

Artificial Intelligence

Home Article for *Van Nostrand Scientific Encyclopedia*, Ninth Edition. New York: Wiley.

To appear, 2002

David B. Leake, Indiana University

<http://www.cs.indiana.edu/~leake>

August 10, 2001

1 Introduction

Artificial intelligence (AI) is a branch of computer science that studies the computational requirements for tasks such as perception, reasoning, and learning, and develops systems to perform those tasks. AI is a diverse field whose researchers address a wide range of problems, use a variety of methods, and pursue a spectrum of scientific goals. For example, some researchers study the requirements for expert performance at specialized tasks, while others model commonsense processes; some researchers explain behaviors in terms of low-level processes, using models inspired by the computation of the brain, while others explain them in terms of higher-level psychological constructs such as plans and goals. Some researchers aim to advance understanding of human cognition, some to understand the requirements for intelligence in general (whether in humans or machines), and some to develop artifacts such as intelligent devices, autonomous agents, and systems that cooperate with people to amplify human abilities.

AI is a young field—even its name, “artificial intelligence,” was only coined in 1956. One of the challenges for AI has been to determine which tasks to study—what constitutes an “AI question”—and how to evaluate progress. Much early AI research focused on tasks commonly thought to require high intelligence in people, such as playing high-quality chess. Skeptics viewed this as an impossible assignment, but AI made rapid progress. By the 1960’s, programs were capable of tournament play. In 1997, in a landmark match, the chess system Deep Blue defeated Gary Kasparov, the world’s human chess champion for the previous twelve years. At the same time, however, AI research was illuminating the enormous difficulty of commonsense tasks that people take for granted, such as understanding stories or conversations. Developing programs that can deal at a human level with rich everyday reasoning remains a fundamental research challenge.

The first half-century of AI has yielded a wide range of results. AI research has illuminated the nature of reasoning problems, and the fundamental requirements for intelligent systems. AI research in the area of *cognitive science* has developed models that have helped

to understand human cognition. Applied AI research has provided high-impact applications systems that are in daily use throughout the world. This chapter provides a brief introduction to the history of AI, sketches some major research areas, and closes by illustrating the practical impact of AI technology.

2 The History of AI

The name “artificial intelligence” dates only to the 1950’s, but its roots stretch back thousands of years, into the earliest studies of the nature of knowledge and reasoning. Intelligent artifacts appear in Greek mythology; the idea of developing ways to perform reasoning automatically, and efforts to build automata to perform tasks such as game-playing, date back hundreds of years. Psychologists have long studied human cognition, helping to build up knowledge about the nature of human intelligence. Philosophers have analyzed the nature of knowledge, have studied the *mind-body problem* of how mental states relate to physical processes, and have explored formal frameworks for deriving conclusions.

The advent of electronic computers, however, provided a revolutionary advance in the ability to study intelligence by actually building *intelligent artifacts*—systems to perform complex reasoning tasks—and observing and experimenting with their behavior to identify fundamental principles. In 1950, a landmark paper by Alan Turing argued for the possibility of building intelligent computing systems [Turing, 1950]. That paper proposed an operational test for comparing the intellectual ability of humans and AI systems, now generally called the “Turing Test.” In the Turing Test, a judge uses a teletype to communicate with two players in other rooms: a person and a computer. The judge knows the players only by anonymous labels, such as “player A” and “player B,” on the text that they send to him. By typing questions to the players and examining their answers, the judge attempts to decide which is which. Both the human and machine try to convince the questioner that they are the human; the goal for the machine is to answer so that the judge cannot reliably distinguish which is which.

The game is intended to provide a rich test of intellectual abilities, separated from physical capabilities. The questions are unrestricted; Turing’s samples range from “Please write me a sonnet on the subject of the Forth Bridge,” to “Add 34957 to 70764.” Turing’s examples of possible responses make clear that the aim is to imitate human intelligence, rather than to demonstrate superhuman capabilities: His sample responses are “Count me out on this one. I never could write poetry,” and, after a 30-second pause, 105,621—which is wrong.

The significance of the Turing Test has been controversial. Some, both inside and outside AI, have believed that building a system to pass the Turing Test should be the goal of AI. Others, however, reject the goal of developing systems to imitate human behavior. Ford and Hayes (1998) illustrate this point with an analogy between developing artificial intelligence and developing mechanical flight. Early efforts at mechanical flight were based on trying to imitate the flight of birds, which at that time were the only available examples of flight. How birds flew was not understood, but their observed features (aspects such as beaks, feathers, and flapping wings) could be imitated, and become models for aircraft (even to the extent of airplanes with beaks being featured in a 1900s textbook on aircraft design!) Success at mechanical flight, however, depended on replacing attempts at imitation with study of the

functional requirements for flight, and the development of aircraft that used all available methods to achieve them. In addition, passing the Turing Test is not a precondition for developing useful practical systems. For example, an intelligent system to aid doctors or to tutor students can have enormous practical impact with only the ability to function in a specific, limited domain.

2.1 The First Decades

Turing’s paper surveys many common arguments against the possibility of AI and provides responses to each one. One of these arguments is that machines “can only do what they are programmed to do,” from which some conclude that programs could never “take us by surprise.” Shortly after the appearance of Turing’s paper, a program provided concrete proof that programs can go beyond their creators: Arthur Samuel wrote the first checkers-playing program, which used learning techniques to develop tournament-level skills, surpassing its creator’s own abilities [Samuel, 1963].

Early AI research rapidly developed systems to perform a wide range of tasks often associated with intelligence in people, including theorem-proving in geometry, symbolic integration, solving equations, and even solving analogical reasoning problems of the types sometimes found on human intelligence tests. However, research also revealed that methods which worked well on small sample domains might not “scale up” to larger and richer tasks, and led to an awareness of the enormous difficulty of the problems that the field aimed to address. A classic example concerns early work in machine translation, which was recognized in the 1960’s to be a far more difficult problem than expected—causing the termination of funding for machine translation research.

Two impediments to wider application of early AI systems were their general methods and lack of knowledge. For small tasks, exhaustively considering possibilities may be practical, but for rich tasks, specialized knowledge is needed to focus reasoning. This observation led to research on knowledge-based systems, which demonstrated that there is an important class of problems requiring deep but narrow knowledge, and that systems capturing this knowledge in the form of rules can achieve expert-level performance for these tasks. An early example, DENDRAL [Feigenbaum and Buchanan, 1993], used rules about mass spectrometry and other data to hypothesize structures for chemical compounds. Using only simple inference methods, it achieved expert-level performance and was the source of results published in the chemical literature. Such systems provided the basis for numerous applied AI systems (See EXPERT SYSTEMS). Continuing research revealed the need to develop additional methods for tasks such as acquiring the knowledge for systems to use, dealing with incomplete or uncertain information, and automatically adapting to new tasks and environments.

The accompanying timeline (<http://www.aaai.org/AITopics/bbhist.html>), prepared by Bruce Buchanan, provides a list of major milestones in the development of AI, and Russell and Norvig provide historical summary of the field in Chapter 1 of their AI textbook [Russell and Norvig, 1995]. An article by Hearst and Hirsh [Hearst and Hirsh, 2000] presents a range of viewpoints on the greatest trends and controversies in AI, collected from leading figures in the development of artificial intelligence.

3 AI Perspectives

Just as AI researchers must select the goals they will pursue, they must select the frameworks within which to pursue them. These frameworks provide a perspective on AI problems, shaping researchers' choices of which questions to address, how to address them, and what constitutes an answer. One perspective, which can be described as *biomorphic*, takes inspiration from biological systems. Neural network models, for example, are inspired by neurons in the brain (See NEURAL NETWORKS). Another example is genetic algorithms, which take their inspiration from evolution, “evolving” promising solutions by a simulated process of natural selection (See GENETIC ALGORITHMS AND EVOLUTIONARY COMPUTATION). Such models may be used not only for the pragmatic goals of solving difficult problems, but also to study the biological processes that they model, in order to increase understanding of the factors affecting living organisms (See ARTIFICIAL LIFE).

Another perspective takes its inspiration from human cognition, focusing on functional constraints rather than on biologically-inspired mechanisms. An illustration is research on case-based reasoning (CBR), which was inspired by the role of memory in human problem-solving. For example, doctors use case-based reasoning when they treat an illness by remembering a similar previous case—the treatment of a previous patient with similar symptoms—and adapting the prior treatment to fit changed circumstances (e.g., adjusting the dosage for a child) (See CASE-BASED REASONING.) This view of problem-solving suggests studying issues such as how a memory of cases must be organized to model the retrievals of human reasoners, which can provide hypotheses about human reasoning as well as useful mechanisms for AI systems. [Leake, 1998] describes how case-based reasoning provides a stance towards cognitive science, and [Leake, 1996] provides an overview of major trends in CBR research and applications.

Yet another perspective is more technological: it studies the requirements and mechanisms for intelligence, without restricting the mechanisms considered. Practitioners seeking to develop useful systems, and researchers interested in understanding the general nature of intelligence, need not be constrained by biology or psychology—the processes that evolved in human reasoners are not necessarily the best ones for achieving high-quality performance in intelligent machines. For example, studies of the psychology of chess suggest that chess masters consider perhaps two moves per second, with their ability to recognize known board patterns playing a key role in their choice of moves. Deep Blue, however, defeated Gary Kasparov by exploiting a special architecture that enabled it to consider 200 million positions per second (See ARTIFICIAL INTELLIGENCE AND GAMES).

4 A Sampling of AI Research Areas

4.1 Search

In 1976, Newell and Simon [Newell and Simon, 1976] proposed that intelligent behavior arises from the manipulation of *symbols*—entities that represent other entities, and that the process by which intelligence arises is *heuristic search*. Search is a process of formulating and examining alternatives. It starts with an initial state, a set of candidate actions, and

5	7	1
4		3
8	2	6

1	2	3
4	5	6
7	8	

Figure 1: Sample initial and goal states for the 8 puzzle.

criteria for identifying the goal state. It is often guided by heuristics, or “rules of thumb,” which are generally useful, but not guaranteed to make the best choices. Starting from the initial state, the search process selects actions to transform that state into new states, which themselves are transformed into more new states, until a goal state is generated. For example, consider a search program to solve the “8-puzzle” for children, which is shown in Figure 1. A child solves the puzzle by sliding the numbered tiles (without lifting them) to reach a configuration in which the tiles are all in numerical order, as shown in the second board in the figure. When the 8 puzzle is seen as a search problem, the initial state is a starting board position, each action is a possible move of one tile up, down, left, or right (when the position it will move to is blank), and the goal state is the second state in Figure 1. Here a heuristic function might suggest candidate moves by comparing their results to the goal, in order to favor those moves that appear to be making progress towards the solution. For this search problem, what is of interest is the solution path—how the solution was generated. However, for some problems, only the final state is important—a designer may only be interested in generating a successful design, rather than how it was generated.

A central problem in search is the *combinatorial explosion* of alternatives to consider. For example, if there are 10 possible actions from each state, after 5 moves there are a million possibilities to consider for the next move. Numerous techniques have been developed to improve search performance, and the combination of intelligent strategies and special-purpose computing hardware has enabled AI systems to rapidly search enormous spaces of alternatives. For examples of the role of search in two specific AI areas, see AUTOMATED REASONING and ARTIFICIAL INTELLIGENCE AND GAMES.

4.2 Knowledge capture, representation and reasoning

In order to guide search—or even to describe problems, actions, and solutions—the relevant domain knowledge must be encoded in a form that can be effectively manipulated by a program. More generally, the usefulness of any reasoning process depends not only on the reasoning process itself, but also on having the right knowledge and representing it in a form the program can use.

In the logicist approach to knowledge representation and reasoning, information is encoded as assertions in a logic, and the system draws conclusions by deduction from those assertions (See AUTOMATED REASONING). Other research studies non-deductive forms of reasoning, such as reasoning by analogy and abductive inference—the process of inferring

the best explanation for a set of facts. Abductive inference does not guarantee sound conclusions, but is enormously useful for tasks such as medical diagnosis, in which a reasoner must hypothesize causes for a set of symptoms.

Capturing the knowledge needed by AI systems has proven to be a challenging task. The knowledge in rule-based expert systems, for example, is represented in the form of rules listing conditions to check for, and conclusions to be drawn if those conditions are satisfied. For example, a rule might state that *IF* certain conditions hold (e.g., the patient has certain symptoms), *THEN* certain conclusions should be drawn (e.g., that the patient has a particular condition or disease). A natural way to generate these rules is to interview experts. Unfortunately, the experts may not be able to adequately explain their decisions in a rule-based way, resulting in a “knowledge-acquisition bottleneck” impeding system development.

One approach to alleviating the knowledge acquisition problem is to develop sharable knowledge sources that represent knowledge in a form that can be re-used across multiple tasks. The CYC project, for example, is a massive ongoing effort to encode the “consensus knowledge” that underlies much commonsense reasoning [Lenat, 1995]. Much current knowledge representation research develops sharable *ontologies* that represent particular domains. Ontologies provide a formal specification of the concepts in the domain and their relationships, to use as a foundation for developing knowledge bases and facilitating knowledge sharing [Chandrasekaran *et al.*, 1999].

4.3 Reasoning under uncertainty

AI systems—like people—must often act despite partial and uncertain information. First, the information received may be unreliable (e.g., a patient may mis-remember when a disease started, or may not have noticed a symptom that is important to a diagnosis). In addition, rules connecting real-world events can never include all the factors that might determine whether their conclusions really apply (e.g., the correctness of basing a diagnosis on a lab test depends whether there were conditions that might have caused a false positive, on the test being done correctly, on the results being associated with the right patient, etc.) Thus in order to draw useful conclusions, AI systems must be able to reason about the probability of events, given their current knowledge (See PROBABILITY). Research on Bayesian reasoning provides methods for calculating these probabilities. *Bayesian networks*, graphical models of the relationships between variables of interest, have been applied to a wide range of tasks, including natural language understanding, user modeling, and medical diagnosis. For example, Intellipath, a commercial system for pathology diagnosis, was approved by the AMA and has been fielded in hundreds of hospitals worldwide. Diagnostic reasoning may also be combined with reasoning about the value of alternative actions, in order to select the course of action with the greatest expected utility. For example, a medical decision-making system might make decisions by considering the probability of a patient having a particular condition, the probability of bad side-effects of a treatment and their severity, and the probability and severity of bad effects if the treatment is not performed.

In addition to dealing with uncertain information, everyday reasoners must be able to deal with vague descriptions, such as those provided in natural language. For example, a doctor who is told that a patient has a “high fever,” must be able to reason about the fuzzy

concept of “high fevers.” Whether a particular fever is “high” is not simply a true or false decision decided by a cutoff point, but rather, a matter of degree. Fuzzy reasoning provides methods for reasoning about vague knowledge (see FUZZY REASONING).

4.4 Planning, Vision, and Robotics

The conclusions of the reasoning process can determine goals to be achieved. Planning addresses the question of how to determine a sequence of actions to achieve those goals. The resulting action sequences may be designed to be applied in many ways, such as by robots in the world, by intelligent agents on the Internet, or even by humans. Planning systems may use a number of techniques to make the planning process practical, such as *hierarchical planning*, reasoning first at higher levels of abstraction and then elaborating details within the high-level framework (e.g., as a person might do when first outlining general plans for a trip, and then considering fine-grained details such as how to get to the airport), and *partial-order* planning, enabling actions to be inserted in the plan in any order, rather than chronologically, and subplans to be merged. Dean and Kambhampati (1997) provide an extensive survey of this area.

In real-world situations, it is seldom possible to generate a complete plan in advance and then execute it without changes. The state of the world may be imperfectly-known, the effects of actions may be uncertain, the world may change while the plan is being generated or executed, and the plan may require the coordination of multiple cooperating agents, or counterplanning to neutralize the interference of agents with opposing goals. Determining the state of the world and guiding action requires the ability to gather information about the world, though sensors such as sonar or cameras, and to interpret that information to draw conclusions (See MACHINE VISION). In addition, carrying out actions in a messy and changing world may require rapid responses to important events (e.g., for a robot-guided vehicle to correct a skid), or an ongoing process of rapidly selecting actions based on the current context (for example, when a basketball player must avoid an opponent). Such problems have led to research on reactive planning, as well as on how to integrate reactive methods with the deliberative methods providing long-term guidance (See ROBOTICS). The RoboCup Federation sponsors an annual series of competitions between robot soccer teams as a testbed for demonstrating new methods and extending the state of the art in robotics (www.robocup.org).

4.5 Natural language processing

Achieving natural interactions between humans and machines requires machines to understand and generate language. Likewise, understanding human communication requires the understanding of how language is processed by people. The nature of human language raises many challenging issues for language processing systems: natural language is elliptic, leaving much unstated, and its meaning is context-dependent (“Mary took aspirin” will have a different meaning when explaining how she recovered from her headache, or her arrest for shoplifting). Some natural language processing approaches investigate algorithms for syntactic parsing, to determine the grammatical structure of textual passages; others take a

cognitively-inspired view, studying the knowledge structures underlying human understanding and modeling the process by which they are applied, or even attempting to directly apply expectations from memory to the parsing process. Other systems apply statistical methods to tasks such as information extraction from newspaper articles. Machine translation systems, though still far from replacing human translators for literature, can now generate useful translations (See NATURAL LANGUAGE PROCESSING).

4.6 Machine Learning

In a complex world, it is difficult to encode all the knowledge that a system may need, and may also be difficult to keep system knowledge up-to-date. Machine learning research focuses on how AI systems can augment or refine their own knowledge to improve their performance. Just as people use different learning techniques, machine learning systems use a wide range of approaches. Some of these are *supervised*, in that they presume that the learner will have access to the correct answers; others are *unsupervised*, requiring the learner to proceed without benefit of feedback.

Inductive learning systems learn by analyzing examples to identify correlations between inputs and outputs. For example, neural network models process inputs according to networks of idealized neurons, and learn by algorithms that adjust the weights of neural connections based on correlations between inputs and outputs in training examples. A neural network system to recognize faces might be trained on a digitized set of photographs of faces (inputs) and the associated identities (outputs), to learn which facial features are correlated with different individuals (See NEURAL NETWORKS). Theory-driven learning approaches use background knowledge to guide generalizations, in order to focus on important types of features. Instance-based learning systems and case-based reasoners perform “lazy learning:” rather than attempting to generalize experiences as they are encountered, case-based reasoning systems store learned cases as-is, adapting or generalizing their lessons only if needed to solve new problems (See MACHINE LEARNING).

5 Practical Impact of AI

AI technology has had broad impact. AI components are embedded in numerous devices, such as copy machines that combine case-based reasoning and fuzzy reasoning to automatically adjust the copier to maintain copy quality. AI systems are also in everyday use for tasks such as identifying credit card fraud, configuring products, aiding complex planning tasks, and advising physicians. AI is also playing an increasing role in corporate knowledge management, facilitating the capture and reuse of expert knowledge. Intelligent tutoring systems make it possible to provide students with more personalized attention, and even for the computer to listen to what children say and respond to it (<http://www.cs.cmu.edu/~listen/>). Cognitive models developed by AI can also suggest principles for effective support for human learning, guiding the design of educational systems [Leake and Kolodner, 2001].

AI technology is being used in autonomous agents that independently monitor their surroundings, make decisions and act to achieve their goals without human intervention. For example, in space exploration, the lag times for communications between earth and

probes make it essential for robotic space probes to be able to perform their own decision-making—Depending on the relative locations of the earth and Mars, one-way communication can take over 20 minutes. In a 1999 experiment, an AI system was given primary control of a spacecraft, NASA’s Deep Space 1, 60,000,000 miles from earth, as a step towards autonomous robotic exploration of space (see rax.arc.nasa.gov). Methods from autonomous systems also promise to provide important technologies to aid humans. For example, in a 1996 experiment called “No Hands Across America,” the RALPH system [Pomerleau and Jochem, 1996], a vision-based adaptive system to learn road features, was used to drive a vehicle for 98 percent of a trip from Washington, D.C., to San Diego, maintaining an average speed of 63 mph in daytime, dusk and night driving conditions. Such systems could be used not only for autonomous vehicles, but also for safety systems to warn drivers if their vehicles deviate from a safe path.

In electronic commerce, AI is providing methods for determining which products buyers want and configuring them to suit buyers’ needs. The explosive growth of the internet has also led to growing interest in internet agents to monitor users’ tasks, seek needed information, and learn which information is most useful [Hendler, 1999]. For example, the Watson system monitors users as they perform tasks using standard software tools such as word processors, and uses the task context to focus search for useful information to provide to them as they work [Budzik and Hammond, 2000].

Continuing investigation of fundamental aspects of intelligence promises broad impact as well. For example, researchers are studying the nature of creativity and how to achieve creative computer systems, providing strong arguments that creativity can be realized by artificial systems [Hofstadter, 1985]. Numerous programs have been developed for tasks that would be considered creative in humans, such as discovering interesting mathematical concepts, in the program AM [Lenat, 1979], making paintings, in Aaron [Cohen, 1995], and performing creative explanation, in SWALE [Schank and Leake, 1989]. The task of AM, for example, was not to prove mathematical theorems, but to discover interesting concepts. The program was provided only with basic background knowledge from number theory (e.g., the definition of sets), and with heuristics for revising existing concepts and selecting promising concepts to explore. Starting from this knowledge, it discovered fundamental concepts such as addition, multiplication, and prime numbers. It even rediscovered a famous mathematical conjecture that was not known to its programmer: Goldbach’s conjecture, the conjecture that every even integer greater than 2 can be written as the sum of two primes. Buchanan (2001) surveys some significant projects in machine creativity and argues for its potential impact on the future of artificial intelligence.

In addition, throughout the history of AI, AI research has provided a wellspring of contributions to computer science in general. For example, the computer language Lisp, developed by John McCarthy in 1958, provided a tool for developing early AI systems using symbolic computation, but has remained in use to the present day, both within and outside AI, and has had significant influence on the area of programming languages. Later AI research also gave rise to the computer language, Prolog, used for logic programming. A key idea of logic programming is that the programmer should specify only the problem to be solved and constraints on its solution, leaving the system itself to determine the details of how the solution should be obtained.

6 Conclusion and Resources

In its short existence, AI has increased understanding of the nature of intelligence and provided an impressive array of applications in a wide range of areas. It has sharpened understanding of human reasoning, and of the nature of intelligence in general. At the same time, it has revealed the complexity of modeling human reasoning, providing new areas and rich challenges for the future.

AAAI, the American Association for Artificial Intelligence, maintains an extensive on-line library of articles on AI, ranging from general introductions to focused articles on specific areas, at <http://www.aaai.org/Pathfinder/>. *AI Magazine*, the official magazine of AAAI, publishes accessible articles on current research and applications, as well as tutorials on important AI areas. After a delay, full-text electronic versions of articles from back issues are freely available from the magazine home page <http://www.aimagazine.org>. The magazines *IEEE Intelligent Systems* and *Intelligence* are additional sources for accessible articles on new developments in AI and its applications.

7 Acknowledgment

I would like to thank Raja Sooriamurthi for helpful comments on a draft of this article.

References

- [Buchanan, 2001] Buchanan, B. 2001. Creativity at the meta-level. *AI Magazine*. In press.
- [Budzik and Hammond, 2000] Budzik, J. and Hammond, K. 2000. User interactions with everyday applications as context for just-in-time information access. In *Proceedings of the 2000 International Conference on Intelligent User Interfaces*. 44–51. <http://lieber.www.media.mit.edu/people/lieber/IUI/Budzik/Budzik.pdf>.
- [Chandrasekaran *et al.*, 1999] Chandrasekaran, B.; Josephson, J.; and Benjamins, R. 1999. What are ontologies, and why do we need them? *IEEE Intelligent Systems* 14(1).
- [Cohen, 1995] Cohen, H. 1995. The further exploits of AARON, painter. *Stanford Humanities Review* 4. <http://www.stanford.edu/group/SHR/4-2/text/cohen.html>.
- [Dean and Kambhampati, 1997] Dean, T. and Kambhampati, S. 1997. Planning and scheduling. In *The Computer Science and Engineering Handbook*. CRC Press, Hillsdale, NJ. 614–636. <http://www.research.salford.ac.uk/plansig/papers/crc-chapter.pdf>.
- [Feigenbaum and Buchanan, 1993] Feigenbaum, E. A. and Buchanan, B. G. 1993. Dendral and meta-dendral: Roots of knowledge systems and expert system applications. *Artificial Intelligence* 59:233–240.
- [Ford and Hayes, 1998] Ford, K. and Hayes, P. 1998. On computational wings: Rethinking the goals of artificial intelligence. *Scientific American Presents* 9(4):78–83.

- [Hearst and Hirsh, 2000] Hearst, M. and Hirsh, H. 2000. AI's greatest trends and controversies. *IEEE Intelligent Systems* 15(1):8–17. <http://www.computer.org/intelligent/ex2000/pdf/x1008.pdf>.
- [Hendler, 1999] Hendler, J. 1999. Is there an intelligent agent in your future? *Nature Webmatters*. <http://helix.nature.com/webmatters/agents/agents.html>.
- [Hofstadter, 1985] Hofstadter, D. 1985. On the seeming paradox of mechanizing creativity. In *Metamagical Themas*. Basic Books, New York. 525–546.
- [Leake and Kolodner, 2001] Leake, D. and Kolodner, J. 2001. Learning through case analysis. In *Encyclopedia of Cognitive Science*. Macmillan, London. In press.
- [Leake, 1996] Leake, D. 1996. CBR in context: The present and future. In Leake, D., editor 1996, *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. AAAI Press, Menlo Park, CA. 3–30. <http://www.cs.indiana.edu/leake/papers/a-96-01.html>.
- [Leake, 1998] Leake, D. 1998. Cognition as case-based reasoning. In Bechtel, W. and Graham, G., editors 1998, *A Companion to Cognitive Science*. Blackwell, Oxford. 465–476.
- [Lenat, 1979] Lenat, D. 1979. On automated scientific theory formation: A case study using the AM program. In Hayes, J.; Mitchie, D.; and Milulich, L., editors 1979, *Machine Intelligence*, volume 9. Halsted Press.
- [Lenat, 1995] Lenat, D. 1995. A large-scale investment in knowledge infrastructure. *Communications of the ACM* 38(11):33–38.
- [Newell and Simon, 1976] Newell, A. and Simon, H. 1976. Computer science as empirical inquiry: Symbols and search. *Communications of the ACM* 19:113–126. Reprinted in Haugeland, J., ed, *Mind Design II*, MIT Press, 1997.
- [Pomerleau and Jochem, 1996] Pomerleau, D. and Jochem, T. 1996. A rapidly adapting machine vision system for automated vehicle steering. *IEEE Expert* 11(2):19–27.
- [Russell and Norvig, 1995] Russell, S. and Norvig, P. 1995. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Englewood Cliffs, NJ.
- [Samuel, 1963] Samuel, A.L. 1963. Some studies in machine learning using the game of checkers. In Feigenbaum, E.A. and Feldman, J., editors 1963, *Computers and thought*. McGraw-Hill. Also in IBM Journal of Research and Development (1959).
- [Schank and Leake, 1989] Schank, R.C. and Leake, D. 1989. Creativity and learning in a case-based explainer. *Artificial Intelligence* 40(1-3):353–385. Also in Carbonell, J., editor, *Machine Learning: Paradigms and Methods*, MIT Press, Cambridge, MA, 1990.
- [Turing, 1950] Turing, A. 1950. Computing machinery and intelligence. *Mind* 59. Reprinted in J. Haugeland, Ed., *Mind Design II*, MIT Press, 1997.