T. J.; Cox, P. A.; Schmidt, J. L.; Kao, S. M.; and Read, J. Y. 1988. Real-time knowledge based system. *AI Magazine* 9:27–45.
- Lakoff, G. 1988. Cognitive semantics. In Eco, U.; Santambrogio, M.; and Violi, P., eds., *Meaning and Mental Representation*. Bloomington, Indiana: Indiana University Press. 119–154.
- Legg, S., and Hutter, M. 2007. Universal intelligence: a definition of machine intelligence. *Minds & Machines* 17(4):391–444.
- McCarthy, J. 1988. Mathematical logic in artificial intelligence. *Dædalus* 117(1):297–311.
- McDermott, D. 1987. A critique of pure reason. *Computational Intelligence* 3:151–160.
- Medin, D. L., and Ross, B. H. 1992. *Cognitive Psychology*. Fort Worth: Harcourt Brace Jovanovich.
- Nilsson, N. J. 1991. Logic and artificial intelligence. *Artificial Intelligence* 47:31–56.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. San Mateo, California: Morgan Kaufmann Publishers.
- Reeke, G. N., and Edelman, G. M. 1988. Real brains and artificial intelligence. *Dædalus* 117(1):143–173.
- Reiter, R. 1987. Nonmonotonic reasoning. *Annual Review of Computer Science* 2:147–186.
- Russell, S., and Wefald, E. H. 1991. *Do the Right Thing: Studies in Limited Rationality*. Cambridge, Massachusetts: MIT Press.
- Simon, H. A. 1957. *Models of Man: Social and Rational*. New York: John Wiley.
- Strosnider, J. K., and Paul, C. J. 1994. A structured view of real-time problem solving. *AI Magazine* 15(2):45–66.
- Tarski, A. 1944. The semantic conception of truth. *Philosophy and Phenomenological Research* 4:341–375.
- Wang, P. 1993. Non-axiomatic reasoning system (version 2.2). Technical Report 75, Center for Research on Concepts and Cognition, Indiana University, Bloomington, Indiana. Full text available online.
- Wang, P. 1995a. *Non-Axiomatic Reasoning System: Exploring the Essence of Intelligence*. Ph.D. Dissertation, Indiana University.
- Wang, P. 1995b. Reference classes and multiple inheritances. *International Journal of Uncertainty, Fuzziness and Knowledge-based Systems* 3(1):79–91.
- Wang, P. 2004a. Cognitive logic versus mathematical logic. In *Lecture notes of the Third International Seminar on Logic and Cognition*. Full text available online.
- Wang, P. 2004b. The limitation of Bayesianism. *Artificial Intelligence* 158(1):97–106.
- Wang, P. 2005. Experience-grounded semantics: a theory for intelligent systems. *Cognitive Systems Research* 6(4):282–302.
- Wang, P. 2006. *Rigid Flexibility: The Logic of Intelligence*. Dordrecht: Springer.
- Wang, P. 2009. Case-by-case problem solving. In *Artificial General Intelligence 2009*, 180–185.