NATIONAL SECURITY AGENCY
CENTRAL SECURITY SERVICE
FORT GEORGE G. MEADE, MARYLAND 20755-6000

FOIA Case: 83602A
15 February 2017

JOHN GREENEWALD
[redacted]

Dear Mr. Greenewald:

This responds to your Freedom of Information Act (FOIA) request of 27 January 2016 for Intellipedia entries on "Artificial Intelligence". As stated in our previous letter, dated 28 January 2016, your request was assigned Case Number 83602. A copy of your request is enclosed. For purposes of this request and based on the information you provided in your letter, you are considered an "all other" requester. As such, you are allowed 2 hours of search and the duplication of 100 pages at no cost. There are no assessable fees for this request. Your request has been processed under the FOIA.

For your information, NSA provides a service of common concern for the Intelligence Community (IC) by serving as the executive agent for Intelink. As such, NSA provides technical services that enable users to access and share information with peers and stakeholders across the IC and DoD. Intellipedia pages are living documents that may be originated by any user organization, and any user organization may contribute to or edit pages after their origination. Intellipedia pages should not be considered the final, coordinated position of the IC on any particular subject. The views and opinions of authors do not necessarily state or reflect those of the U.S. Government.

We conducted a search of all three levels of Intellipedia for the requested topic, and located three documents that are responsive to your request. These documents are enclosed. Certain information, however, has been deleted from the enclosures.

This Agency is authorized by statute to protect certain information concerning its activities (in this case, internal URLs) as well as the names of its employees. Such information is exempt from disclosure pursuant to the third exemption of the FOIA, which provides for the withholding of information specifically protected from disclosure by statute. The specific statute applicable in this case is Section 6, Public Law 86-36 (50 U.S. Code 3605). We have determined that such information exists in this record, and we have excised it accordingly.

In addition, personal information regarding individuals has been deleted from the enclosures in accordance with 5 U.S.C. 552 (b)(6). This exemption protects from disclosure information that would constitute a clearly unwarranted invasion of personal privacy. In balancing the public interest for the information you request against the privacy interests involved, we have determined that the privacy interests sufficiently satisfy the requirements for the application of the (b)(6) exemption.

Since these deletions may be construed as a partial denial of your request, you are hereby advised of this Agency's appeal procedures. You may appeal this decision. If you decide to appeal, you should do so in the manner outlined below.

- The appeal must be in writing and addressed to:

  NSA/CSS FOIA/PA Appeal Authority (P132),
  National Security Agency
  9800 Savage Road STE 6932
  Fort George G. Meade, MD 20755-6932

- The request must be postmarked no later than 90 calendar days of the date of this letter. Decisions not appealed within 90 days will not be addressed.
- Please include the case number provided above.
- Please describe with sufficient detail why you believe the denial of requested information was unwarranted.
- NSA will endeavor to respond within 20 working days of receiving your appeal, absent any unusual circumstances.

You may also contact our NSA FOIA Public Liaison at foialo@nsa.gov for any further assistance and to discuss any aspect of your request. Additionally, you may contact the Office of Government Information Services (OGIS) at the National Archives and Records Administration to inquire about the FOIA mediation services they offer. OGIS contact information is: Office of Information Services, National Archives and Records Administration, 8601 Adelphi Road-OGIS, College Park, MD 20740-6001; e-mail ogis@nara.gov; 202-741-5770; toll free 1-877-684-6448; or fax 202-741-5769.

Sincerely,

Paul N
for

JOHN R. CHAPMAN
Chief, FOIA/PA Office
NSA Initial Denial Authority

Encls:
a/s

Archer, Lynn M

From: donotreply@nsa.gov
Sent: Wednesday, January 27, 2016 7:17 PM
To: donotreply@nsa.gov
Cc: john@greenewald.com
Subject: FOIA Request (Web form submission)

Name: John Greenewald

Title: Mr.

Email: john@greenewald.com

Company: The Black Vault

Postal Address: [redacted]

Postal City: [redacted]

Postal State-prov: [redacted]

Zip Code: [redacted]

Country: United States of America

Home Phone: [redacted]

Work Phone: [redacted]

Records Requested: To whom it may concern,

This is a non-commercial request made under the provisions of the Freedom of Information Act 5 U.S.C. S 552. My FOIA requester status as a "representative of the news media" however due to your agency's denial of this status, I hereby submit this request as an "All other" requester.

I prefer electronic delivery of the requested material either via email to john@greenewald.com or via CD-ROM or DVD via postal mail. Please contact me should this FOIA request should incur a charge.

I respectfully request a copy of the Intellipedia entry (from all three Wikis that make up the Intellipedia) for the following entry(s) (Or whatever similar topic may pertain if it is slightly worded differently):

ARTIFICIAL INTELLIGENCE

Thank you so much for your time, and I am very much looking forward to your response.

Sincerely,

John Greenewald, Jr.
[redacted]

[redacted]

(U//~~FOUO~~) Artificial intelligence

UNCLASSIFIED

(b)(3) - P.L. 86-36

From Intellipedia

The discipline concerned with the technology, the background, and the theory of building computer systems that perform functions that would be said to require intelligence if performed by humans or animals (or functions that might otherwise be described as intelligent). AI is an interdisciplinary area with inputs from computer science, psychology, neuroscience, cybernetics, linguistics and philosophy. AI is concerned with such things as natural language understanding, machine learning, computer vision, machine reasoning, expert systems. There are two main paradigms of AI - the symbolic paradigm where information is explicitly represented in a symbolic form (this corresponds to the knowledge based approach - and the subsymbolic paradigm (corresponding to the neural networks approach).

Acronym: AI

Broader Terms:

- machine theory

Narrower Terms:

- machine translation
- natural language processing

Related Terms:

- character recognition
- cognitive science
- neural computing
- object oriented programming
- robotics
- voice recognition

Retrieved from "[redacted]"
Category: Technology

UNCLASSIFIED

- This page has been accessed 1,458 times.
- 0 watching users
- This page was last modified 17:53, 5 March 2007 by [redacted] Most recent editors: [redacted]

(b)(6)

linipedweb9s

[redacted]

criminal prosecution.
Evidence of unauthorized use collected during monitoring may be used for administrative, criminal, or other adverse actions.

This page contains dynamic content -- Highest Possible Classification is **SECRET//NOFORN**

(b)(3) - P.L. 86-36

(b)(3) - P.L. 86-36

(U) Artificial intelligence

UNCLASSIFIED

From Intellipedia

You have new messages (last change).

For other uses, see AI.

(U) Be bold in modifying this **Wikipedia** import .

(U) Correct mistakes; remove bias; categorize; delete superfluous links, templates, and passages; add classified information and citations.

(U) When assimilation into Intellipedia is complete, remove this template and add {{From Wikipedia}}.

Contents

(b)(3) - P.L. 86-36

Definition

(U) A computer science discipline concerned with building systems that exhibit the characteristics associated with intelligence in human behavior, e.g., understanding language, learning from experience, logical reasoning, solving problems, and explaining one's own behavior.[1]

Major AI textbooks define the field as "the study and design of intelligent agents,"[2] where an intelligent agent is a system that perceives its environment and takes actions which maximize its chances of success.[3] John McCarthy, who coined the term in 1956,[4] defines it as "the science and engineering of making intelligent machines."[5]

The field was founded on the claim that a central property of human beings, intelligence—the sapience of *Homo sapiens*—can be so precisely described that it can be simulated by a machine.[6] This raises philosophical issues about the nature of the mind and limits of scientific hubris, issues which have been addressed by myth, fiction and philosophy since antiquity.[7] Artificial intelligence has been the subject of breathtaking optimism,[8] has suffered stunning setbacks[9] and, today, has become an essential part of the technology industry, providing the heavy lifting for many of the most difficult problems in computer science.

AI research is highly technical and specialized, so much so that some critics decry the "fragmentation" of the field.[10] Subfields of AI are organized around particular problems, the application of particular tools and around long standing theoretical differences of opinion. The central problems of AI include such traits as reasoning, knowledge, planning, learning, communication, perception and the ability to move and manipulate objects.[11] General intelligence (or "strong AI") is still a long term goal of (some) research.[12]

Perspectives on AI

AI in myth, fiction and speculation

Ethics of artificial intelligence Main article: Artificial intelligence in fiction

Thinking machines and artificial beings appear in Greek myths, such as Talos of Crete, the golden robots of Hephaestus and Pygmalion's Galatea.[13] Human likenesses believed to have intelligence were built in many ancient societies; some of the earliest being the sacred statues worshipped in Egypt and Greece,

(b)(3) - P.L. 86-36

[14][15] and including the machines of Yan Shi,[16] Hero of Alexandria,[17] Al-Jazari[18] or Wolfgang von Kempelen.[19] It was widely believed that artificial beings had been created by Geber,[20] Judah Loew[21] and Paracelsus.[22] Stories of these creatures and their fates discuss many of the same hopes, fears and ethical concerns that are presented by artificial intelligence.[7]

Mary Shelley's *Frankenstein*,[23] considers a key issue in the ethics of artificial intelligence: if a machine can be created that has intelligence, could it also *feel*? If it can feel, does it have the same rights as a human being? The idea also appears in modern science fiction: the film *Artificial Intelligence: A.I.* considers a machine in the form of a small boy which has been given the ability to feel human emotions, including, tragically, the capacity to suffer. This issue, now known as "robot rights", is currently being considered by, for example, California's Institute for the Future,[24] although many critics believe that the discussion is premature.[25]

Another issue explored by both science fiction writers and futurists is the impact of artificial intelligence on society. In fiction, AI has appeared as a servant (R2D2 in *Star Wars*), a law enforcer (K.I.T.T. "Knight Rider"), a comrade (Lt. Commander Data in *Star Trek*), a conqueror (*The Matrix*), a dictator (*With Folded Hands*), an exterminator (*Terminator*, *Battlestar Galactica*), an extension to human abilities (*Ghost in the Shell*) and the saviour of the human race (R. Daneel Olivaw in the *Foundation Series*). Academic sources have considered such consequences as: a decreased demand for human labor;[26] the enhancement of human ability or experience;[27] and a need for redefinition of human identity and basic values.[28]

Several futurists argue that artificial intelligence will transcend the limits of progress and fundamentally transform humanity. Ray Kurzweil has used Moore's law (which describes the relentless exponential improvement in digital technology with uncanny accuracy) to calculate that desktop computers will have the same processing power as human brains by the year 2029, and that by 2045 artificial intelligence will reach a point where it is able to improve *itself* at a rate that far exceeds anything conceivable in the past, a scenario that science fiction writer Vernor Vinge named the "technological singularity".[27] Edward Fredkin argues that "artificial intelligence is the next stage in evolution,"[29] an idea first proposed by Samuel Butler's *Darwin Among the Machines* (1863), and expanded upon by George Dyson in his book of the same name in 1998. Several futurists and science fiction writers have predicted that human beings and machines will merge in the future into cyborgs that are more capable and powerful than either. This idea, called transhumanism, which has roots in Aldous Huxley and Robert Ettinger, is now associated with robot designer Hans Moravec, cyberneticist Kevin Warwick and inventor Ray Kurzweil.[27] Transhumanism has been illustrated in fiction as well, for example in the manga *Ghost in the Shell* and the science fiction series *Dune*. Pamela McCorduck writes that these scenarios are expressions of an ancient human desire to, as she calls it, "forge the gods."[7]

History of AI research

timeline of artificial intelligence Main article: history of artificial intelligence

In the middle of the 20th century, a handful of scientists began a new approach to building intelligent machines, based on recent discoveries in neurology, a new mathematical theory of information, an understanding of control and stability called cybernetics, and above all, by the invention of the digital computer, a machine based on the abstract essence of mathematical reasoning.[30]

(b)(3) - P.L. 86-36

The field of modern AI research was founded at a conference on the campus of Dartmouth College in the summer of 1956.[31] Those who attended would become the leaders of AI research for many decades, especially John McCarthy, Marvin Minsky, Allen Newell and Herbert Simon, who founded AI laboratories at MIT, CMU and Stanford. They and their students wrote programs that were, to most people, simply astonishing:[32] computers were solving word problems in algebra, proving logical theorems and speaking English.[33] By the middle 60s their research was heavily funded by the U.S. Department of Defense[34] and they were optimistic about the future of the new field:

- 1965, H. A. Simon: "[M]achines will be capable, within twenty years, of doing any work a man can do"[35]
- 1967, Marvin Minsky: "Within a generation ... the problem of creating 'artificial intelligence' will substantially be solved."[36]

These predictions, and many like them, would not come true. They had failed to recognize the difficulty of some of the problems they faced.[37] In 1974, in response to the criticism of England's Sir James Lighthill and ongoing pressure from Congress to fund more productive projects, the U.S. and British governments cut off all undirected, exploratory research in AI. This was the first AI winter.[38]

In the early 80s, AI research was revived by the commercial success of expert systems[39] (a form of AI program that simulated the knowledge and analytical skills of one or more human experts). By 1985 the market for AI had reached more than a billion dollars and governments around the world poured money back into the field.[40] However, just a few years later, beginning with the collapse of the Lisp Machine market in 1987, AI once again fell into disrepute, and a second, more lasting AI winter began.[41]

In the 90s and early 21st century AI achieved its greatest successes, albeit somewhat behind the scenes. Artificial intelligence was adopted throughout the technology industry, providing the heavy lifting for logistics, data mining, medical diagnosis and many other areas.[42] The success was due to several factors: the incredible power of computers today (see Moore's law), a greater emphasis on solving specific subproblems, the creation of new ties between AI and other fields working on similar problems, and above all a new commitment by researchers to solid mathematical methods and rigorous scientific standards.[43]

Philosophy of AI

Artificial intelligence, by claiming to be able to recreate the capabilities of the human mind, is both a challenge and an inspiration for philosophy. Are there limits to how intelligent machines can be? Is there an essential difference between human intelligence and artificial intelligence? Can a machine have a mind and consciousness? A few of the most influential answers to these questions are given below.[44]

Turing's "polite convention"
: *If a machine acts as intelligently as a human being, then it is as intelligent as a human being.* Alan Turing theorized that, ultimately, we can only judge the intelligence of machine based on its behavior. This theory forms the basis of the Turing test.[45]

The Dartmouth proposal
: *"Every aspect of learning or any other feature of intelligence can be so precisely described that a machine can be made to simulate it."* This assertion was printed in the proposal for the Dartmouth

Conference of 1956, and represents the position of most working AI researchers.[6]

Newell and Simon's physical symbol system hypothesis

"A physical symbol system has the necessary and sufficient means of general intelligent action." This statement claims that the essence of intelligence is symbol manipulation.[46] Hubert Dreyfus argued that, on the contrary, human expertise depends on unconscious instinct rather than conscious symbol manipulation and on having a "feel" for the situation rather than explicit symbolic knowledge.[47][48]

Gödel's incompleteness theorem

A formal system (such as a computer program) can not prove all true statements. Roger Penrose is among those who claim that Gödel's theorem limits what machines can do.[49][50]

Searle's strong AI hypothesis

"The appropriately programmed computer with the right inputs and outputs would thereby have a mind in exactly the same sense human beings have minds."[51] Searle counters this assertion with his Chinese room argument, which asks us to look *inside* the computer and try to find where the "mind" might be.[52]

The artificial brain argument

The brain can be simulated. Hans Moravec, Ray Kurzweil and others have argued that it is technologically feasible to copy the brain directly into hardware and software, and that such a simulation will be essentially identical to the original. This argument combines the idea that a suitably powerful machine can simulate any process, with the materialist idea that the mind is the result of physical processes in the brain.[53]

AI research

In the 21st century, AI research has become highly specialized and technical. It is deeply divided into subfields that often fail to communicate with each other.[10] Subfields have grown up around particular institutions, the work of particular researchers, particular problems, long standing differences of opinion about how AI should be done) and the application of widely differing tools.

Problems of AI

The problem of simulating (or creating) intelligence has been broken down into a number of specific sub-problems. These consist of particular traits or capabilities that researchers would like an intelligent system to display. The traits described below have received the most attention.[11]

Deduction, reasoning, problem solving

Early AI researchers developed algorithms that imitated the step-by-step reasoning that human beings use when they solve puzzles, play board games or make logical deductions.[54] By the late 80s and 90s, AI research had also developed highly successful methods for dealing with uncertain or incomplete information, employing concepts from probability and economics.[55]

For difficult problems, most of these algorithms can require enormous computational resources — most experience a "combinatorial explosion": the amount of memory or computer time required becomes astronomical when the problem goes beyond a certain size. The search for more efficient problem solving

(b)(3) - P.L. 86-36

algorithms is a high priority for AI research.[56]

Human beings solve most of their problems using fast, intuitive judgments rather than the conscious, step-by-step deduction that early AI research was able to model.[57] AI has made some progress at imitating this kind of "sub-symbolic" problem solving: embodied approaches emphasize the importance of sensorimotor skills to higher reasoning; neural net research attempts to simulate the structures inside human and animal brains that gives rise to this skill.

Knowledge representation

commonsense knowledge Main article: knowledge representation

Knowledge representation[58] and knowledge engineering[59] are central to AI research. Many of the problems machines are expected to solve will require extensive knowledge about the world. Among the things that AI needs to represent are: objects, properties, categories and relations between objects;[60] situations, events, states and time;[61] causes and effects;[62] knowledge about knowledge (what we know about what other people know);[63] and many other, less well researched domains. A complete representation of "what exists" is an ontology[64] (borrowing a word from traditional philosophy), of which the most general are called upper ontologies.

Among the most difficult problems in knowledge representation are:

Default reasoning and the qualification problem
: Many of the things people know take the form of "working assumptions." For example, if a bird comes up in conversation, people typically picture an animal that is fist sized, sings, and flies. None of these things are true about all birds. John McCarthy identified this problem in 1969[65] as the qualification problem: for any commonsense rule that AI researchers care to represent, there tend to be a huge number of exceptions. Almost nothing is simply true or false in the way that abstract logic requires. AI research has explored a number of solutions to this problem.[66]

The breadth of commonsense knowledge
: The number of atomic facts that the average person knows is astronomical. Research projects that attempt to build a complete knowledge base of commonsense knowledge (e.g., Cyc) require enormous amounts of laborious ontological engineering — they must be built, by hand, one complicated concept at a time.[67]

The subsymbolic form of some commonsense knowledge
: Much of what people know isn't represented as "facts" or "statements" that they could actually say out loud. For example, a chess master will avoid a particular chess position because it "feels too exposed"[68] or an art critic can take one look at a statue and instantly realize that it is a fake.[69] These are intuitions or tendencies that are represented in the brain non-consciously and sub-symbolically. Knowledge like this informs, supports and provides a context for symbolic, conscious knowledge. As with the related problem of sub-symbolic reasoning, it is hoped that situated AI or computational intelligence will provide ways to represent this kind of knowledge.[70]

(b)(3) - P.L. 86-36

Planning

Main article: automated planning and scheduling

Intelligent agents must be able to set goals and achieve them.[71] They need a way to visualize the future (they must have a representation of the state of the world and be able to make predictions about how their actions will change it) and be able to make choices that maximize the utility (or "value") of the available choices.[72]

In some planning problems, the agent can assume that it is the only thing acting on the world and it can be certain what the consequences of its actions may be.[73] However, if this is not true, it must periodically check if the world matches its predictions and it must change its plan as this becomes necessary, requiring the agent to reason under uncertainty.[74]

Multi-agent planning uses the cooperation and competition of many agents to achieve a given goal. Emergent behavior such as this is used by evolutionary algorithms and swarm intelligence.[75]

Learning

Main article: machine learning

Machine learning[76] has been central to AI research from the beginning.[77] Unsupervised learning is the ability to find patterns in a stream of input. Supervised learning includes both classification (be able to determine what category something belongs in, after seeing a number of examples of things from several categories) and regression (given a set of numerical input/output examples, discover a continuous function that would generate the outputs from the inputs). In reinforcement learning[78] the agent is rewarded for good responses and punished for bad ones. These can be analyzed in terms of decision theory, using concepts like utility. The mathematical analysis of machine learning algorithms and their performance is a branch of theoretical computer science known as computational learning theory.

Natural language processing

Main article: natural language processing

Natural language processing[79] gives machines the ability to read and understand the languages that the human beings speak. Many researchers hope that a sufficiently powerful natural language processing system would be able to acquire knowledge on its own, by reading the existing text available over the internet. Some straightforward applications of natural language processing include information retrieval (or text mining) and machine translation.[80]

Other Fields of Study

- Machine perception
  - Speech recognition:

(b)(3) - P.L. 86-36

- facial recognition and object recognition.**Cite error: Closing </ref> missing for <ref> tag**

Should artificial intelligence simulate natural intelligence, by studying human psychology or animal neurobiology? Or is human biology as irrelevant to AI research as bird biology is to aeronautical engineering?[81] Can intelligent behavior be described using simple, elegant principles (such as logic or optimization)? Or does artificial intelligence necessarily require solving many unrelated problems?[82]

Cybernetics and brain simulation

In the 40s and 50s, a number of researchers explored the connection between neurology, information theory, and cybernetics. Some of them built machines that used electronic networks to exhibit rudimentary intelligence, such as W. Grey Walter's turtles and the Johns Hopkins Beast. Many of these researchers gathered for meetings of the Teleological Society at Princeton University and the Ratio Club in England.[30]

Intelligent agent paradigm

The "intelligent agent" paradigm became widely accepted during the 1990s.[83] An intelligent agent is a system that perceives its environment and takes actions which maximizes its chances of success. The simplest intelligent agents are programs that solve specific problems. The most complicated intelligent agents are rational, thinking human beings.[84] The paradigm gives researchers license to study isolated problems and find solutions that are both verifiable and useful, without agreeing on one single approach. An agent that solves a specific problem can use any approach that works — some agents are symbolic and logical, some are sub-symbolic neural networks and others may use new approaches. The paradigm also gives researchers a common language to communicate with other fields—such as decision theory and economics—that also use concepts of abstract agents.

Integrating the approaches

An agent architecture or cognitive architecture allows researchers to build more versatile and intelligent systems out of interacting intelligent agents in a multi-agent system.[85] A system with both symbolic and sub-symbolic components is a hybrid intelligent system, and the study of such systems is artificial intelligence systems integration. A hierarchical control system provides a bridge between sub-symbolic AI at its lowest, reactive levels and traditional symbolic AI at its highest levels, where relaxed time constraints permit planning and world modelling.[86]

Tools of AI research

In the course of 50 years of research, AI has developed a large number of tools to solve the most difficult problems in computer science. A few of the most general of these methods are discussed below.

Search and optimization

(b)(3) - P.L. 86-36

Many problems in AI can be solved in theory by intelligently searching through many possible solutions:[87] Reasoning can be reduced to performing a search. For example, logical proof can be viewed as searching for a path that leads from premises to conclusions, where each step is the application of an inference rule.[88] Planning algorithms search through trees of goals and subgoals, attempting to find a path to a target goal, a process called means-ends analysis.[89] Robotics algorithms for moving limbs and grasping objects use local searches in configuration space.[90] Many learning algorithms use search algorithms based on optimization.

Simple exhaustive searches[91] are rarely sufficient for most real world problems: the search space (the number of places to search) quickly grows to astronomical numbers. The result is a search that is too slow or never completes. The solution, for many problems, is to use "heuristics" or "rules of thumb" that eliminate choices that are unlikely to lead to the goal (called "pruning the search tree"). Heuristics supply the program with a "best guess" for what path the solution lies on.[92]

A very different kind of search came to prominence in the 1990s, based on the mathematical theory of optimization. For many problems, it is possible to begin the search with some form of a guess and then refine the guess incrementally until no more refinements can be made. These algorithms can be visualized as blind hill climbing: we begin the search at a random point on the landscape, and then, by jumps or steps, we keep moving our guess uphill, until we reach the top. Other optimization algorithms are simulated annealing, beam search and random optimization.[93]

Evolutionary computation uses a form of optimization search. For example, they may begin with a population of organisms (the guesses) and then allow them to mutate and recombine, selecting only the fittest to survive each generation (refining the guesses). Forms of evolutionary computation include swarm intelligence algorithms (such as ant colony or particle swarm optimization)[94] and evolutionary algorithms (such as genetic algorithms[95] and genetic programming[96][97]).

Classifiers and statistical learning methods

The simplest AI applications can be divided into two types: classifiers ("if shiny then diamond") and controllers ("if shiny then pick up"). Controllers do however also classify conditions before inferring actions, and therefore classification forms a central part of many AI systems.

Classifiers[98] are functions that use pattern matching to determine a closest match. They can be tuned according to examples, making them very attractive for use in AI. These examples are known as observations or patterns. In supervised learning, each pattern belongs to a certain predefined class. A class can be seen as a decision that has to be made. All the observations combined with their class labels are known as a data set.

When a new observation is received, that observation is classified based on previous experience. A classifier can be trained in various ways; there are many statistical and machine learning approaches.

A wide range of classifiers are available, each with its strengths and weaknesses. Classifier performance depends greatly on the characteristics of the data to be classified. There is no single classifier that works best on all given problems; this is also referred to as the "no free lunch" theorem. Various empirical tests have been performed to compare classifier performance and to find the characteristics of data that

(b)(3) - P.L. 86-36

determine classifier performance. Determining a suitable classifier for a given problem is however still more an art than science.

The most widely used classifiers are the neural network,[99] kernel methods such as the support vector machine,[100] k-nearest neighbor algorithm,[101] Gaussian mixture model,[102] naive Bayes classifier,[103] and decision tree.[104] The performance of these classifiers have been compared over a wide range of classification tasks[105] in order to find data characteristics that determine classifier performance.

Applications of artificial intelligence

Main article: Applications of artificial intelligence

Artificial intelligence has successfully been used in a wide range of fields including medical diagnosis, stock trading, robot control, law, scientific discovery, video games and toys. Frequently, when a technique reaches mainstream use it is no longer considered artificial intelligence, sometimes described as the AI effect.[106] It may also become integrated into artificial life.

See also

- List of AI projects
- List of AI researchers
- List of emerging technologies
- List of basic artificial intelligence topics
- List of important AI publications

Notes

1. ↑ Glossary of Intelligence Terms and Definitions, June 1989 Section 1 Info (U//FOUO)
2. ↑ Poole, Mackworth & Goebel 1998, p. 1 (who use the term "computational intelligence" as a synonym for artificial intelligence). Other textbooks that define AI this way include Template:Harvtxt, and Template:Harvtxt (who prefer the term "rational agent") and write "The whole-agent view is now widely accepted in the field" Template:Harvy
3. ↑ This definition, in terms of goals, actions, perception and environment, is due to Template:Harvtxt. Other definitions also include knowledge and learning as additional criteria.
4. ↑ Although there is some controversy on this point (see Crevier 1993, p. 50), McCarthy states unequivocally "I came up with the term" in a c|net interview. (See Getting Machines to Think Like Us (http://news.com.com/Getting+machines+to+think+like+us/2008-11394_3-6090207.html) .)
5. ↑ See John McCarthy, What is Artificial Intelligence? (http://www-formal.stanford.edu/jmc/whatisai/whatisai.html)
6. ↑ 6.0 6.1 Dartmouth proposal:
   - McCarthy et al. 1955
7. ↑ 7.0 7.1 7.2 This is a central idea of Pamela McCorduck's *Machines That Think*. She writes: "I like to think of artificial intelligence as the scientific apotheosis of a veneralbe cultural tradition." *(McCorduck 2004, p. 34)* "Artificial intelligence in one form or another is an idea that has pervaded

 (b)(3) - P.L. 86-36

Western intellectual history, a dream in urgent need of being realized." *(McCorduck 2004, p. xviii)* "Our history is full of attempts—nutty, eerie, comical, earnest, legendary and real—to make artificial intelligences, to reproduce what is the essential us—bypassing the ordinary means. Back and forth between myth and reality, our imaginations supplying what our workshops couldn't, we have engaged for a long time in this odd form of self-reproduction." *(McCorduck 2004, p. 3)* She traces the desire back to its Hellenistic roots and calls it the urge to "forge the Gods." *(McCorduck 2004, p. 340-400)*

8. ↑ The optimism referred to includes the predictions of early AI researchers (see optimism in the history of AI) as well as the ideas of modern transhumanists such as Ray Kurzweil.
9. ↑ The "setbacks" referred to include the ALPAC report of 1966, the abandonment of perceptrons in 1970, the the Lighthill Report of 1973 and the collapse of the lisp machine market in 1987.
10. ↑ [10.0] [10.1] Fractioning of AI into subfields:
    - McCorduck 2004, pp. 421-425
11. ↑ [11.0] [11.1] This list of intelligent traits is based on the topics covered by the major AI textbooks, including:
    - Russell & Norvig 2003
    - Luger & Stubblefield 2004
    - Poole, Mackworth & Goebel 1998
    - Nilsson 1998.
12. ↑ General intelligence (strong AI) is discussed in popular introductions to AI:
    - Kurzweil 1999 and Kurzweil 2005
13. ↑ AI in Myth:
    - McCorduck 2004, p. 4-5
    - Russell & Norvig 2003, p. 939
14. ↑ Sacred statues as artificial intelligence:
    - Template:Harvtxt (statue of Amun)
    - Template:Harvtxt
15. ↑ These were the first machines to be believed to have true intelligence and consciousness. Hermes Trismegistus expressed the common belief that with these statues, craftsman had reproduced "the true nature of the gods", their *sensus* and *spiritus*. McCorduck makes the connection between sacred automatons and Mosaic law (developed around the same time), which expressly forbids the worship of robots *(McCorduck 2004, pp. 6-9)*
16. ↑ Needham 1986, p. 53
17. ↑ McCorduck 2004, p. 6
18. ↑ A Thirteenth Century Programmable Robot (http://www.shef.ac.uk/marcoms/eview/articles58/robot.html)
19. ↑ McCorduck 2004, p. 17
20. ↑ Takwin: O'Connor, Kathleen Malone. "*The alchemical creation of life (takwin) and other concepts of Genesis in medieval Islam (http://repository.upenn.edu/dissertations/AAI9503804)* ". University of Pennsylvania. URL accessed on 2007-01-10.
21. ↑ Golem: McCorduck 2004, p. 15-16, Buchanan 2005, p. 50
22. ↑ McCorduck 2004, p. 13-14
23. ↑ Template:Harvtxt discusses *Frankenstein* and identifies the key ethical issues as scientific hubris and the suffering of the monster, i.e. robot rights.
24. ↑ Robot rights:

  - Russell & Norvig 2003, p. 964
  - Robots could demand legal rights (http://news.bbc.co.uk/2/hi/technology/6200005.stm)

25. ↑ See the Times Online, Human rights for robots? We're getting carried away (http://www.timesonline.co.uk/tol/news/uk/science/article1695546.ece)
26. ↑ Template:Harvtxt
27. ↑ 27.0 27.1 27.2 Singularity, transhumanism:
    - Kurzweil 2005
    - Russell & Norvig 2003, p. 963
28. ↑ Joseph Weizenbaum's critique of AI:
    - Weizenbaum 1976
    - Crevier 1993, pp. 132–144
    - McCorduck 2004, pp. 356-373
    - Russell & Norvig 2003, p. 961

    Weizenbaum (the AI researcher who developed the first chatterbot program, ELIZA) argued in 1976 that the misuse of artificial intelligence has the potential to devalue human life.
29. ↑ Quoted in Template:Harvtxt
30. ↑ 30.0 30.1 AI's immediate precursors:
    - McCorduck 2004, pp. 51-107
    - Crevier 1993, pp. 27-32
    - Russell & Norvig 2003, pp. 15,940
    - Moravec 1988, p. 3

    Among the researchers who laid the foundations of the theory of computation, cybernetics, information theory and neural networks were Alan Turing, John Von Neumann, Norbert Weiner, Claude Shannon, Warren McCullough, Walter Pitts and Donald Hebb
31. ↑ Dartmouth conference:
    - McCorduck, pp. 111-136
    - Crevier 1993, pp. 47-49
    - Russell & Norvig 2003, p. 17
    - NRC 1999, pp. 200-201
32. ↑ Russell and Norvig write "it was astonishing whenever a computer did anything kind of smartish." Russell & Norvig 2003, p. 18
33. ↑ "Golden years" of AI (successful symbolic reasoning programs 1956-1973):
    - McCorduck, pp. 243-252
    - Crevier 1993, pp. 52-107
    - Moravec 1988, p. 9
    - Russell & Norvig 2003, p. 18-21

    The programs described are Daniel Bobrow's STUDENT, Newell and Simon's Logic Theorist and Terry Winograd's SHRDLU.
34. ↑ DARPA pours money into undirected pure research into AI during the 1960s:
    - McCorduck 2005, pp. 131
    - Crevier 1993, pp. 51, 64-65
    - NRC 1999, pp. 204-205
35. ↑ Simon 1965, p. 96 quoted in Crevier 1993, p. 109
36. ↑ Minsky 1967, p. 2 quoted in Crevier 1993, p. 109
37. ↑ See History of artificial intelligence — the problems.

 (b)(3) - P.L. 86-36

38. ↑ First AI Winter:
    - Crevier 1993, pp. 115-117
    - Russell & Norvig 2003, p. 22
    - NRC 1999, pp. 212-213
    - Howe 1994
39. ↑ Expert systems:
    - ACM 1998, I.2.1,
    - Russell & Norvig 2003, pp. 22–24
    - Luger & Stubblefield 2004, pp. 227-331,
    - Nilsson 1998, chpt. 17.4
    - McCorduck 2004, pp. 327-335, 434-435
    - Crevier 1993, pp. 145-62, 197–203
40. ↑ Boom of the 1980s: rise of expert systems, Fifth Generation Project, Alvey, MCC, SCI:
    - McCorduck 2004, pp. 426-441
    - Crevier 1993, pp. 161-162,197-203, 211, 240
    - Russell & Norvig 2003, p. 24
    - NRC 1999, pp. 210-211
41. ↑ Second AI Winter:
    - McCorduck 2004, pp. 430-435
    - Crevier 1993, pp. 209-210
    - NRC 1999, pp. 214-216
42. ↑ AI applications widely used behind the scenes:
    - Russell & Norvig 2003, p. 28
    - Kurzweil 2005, p. 265
    - NRC 1999, pp. 216-222
43. ↑ Formal methods are now preferred ("Victory of the neats"):
    - Russell & Norvig 2003, pp. 25-26
    - McCorduck 2004, pp. 486-487
44. ↑ All of these positions below are mentioned in standard discussions of the subject, such as:
    - Russell & Norvig 2003, pp. 947-960
    - Fearn 2007, pp. 38-55
45. ↑ Philosophical implications of the Turing test:
    - Turing 1950,
    - Haugeland 1985, pp. 6-9,
    - Crevier 1993, p. 24,
    - Russell & Norvig 2003, pp. 2-3 and 948
46. ↑ The physical symbol systems hypothesis:
    - Newell & Simon 1976, p. 116
    - Russell & Norvig 2003, p. 18
47. ↑ Dreyfus criticized the necessary condition of the physical symbol system hypothesis, which he called the "psychological assumption": "The mind can be viewed as a device operating on bits of information according to formal rules". *(Dreyfus 1992, p. 156)*
48. ↑ Dreyfus' Critique of AI:
    - Dreyfus 1972,
    - Dreyfus & Dreyfus 1986,
    - Russell & Norvig 2003, pp. 950-952,

  - Crevier 1993, pp. 120-132 and
49. ↑ This is a paraphrase of the important implication of Gödel's theorems.
50. ↑ The Mathematical Objection:
    - Russell & Norvig 2003, p. 949
    - McCorduck 2004, p. 448-449

    Refuting Mathematical Objection:
    - Turing 1950 under "(2) The Mathematical Objection"
    - Hofstadter 1979,

    Making the Mathematical Objection:
    - Lucas 1961,
    - Penrose 1989.

    Background:
    - Gödel 1931, Church 1936, Kleene 1935, Turing 1937,
51. ↑ This version is from Template:Harvtxt, and is also quoted in Dennett 1991, p. 435. Searle's original formulation was "The appropriately programmed computer really is a mind, in the sense that computers given the right programs can be literally said to understand and have other cognitive states." *(Searle 1980, p. 1)*. Strong AI is defined similarly by Template:Harvtxt: "The assertion that machines could possibly act intelligently (or, perhaps better, act as if they were intelligent) is called the 'weak AI' hypothesis by philosophers, and the assertion that machines that do so are actually thinking (as opposed to simulating thinking) is called the 'strong AI' hypothesis."
52. ↑ Searle's Chinese Room argument:
    - Searle 1980, Searle 1991
    - Russell & Norvig 2003, pp. 958-960
    - McCorduck 2004, pp. 443-445
    - Crevier 1993, pp. 269-271
53. ↑ Artificial brain:
    - Moravec 1988
    - Kurzweil 2005, p. 262
    - Russell Norvig, p. 957
    - Crevier 1993, pp. 271 and 279

    The most extreme form of this argument (the brain replacement scenario) was put forward by Clark Glymour in the mid-70s and was touched on by Zenon Pylyshyn and John Searle in 1980. Daniel Dennett sees human consciousness as multiple functional thought patterns; see "Consciousness Explained."
54. ↑ Problem solving, puzzle solving, game playing and deduction:
    - Russell & Norvig 2003, chpt. 3-9,
    - Poole et al. chpt. 2,3,7,9,
    - Luger & Stubblefield 2004, chpt. 3,4,6,8,
    - Nilsson, chpt. 7-12.
55. ↑ Uncertain reasoning:
    - Russell & Norvig 2003, pp. 452-644,
    - Poole, Mackworth & Goebel 1998, pp. 345-395,
    - Luger & Stubblefield 2004, pp. 333-381,
    - Nilsson 1998, chpt. 19
56. ↑ Intractability and efficiency and the combinatorial explosion:
    - Russell & Norvig 2003, pp. 9, 21-22

57. ↑ Cognitive science has provided several famous examples:
    - Template:Harvtxt showed that people do poorly on completely abstract problems, but if the problem is restated to allow the use of intuitive social intelligence, performance dramatically improves. (See Wason selection task)
    - Template:Harvtxt have shown that people are terrible at elementary problems that involve uncertain reasoning. (See list of cognitive biases for several examples).
    - Template:Harvtxt have controversially argued that even our skills at mathematics depend on knowledge and skills that come from "the body", i.e. sensorimotor and perceptual skills. (See Where Mathematics Comes From)
58. ↑ Knowledge representation:
    - ACM 1998, I.2.4,
    - Russell & Norvig 2003, pp. 320-363,
    - Poole, Mackworth & Goebel 1998, pp. 23-46, 69-81, 169-196, 235-277, 281-298, 319-345,
    - Luger & Stubblefield 2004, pp. 227-243,
    - Nilsson 1998, chpt. 18
59. ↑ Knowledge engineering:
    - Russell & Norvig 2003, pp. 260-266,
    - Poole, Mackworth & Goebel 1998, pp. 199-233,
    - Nilsson 1998, chpt. ~17.1-17.4
60. ↑ Representing categories and relations: Semantic networks, description logics, inheritance (including frames and scripts):
    - Russell & Norvig 2003, pp. 349-354,
    - Poole, Mackworth & Goebel 1998, pp. 174-177,
    - Luger & Stubblefield 2004, pp. 248-258,
    - Nilsson 1998, chpt. 18.3
61. ↑ Representing events and time:Situation calculus, event calculus, fluent calculus (including solving the frame problem):
    - Russell & Norvig 2003, pp. 328-341,
    - Poole, Mackworth & Goebel 1998, pp. 281-298,
    - Nilsson 1998, chpt. 18.2
62. ↑ Causal calculus:
    - Poole, Mackworth & Goebel 1998, pp. 335-337
63. ↑ Representing knowledge about knowledge: Belief calculus, modal logics:
    - Russell & Norvig 2003, pp. 341-344,
    - Poole, Mackworth & Goebel 1998, pp. 275-277
64. ↑ Ontology:
    - Russell & Norvig 2003, pp. 320-328
65. ↑ McCarthy & Hayes 1969. While McCarthy was primarily concerned with issues in the logical representation of actions, Russell & Norvig 2003 apply the term to the more general issue of default reasoning in the vast network of assumptions underlying all our commonsense knowledge.
66. ↑ Default reasoning and default logic, non-monotonic logics, circumscription, closed world assumption, abduction (Poole *et al.* places abduction under "default reasoning". Luger *et al.* places this under "uncertain reasoning"):
    - Russell & Norvig 2003, pp. 354-360,
    - Poole, Mackworth & Goebel 1998, pp. 248-256, 323-335,
    - Luger & Stubblefield 2004, pp. 335-363,

 (b)(3) - P.L. 86-36 [redacted]

- Nilsson 1998, ~18.3.3

67. ↑ Breadth of commonsense knowledge:
    - Russell & Norvig 2003, p. 21,
    - Crevier 1993, pp. 113-114,
    - Moravec 1988, p. 13,
    - Lenat & Guha 1989 (Introduction)
68. ↑ Dreyfus & Dreyfus 1986
69. ↑ Gladwell 2005
70. ↑ Expert knowledge as embodied intuition:
    - Dreyfus & Dreyfus 1986 (Hubert Dreyfus is a philosopher and critic of AI who was among the first to argue that most useful human knowledge was encoded sub-symbolically.)
    - Gladwell 2005 (Gladwell's *Blink* is a popular introduction to sub-symbolic reasoning and knowledge.)
    - Hawkins 2005 (Hawkins argues that sub-symbolic knowledge should be the primary focus of AI research.)
71. ↑ Planning:
    - ACM 1998, ~I.2.8,
    - Russell & Norvig 2003, pp. 375-459,
    - Poole, Mackworth & Goebel 1998, pp. 281-316,
    - Luger & Stubblefield 2004, pp. 314-329,
    - Nilsson 1998, chpt. 10.1-2, 22
72. ↑ **Cite error: Invalid <ref> tag; no text was provided for refs named IVT**
73. ↑ Classical planning:
    - Russell & Norvig 2003, pp. 375-430,
    - Poole, Mackworth & Goebel 1998, pp. 281-315,
    - Luger & Stubblefield 2004, pp. 314-329,
    - Nilsson 1998, chpt. 10.1-2, 22
74. ↑ Planning and acting in non-deterministic domains: conditional planning, execution monitoring, replanning and continuous planning:
    - Russell & Norvig 2003, pp. 430-449
75. ↑ Multi-agent planning and emergent behavior:
    - Russell & Norvig 2003, pp. 449-455
76. ↑ Learning:
    - ACM 1998, I.2.6,
    - Russell & Norvig 2003, pp. 649-788,
    - Poole, Mackworth & Goebel 1998, pp. 397-438,
    - Luger & Stubblefield 2004, pp. 385-542,
    - Nilsson 1998, chpt. 3.3 , 10.3, 17.5, 20
77. ↑ Alan Turing discussed the centrality of learning as early as 1950, in his classic paper Computing Machinery and Intelligence. *(Turing 1950)*
78. ↑ Reinforcement learning:
    - Russell & Norvig 2003, pp. 763-788
    - Luger & Stubblefield 2004, pp. 442-449
79. ↑ Natural language processing:
    - ACM 1998, I.2.7
    - Russell & Norvig 2003, pp. 790-831

 (b)(3) - P.L. 86-36

- Poole, Mackworth & Goebel 1998, pp. 91-104
- Luger & Stubblefield 2004, pp. 591-632

80. ↑ Applications of natural language processing, including information retrieval (i.e. text mining) and machine translation:
    - Russell & Norvig 2003, pp. 840-857,
    - Luger & Stubblefield 2004, pp. 623-630
81. ↑ The analogy with aeronautical engineering is due to Template:Harvtxt.
82. ↑ Neats vs. scruffies:
    - McCorduck 2004, pp. 421-424, 486-489
    - Crevier 1993, pp. 168
83. ↑ "The whole-agent view is now widely accepted in the field" Russell & Norvig 2003, p. 55.
84. ↑ The intelligent agent paradigm:
    - Russell & Norvig 2003, pp. 27, 32-58, 968-972,
    - Poole, Mackworth & Goebel 1998, pp. 7-21,
    - Luger & Stubblefield 2004, pp. 235-240
85. ↑ Agent architectures, hybrid intelligent systems:
    - Template:Harvtxt
    - Template:Harvtxt
86. ↑ Albus, J. S. 4-D/RCS reference model architecture for unmanned ground vehicles. In G Gerhart, R Gunderson, and C Shoemaker, editors, Proceedings of the SPIE AeroSense Session on Unmanned Ground Vehicle Technology, volume 3693, pages 11—20
87. ↑ Search algorithms:
    - Russell & Norvig 2003, pp. 59-189
    - Poole, Mackworth & Goebel 1998, pp. 113-163
    - Luger & Stubblefield 2004, pp. 79-164, 193-219
    - Nilsson 1998, chpt. 7-12
88. ↑ Forward chaining, backward chaining, Horn clauses, and logical deduction as search:
    - Russell & Norvig 2003, pp. 217-225, 280-294
    - Poole, Mackworth & Goebel 1998, pp. ~46-52
    - Luger & Stubblefield 2004, pp. 62-73
    - Nilsson 1998, chpt. 4.2, 7.2
89. ↑ State space search and planning:
    - Russell & Norvig 2003, pp. 382-387
    - Poole, Mackworth & Goebel 1998, pp. 298-305
    - Nilsson 1998, chpt. 10.1-2
90. ↑ **Cite error: Invalid `<ref>` tag; no text was provided for refs named `cs`**
91. ↑ Uninformed searches (breadth first search, depth first search and general state space search):
    - Russell & Norvig 2003, pp. 59-93
    - Poole, Mackworth & Goebel 1998, pp. 113-132
    - Luger & Stubblefield 2004, pp. 79-121
    - Nilsson 1998, chpt. 8
92. ↑ Heuristic or informed searches (e.g., greedy best first and A*):
    - Russell & Norvig 2003, pp. 94-109,
    - Poole, Mackworth & Goebel 1998, pp. pp. 132-147,
    - Luger & Stubblefield 2004, pp. 133-150,
    - Nilsson 1998, chpt. 9

93. ↑ Optimization searches:
    - Russell & Norvig 2003, pp. 110-116,120-129
    - Poole, Mackworth & Goebel 1998, pp. 56-163
    - Luger & Stubblefield 2004, pp. 127-133
94. ↑ Artificial life and society based learning:
    - Luger & Stubblefield 2004, pp. 530-541
95. ↑ Genetic algorithms for learning:
    - Luger & Stubblefield 2004, pp. 509-530,
    - Nilsson 1998, chpt. 4.2.

    See also: Holland, John H. (1975). *Adaptation in Natural and Artificial Systems.* University of Michigan Press.
96. ↑ Koza, John R. (1992). *Genetic Programming.* MIT Press.
97. ↑ Poli, R., Langdon, W. B., McPhee, N. F. (2008). *A Field Guide to Genetic Programming.* Lulu.com, freely available from http://www.gp-field-guide.org.uk/.
98. ↑ Statistical learning methods and classifiers:
    - Russell & Norvig 2003, pp. 712-754,
    - Luger & Stubblefield 2004, pp. 453-541
99. ↑ **Cite error: Invalid `<ref>` tag; no text was provided for refs named NN**
100. ↑ Kernel methods:
    - Russell & Norvig 2003, pp. 749-752
101. ↑ K-nearest neighbor algorithm:
    - Russell & Norvig 2003, pp. 733-736
102. ↑ Gaussian mixture model:
    - Russell & Norvig 2003, pp. 725-727
103. ↑ Naive Bayes classifier:
    - Russell & Norvig 2003, pp. 718
104. ↑ Decision tree:
    - Russell & Norvig 2003, pp. 653-664,
    - Poole, Mackworth & Goebel 1998, pp. 403-408,
    - Luger & Stubblefield 2004, pp. 408-417
105. ↑ Template:Cite-web
106. ↑ AI set to exceed human brain power (http://www.cnn.com/2006/TECH/science/07/24/ai.bostrom/) (web article). CNN.com (2006-07-26). Retrieved on 2008-02-26.

References

Major AI textbooks

See also A.I. Textbook survey

- Luger, George & Stubblefield, William (2004), *Artificial Intelligence: Structures and Strategies for Complex Problem Solving (http://www.cs.unm.edu/~luger/ai-final/tocfull.html)* (5th ed.), The Benjamin/Cummings Publishing Company, Inc., ISBN 0-8053-4780-1, <http://www.cs.unm.edu/~luger/ai-final/tocfull.html>
- Nilsson, Nils (1998), *Artificial Intelligence: A New Synthesis*, Morgan Kaufmann Publishers, ISBN 978-1-55860-467-4

- Template:Russell Norvig 2003
- Poole, David; Mackworth, Alan & Goebel, Randy (1998), *Computational Intelligence: A Logical Approach (http://www.cs.ubc.ca/spider/poole/ci.html)* , New York: Oxford University Press, <http://www.cs.ubc.ca/spider/poole/ci.html>
- Winston, Patrick Henry (1984), *Artificial Intelligence*, Reading, Massachusetts: Addison-Wesley, ISBN 0201082594

History of AI

- Template:Crevier 1993
- Template:McCorduck 2004

Other sources

- ACM, (Association of Computing Machinery) (1998), *ACM Computing Classification System: Artificial intelligence (http://www.acm.org/class/1998/I.2.html)* , <http://www.acm.org/class/1998/I.2.html>
- Brooks, Rodney (1990), "Elephants Don't Play Chess " (PDF), *Robotics and Autonomous Systems* 6: 3–15, doi:10.1016/S0921-8890(05)80025-9 (http://dx.doi.org/10.1016%2FS0921-8890%2805%2980025-9) , <http://people.csail.mit.edu/brooks/papers/elephants.pdf >. Retrieved on 30 August 2007
- Buchanan, Bruce G. (Winter 2005), "A (Very) Brief History of Artificial Intelligence " (PDF), *AI Magazine*: 53–60, <http://www.aaai.org/AITopics/assets/PDF/AIMag26-04-016.pdf >. Retrieved on 30 August 2007
- Dreyfus, Hubert (1972), *What Computers Can't Do*, New York: MIT Press, ISBN 0060110821
- Dreyfus, Hubert (1979), *What Computers* Still *Can't Do*, New York: MIT Press.
- Dreyfus, Hubert & Dreyfus, Stuart (1986), *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*, Oxford, UK: Blackwell.
- Gladwell, Malcolm (2005), *Blink*, New York: Little, Brown and Co., ISBN 0-316-17232-4.
- Haugeland, John (1985), *Artificial Intelligence: The Very Idea*, Cambridge, Mass.: MIT Press, ISBN 0-262-08153-9.
- Hawkins, Jeff & Blakeslee, Sandra (2004), *On Intelligence*, New York, NY: Owl Books, ISBN 0-8050-7853-3.
- Hofstadter, Douglas (1979), *Gödel, Escher, Bach: an Eternal Golden Braid.*
- Howe, J. (November 1994), *Artificial Intelligence at Edinburgh University: a Perspective (http://www.inf.ed.ac.uk/about/AIhistory.html)* , <http://www.inf.ed.ac.uk/about/AIhistory.html>.
- Kahneman, Daniel; Slovic, D. & Tversky, Amos (1982), *Judgment under uncertainty: Heuristics and biases*, New York: Cambridge University Press.
- Kurzweil, Ray (1999), *The Age of Spiritual Machines*, Penguin Books, ISBN 0-670-88217-8
- Kurzweil, Ray (2005), *The Singularity is Near*, Penguin Books, ISBN 0-670-03384-7
- Lakoff, George (1987), *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*, University of Chicago Press., ISBN 0-226-46804-6
- Lakoff, George & Núñez, Rafael E. (2000), *Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being*, Basic Books, ISBN 0-465-03771-2.
- Lenat, Douglas & Guha, R. V. (1989), *Building Large Knowledge-Based Systems*, Addison-Wesley

 (b)(3) - P.L. 86-36

- Lighthill, Professor Sir James (1973), "Artificial Intelligence: A General Survey", *Artificial Intelligence: a paper symposium*, Science Research Council
- Lucas, John (1961), "Minds, Machines and Gödel" (http://users.ox.ac.uk/~jrlucas/Godel/mmg.html) , in Anderson, A.R., *Minds and Machines*, <http://users.ox.ac.uk/~jrlucas/Godel/mmg.html>.
- McCarthy, John; Minsky, Marvin; Rochester, Nathan & Shannon, Claude (1955), *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence (http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html)* , <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>.
- McCarthy, John & Hayes, P. J. (1969), "Some philosophical problems from the standpoint of artificial intelligence (http://www-formal.stanford.edu/jmc/mcchay69.html) ", *Machine Intelligence* **4**: 463–502, <http://www-formal.stanford.edu/jmc/mcchay69.html>
- Minsky, Marvin (1967), *Computation: Finite and Infinite Machines*, Englewood Cliffs, N.J.: Prentice-Hall
- Minsky, Marvin (2006), *The Emotion Machine*, New York, NY: Simon & Schusterl, ISBN 0-7432-7663-9
- Moravec, Hans (1976), *The Role of Raw Power in Intelligence (http://www.frc.ri.cmu.edu/users/hpm/project.archive/general.articles/1975/Raw.Power.html)* , <http://www.frc.ri.cmu.edu/users/hpm/project.archive/general.articles/1975/Raw.Power.html>
- Moravec, Hans (1988), *Mind Children*, Harvard University Press
- NRC (1999), "Developments in Artificial Intelligence", *Funding a Revolution: Government Support for Computing Research*, National Academy Press
- Needham, Joseph (1986), *Science and Civilization in China: Volume 2*, Caves Books Ltd.
- Newell, Allen & Simon, H. A. (1963), "GPS: A Program that Simulates Human Thought", in Feigenbaum, E.A. & Feldman, J., *Computers and Thought*, McGraw-Hill
- Newell, Allen & Simon, H. A. (1976), "Computer Science as Empirical Inquiry: Symbols and Search" (http://www.rci.rutgers.edu/~cfs/472_html/AI_SEARCH/PSS/PSSH4.html) , *Communications of the ACM*, **19**, <http://www.rci.rutgers.edu/~cfs/472_html/AI_SEARCH/PSS/PSSH4.html>
- Searle, John (1980), "Minds, Brains and Programs (http://www.bbsonline.org/documents/a/00/00/04/84/bbs00000484-00/bbs.searle2.html) ", *Behavioral and Brain Sciences* **3**(3): 417–457, <http://www.bbsonline.org/documents/a/00/00/04/84/bbs00000484-00/bbs.searle2.html>
- Searle, John (1999), *Mind, language and society*, New York, NY: Basic Books, ISBN 0465045219, OCLC 231867665 43689264 (http://worldcat.org/oclc/231867665+43689264)
- Shapiro, Stuart C. (1992), "Artificial Intelligence" (http://www.cse.buffalo.edu/~shapiro/Papers/ai.ps) , in Shapiro, Stuart C., *Encyclopedia of Artificial Intelligence* (2nd ed.), New York: John Wiley, pp. 54–57, <http://www.cse.buffalo.edu/~shapiro/Papers/ai.ps>.
- Simon, H. A. (1965), *The Shape of Automation for Men and Management*, New York: Harper & Row
- Template:Turing 1950
- Wason, P. C. (1966), "Reasoning", in Foss, B. M., *New horizons in psychology*, Harmondsworth: Penguin
- Weizenbaum, Joseph (1976), *Computer Power and Human Reason*, San Francisco: W.H. Freeman & Company, ISBN 0716704641

Further reading

- R. Sun & L. Bookman, (eds.), *Computational Architectures: Integrating Neural and Symbolic Processes*. Kluwer Academic Publishers, Needham, MA. 1994.
- Margaret Boden, Mind As Machine, Oxford University Press, 2006
- John Johnston, (2008) "The Allure of Machinic Life: Cybernetics, Artificial Life, and the New AI", MIT Press
- Michaela Ong, pioneer of Artificial Intelligence. "Origin of AI through cs191"

External links

- The Association for the Advancement of Artificial Intelligence (AAAI) - AI Topics (http://www.aaai.org/AITopics/)
- Freeview Video 'Machines with Minds' by the Vega Science Trust and the BBC/OU (http://www.vega.org.uk/video/programme/16)
- John McCarthy's frequently asked questions about AI (http://www-formal.stanford.edu/jmc/whatisai/whatisai.html)
- The Futurist magazine interviews "Ai chasers" Rodney Brooks, Peter Norvig, Barney Pell, et al. (http://www.wfs.org/Dec-janfiles/AIInt.htm)
- Jonathan Edwards looks at AI (BBC audio) (http://www.aiai.ed.ac.uk/events/jonathanedwards2007/bbc-r4-jonathan-edwards-2007-03-28.mp3) C
- Ray Kurzweil's website dedicated to AI including prediction of future development in AI (http://www.kurzweilai.net/)
- Template:Sep entry

(b)(3) - P.L. 86-36

Retrieved from [redacted]

Categories: Computer Science | Intelligence Terms and Definitions | Artificial intelligence

UNCLASSIFIED

(b)(6)

- This page has been accessed 2,424 times.
- 4 [redacted] watching users
- This page was last modified 14:05, 16 April 2014 by [redacted] Most recent editors: [redacted] and others.

(U) Artificial intelligence

UNCLASSIFIED

From Intellipedia

You have new messages (last change).

"AI" redirects here; for other uses of "AI" and "Artificial intelligence", see AI (disambiguation).

Contents

Background

The modern definition of **artificial intelligence** (or **AI**) is "the study and design of intelligent agents where an intelligent agent is a system that perceives its environment and takes actions which maximizes its chances of success. One of the most popular textbooks on AI by Russell and Norvig (2002) defines AI as "the science and engineering of making intelligent machines.

The term **artificial intelligence** is also used to describe the intelligent properties that a system demonstrates. Among the traits that researchers hope machines will exhibit are reasoning, knowledge, planning, learning, communication, perception and the ability to move and manipulate objects.

When AI was born, researchers had hopes of developing a computer with *general intelligence*. This led to a number of attempts to define general (human-like) intelligence in various ways (see, for example, the Turing Test). This "strong AI" was not achieved, and the field of AI has moved on to contribute applications across a wide range of areas including augmented cognition, ubiquitous technology, e-learning, semantic web, agent-based technologies, and intelligent distributed collaborative technology.

(b)(3) - P.L. 86-36

AI research uses tools and insights from many fields, including computer science, psychology, philosophy, neuroscience, cognitive science, linguistics, operations research, economics, control theory, probability, optimization and logic.[1] AI research also overlaps with tasks such as robotics, control systems, scheduling, data mining, logistics, speech recognition, facial recognition and many others.[2]

Perspectives on AI

History of AI

Some of the leaders of AI research include John McCarthy, Marvin Minsky, Allen Newell and Herbert Simon, who founded AI laboratories at MIT, CMU and Stanford. They and their students wrote programs that were, to most people, simply astonishing. These programs solved word problems in algebra, proved logical theorems, and even spoke English!

By the mid 60s, AI research was heavily funded by DARPA[3] and the researchers were optimistic about the future of their new field:

- 1965, Herb A. Simon: "[M]achines will be capable, within twenty years, of doing any work a man can do"[4]
- 1967, Marvin Minsky: "Within a generation ... the problem of creating 'artificial intelligence' will substantially be solved."[5]

These predictions, and many like them, would not come true because the field had failed to recognize the difficulty of some of the problems they faced. In the early 80s, the field was revived by the commercial success of expert systems and by 1985 the market for AI had reached more than a billion dollars.[6]

In the 90s and early 21st century AI achieved its greatest successes, albeit somewhat behind the scenes. Artificial intelligence was adopted throughout the technology industry. It evolved through its application in fields as diverse as logistics, data mining, medical diagnosis, intelligent tutoring systems, neuroscience, agent-based technology and many other areas.

The success was due to several factors: the incredible power of computers today (see Moore's law), a greater emphasis on solving specific subproblems, the creation of new ties between AI and other fields working on similar problems, and above all a new commitment by researchers to solid mathematical methods and rigorous scientific standards.[7]

Philosophy of AI

The philosophy of artificial intelligence considers the question "Can machines think?" Alan Turing, in his classic 1950 paper, *Computing Machinery and Intelligence*, was the first to try to answer it.

- The premise of the Turing Test is as follows: *if a machine acts as intelligently as a human being, then it is as intelligent as a human being.* The test involved a human judge who engages in a natural language conversation with both another human and a machine. If the judge cannot reliably tell which is the human and which is the computer, then the computer has passed the test.
- The Dartmouth proposal: *Every aspect of learning or any other feature of intelligence can be so precisely described that a machine can be made to simulate it.* This assertion was printed in the program for the Dartmouth Conference of 1956, and represents the position of most working AI researchers.[8]

(b)(3) - P.L. 86-36

- Newell and Simon's physical symbol system hypothesis: *A physical symbol system has the necessary and sufficient means of general intelligent action.* This statement claims that essence of intelligence is symbol manipulation.[9]
- The artificial brain argument: *The brain can be simulated.* This argument combines the idea that a Turing complete machine can simulate any process, with the materialist idea that the mind is the result of a physical process in the brain.[10]
- Gödel's incompleteness theorem: *There are statements that no physical symbol system can prove.* Roger Penrose is among those who claim that Gödel's theorem limits what machines can do.[11]
- Dreyfus' "psychological assumption": *The mind can be viewed as a device operating on bits of information according to formal rules.* Dreyfus refuted this statement by showing that human expertise depends on unconscious instinct rather than conscious symbol manipulation and on having a "feel" for the situation rather explicit symbolic knowledge.[12]
- Searle's "strong AI position": *A physical symbol system can have a mind and mental states.* Searle refuted this with his Chinese room argument, which asks us to look *inside* the computer and try to find where the "mind" might be.[13]

AI Research Focus Areas

This section describes the many different areas in which AI work is done. Over the years, each of these areas has grown tremendously, and today one can find international conferences devoted to each topic. Often however large projects in interdisciplinary areas such as robotics, adaptive software, and collaboration technology combine research from many of these areas.

Deduction, reasoning, problem solving

Early AI researchers developed algorithms that imitated the process of conscious, step-by-step reasoning that human beings use when they solve puzzles, play board games, or make logical deductions.[14] These early methods were unable to handle incomplete or imprecise information but by the late 80s and 90s, AI research developed highly successful methods for dealing with uncertainty, employing concepts from probability and economics.[15]

For difficult problems, most of these algorithms can require enormous computational resources — most experience a "combinatorial explosion": the amount of memory or computer time required becomes astronomical when the problem goes beyond a certain size. The search for more efficient problem solving algorithms is a high priority for AI research.[16]

It is not clear, however, that conscious human reasoning is any more efficient when faced with a difficult abstract problem. Cognitive scientists have demonstrated that human beings solve most of their problems using unconscious reasoning, rather than the conscious, step-by-step deduction that early AI research was able to model.[17] For many problems, people seem to simply jump to the correct solution: they think "instinctively" and "unconsciously". These instincts seem to involve skills usually applied to other problems, such as motion and manipulation (our so-called "embodied" skills that allow us deal with the physical world) or perception (for example, our skills at pattern matching). It is hoped that sub-symbolic methods, like computational intelligence and situated AI, will be able to model these instinctive skills. The problem of unconscious problem solving, which forms part of our commonsense reasoning, is largely unsolved.

(b)(3) - P.L. 86-36

Knowledge representation

Knowledge representation[18] and knowledge engineering[19] are central to AI research. Many of the problems machines are expected to solve will require extensive knowledge about the world. Among the things that AI needs to represent are: objects, properties, categories and relations between objects;[20] situations, events, states and time;[21] causes and effects;[22] knowledge about knowledge (what we know about what other people know);[23] and many other, less well researched domains. A complete representation of "what exists" is an ontology[24] (borrowing a word from traditional philosophy). Ontological engineering is the science of finding a general representation that can handle all of human knowledge.

Among the most difficult problems in knowledge representation are:

- *Default reasoning and the qualification problem*: Many of the things people know take the form of "working assumptions." For example, if a bird comes up in conversation, people typically picture a animal that is fist sized, sings, and flies. None of these things are true about birds in general. John McCarthy identified this problem in 1969[25] as the qualification problem: for any commonsense rule that AI researchers care to represent, there tend to be a huge number of exceptions. Almost nothing is simply true or false in the way that abstract logic requires. AI research has explored a number of solutions to this problem.[26]
- *Unconscious knowledge*: Much of what people know isn't represented as "facts" or "statements" that they could actually say out loud. They take the form of intuitions or tendencies and are represented in the brain unconsciously and sub-symbolically. This unconscious knowledge informs, supports and provides a context for our conscious knowledge. As with the related problem of unconscious reasoning, it is hoped that situated AI or computational intelligence will provide ways to represent this kind of knowledge.
- *The breadth of common sense knowledge*: The number of atomic facts that the average person knows is astronomical. Research projects that attempt to build a complete knowledge base of commonsense knowledge, such as Cyc, require enormous amounts of tedious step-by-step ontological engineering — they must be built, by hand, one complicated concept at a time.[27]

Planning

Intelligent agents must be able set goals and achieve them.[28] They need a way to visualize the future: they must have a representation of the state of the world and be able to make predictions about how their actions will change it. There are several types of planning problems:

- Classical planning problems assume that the agent is the only thing acting on the world, and that the agent can be certain what the consequences of it's actions may be.[29] Partial order planning problems take into account the fact that sometimes it's not important which sub-goal the agent achieves first.[30]
- If the environment is changing, or if the agent can't be sure of the results of its actions, it must periodically check if the world matches its predictions (conditional planning and execution monitoring) and it must change its plan as this becomes necessary (replanning and continuous planning).[31]
- Some planning problems take into account the utility or "usefulness" of a given outcome. These problems can be analyzed using tools drawn from economics, such as decision theory or decision analysis[32] and information value theory.[33]
- Multi-agent planning problems try to determine the best plan for a community of agents, using cooperation and competition to achieve a given goal.[34] These problems are related to emerging

(b)(3) - P.L. 86-36

fields like evolutionary algorithms and swarm intelligence.

Learning

Important machine learning[35] problems are:

- Unsupervised learning: find a model that matches a stream of input "experiences", and be able to predict what new "experiences" to expect.
- Supervised learning, such as classification (be able to determine what category something belongs in, after being a number of examples of things from each category), or regression (given a set of numerical input/output examples, discover a continuous function that would generate the outputs from the inputs).
- Reinforcement learning:[36] the agent is rewarded for good responses and punished for bad ones. (These can be analyzed in terms decision theory, using concepts like utility).

Natural language processing

Natural language processing[37] gives machines the ability to be read and understand the languages human beings speak. The problem of natural language processing involves such subproblems as: syntax and parsing;[38] semantics and disambiguation;[39] and discourse understanding.[40] Many researchers hope that a sufficiently powerful natural language processing system would be able to acquire knowledge on its own, by reading the existing text available over the internet.

Some straightforward applications of natural language processing include information retrieval (or text mining) and machine translation.[41]

Machine Perception & Computer Vision

Machine perception is the ability to use input from sensors (such as cameras, microphones, sonar and others more exotic) to deduce aspects of the world. Computer vision is the ability to analyze visual input. A few selected subproblems are speech recognition, facial recognition and object recognition.

Motion and manipulation

The field of robotics[42] is closely related to AI. Intelligence is required for robots to be able to handle such tasks as:

- navigate, referred to as robotic mapping including the sub-problems of localization (knowing where you are), mapping (learning what is around you) and path planning (figuring out how to get there).[43]
- manipulate objects (usually described in terms of configuration space).[44]

Social intelligence

Emotion and social skills play two roles for an intelligent agent:[45]

- It must be able to predict the actions of others, by understanding their motives and emotional states. (This involves elements of game theory, decision theory, as well as the ability model human emotions and the perceptual skills to detect emotions.)

(b)(3) - P.L. 86-36

- For good human-computer interaction, an intelligent machine also needs to *display* emotions — at the very least it must appear polite and sensitive to the humans it interacts with. At best, it should appear to have normal emotions itself.

General intelligence

Most researchers hope that their work will eventually be incorporated into a machine with *general* intelligence (known as strong AI), combining all the skills above and exceeding human abilities at most or all of them.[46] A few believe that anthropomorphic features like artificial consciousness or an artificial brain may be required for such a project.

Many of the problems above are considered AI-complete: to solve one problem, you must solve them all. For example, even a straightforward, specific task like machine translation requires that the machine follow the author's argument (reason), know what it's talking about (knowledge), and faithfully reproduce the author's intention (social intelligence). Machine translation, therefore, is believed to be AI-complete: it may require strong AI to be done as well as humans can do it.[47]

Theoretical Approaches

The field of artificial intelligence contains a number of subfields and is integral to many other disciplines such as computer science, engineering, human factors, statistics, biological sciences, physics, and other sciences. This section describes some of these subfields and their approaches. Often these approaches are applied conjunction with each other. Agent-based or cognitive architectures allow researchers to build more versatile and intelligent systems out of intelligent agents that interact in a multi-agent system.

Cognitive simulation

Economist Herbert Simon and Alan Newell studied human problem solving skills and attempted to formalize them, and their work laid the foundations of the field of artificial intelligence, as well as cognitive science, operations research and management science. Their research team performed psychological experiments to demonstrate the similarities between human problem solving and the programs (such as their "General Problem Solver") they were developing. This tradition, centered at Carnegie Mellon University, would eventually culminate in the development of the Soar architecture in the middle 80s.

Logical AI

Unlike Newell and Simon, John McCarthy felt that machines did not need to simulate human thought, but should instead try find the essence of abstract reasoning and problem solving, regardless of whether people used the same algorithms.[48] His laboratory at Stanford (SAIL) focussed on using formal logic to solve wide variety of problems, including knowledge representation, planning and learning. Work in logic led to the development of the programming language Prolog and the science of logic programming.

Symbolic AI

Researchers at MIT (such as Marvin Minsky and Seymour Papert) found that solving difficult problems in vision and natural language processing required ad-hoc solutions -- they argued that there was no silver bullet, no simple and general principle (like logic) that would capture all the aspects of intelligent behavior. An important realization was that AI required large amounts of commonsense knowledge, and

(b)(3) - P.L. 86-36

that this had to be engineered one complicated concept at time.

Knowledge based AI

When computers with large memories became available around 1970, researchers from all three traditions began to build knowledge into AI applications. This "knowledge revolution" led to the development and deployment of expert systems, the first truly successful form of AI software.

Evolutionary Computing

In evolutionary computing, a population is simulated, and algorithms iteratively simulate processes such as growth or development over time to gradually change the population. The way that the population develops is based on a guided random search to achieve a set of desired goals. Such processes are often inspired by biological mechanisms of evolution.

Intelligent agent paradigm

An intelligent agent is a system that perceives its environment and takes actions which maximizes its chances of success. The simplest intelligent agents are programs that solve specific problems. The most complicated intelligent agents would be rational, thinking human beings.[49]

The "intelligent agent" paradigm became widely accepted during the 1990s. Although earlier researchers had proposed modular "divide and conquer" approaches to AI, intelligent agents did not reach their modern form until Judea Pearl, Alan Newell, and others brought concepts from decision theory and economics into the study of AI.

AI Research Tools

In the course of 50 years of research, AI has developed a large number of tools to solve the most difficult problems in Computer Science. A few of the most general of these methods are discussed below.

Search

Many problems in AI can be solved in theory by intelligently searching through many possible solutions: [50]

- Reasoning can be reduced to performing a search. For example, in game playing, the agent can search through a tree of possible moves and counter moves to find a strategy that improves its position. (Tools for two person games include minimax and alpha-beta pruning.)[51] Logical proof can be viewed as searching for a path that leads from premises to conclusions, where each step is the application of an inference rule.[52] Many other reasoning problems, such as constraint satisfaction[53] and dynamic programming[54] are solved using a form of search.
- Planning algorithms search through trees of goals and subgoals, attempting to find a path to a target goal.[55] These sets of goals and subgoals can be represented with graphs (as in the graphplan algorithm),[56] or in a hierarchical task network.[57]
- Robotics algorithms for moving limbs and grasping objects use local searches in configuration space.[44]

(b)(3) - P.L. 86-36

There are several types of search algorithms:

- "Naive" search algorithms, such as breadth first search, depth first search and general state space search.[58]
- Heuristic or "informed" search. The naive algorithms quickly run into problems: the search space quickly can expand to astronomical size in a "combinatorial explosion. "AI researchers realized from the beginning that they would need to trim the search space using heuristics or "rules of thumb".[59] The use of heuristics led to the development of intelligent searches such as greedy best first and A*.[60]
- Local searches, such as hill climbing, simulated annealing and beam search, use techniques borrowed from optimization theory.[61]
- Genetic algorithms are a form of optimization search that imitates the process of natural selection, searching for an artificial phenotype (i.e. any sort of pattern) which passes a fitness measure by producing many copies of the most successful versions (imitating inheritance) and modifying them slightly (imitating mutation).[62]

Logic

Logic[63] was introduced into AI research by John McCarthy in his 1958 Advice Taker proposal. The most important technical development was J. Alan Robinson's discovery of the resolution and unification algorithm for logical deduction in 1963. This procedure is simple, complete and entirely algorithmic, and can easily be performed by digital computers.[64] However, a naive implementation of the algorithm quickly leads to a combinatorial explosion or an infinite loop. In 1974, Robert Kowalski suggested representing logical expressions as Horn clauses (statements in the form of rules: "if *p* then *q*"), which reduced logical deduction to backward chaining or forward chaining. This greatly alleviated (but did not eliminate) the problem.[52][65]

Logic is used for knowledge representation and problem solving, but it can be applied to other problems as well. For example, the satplan algorithm uses logic for planning,[66] and inductive logic programming is a method for learning.[67]

There are several different forms of logic used in AI research.

- Propositional logic[68] or sentential logic is the logic of statements which can be true or false.
- First order logic[69] also allows the use of quantifiers and predicates, and can express facts about objects, their properties, and their relations with each other.
- Fuzzy logic, a version of first order logic which allows the truth of statement to represented as a value between 0 and 1, rather than simply True (1) or False (0). Fuzzy systems can be used for uncertain reasoning and have been widely used in modern industrial and consumer product control systems.[70]
- Default logics, non-monotonic logics and circumscription are forms of logic designed to help with default reasoning and the qualification problem.[26]
- Several extensions of logic have been designed to handle specific domains of knowledge, such as: description logics;[20] situation calculus, event calculus and fluent calculus (for representing events and time);[21] causal calculus;[22] belief calculus; and modal logics.[23]

(b)(3) - P.L. 86-36

Stochastic methods

Starting in the late 80s and early 90s, Judea Pearl and others championed the use of stochastic or probabilistic methods in artificial intelligence.[71] Researchers have used principles from probability theory[72] to devise a number of powerful tools.

Bayesian networks[73] have been applied to a large number of problems, such as: uncertain reasoning (using the Bayesian inference algorithm),[74] learning (using the expectation-maximization algorithm),[75] and planning (using decision networks).[76]

Probabilistic methods have been particularly successful at dealing with processes that occur over time. Several successful algorithms have been developed for filtering, prediction, smoothing and finding explanations for streams of data,[77] such as hidden Markov models,[78] Kalman filters[79] and dynamic Bayesian networks.[80] These tools are used for the problems of perception (such as pattern matching) and learning.

Economic models

AI has been able to use tools drawn from economics, such as decision theory and decision analysis,[32] Bayesian decision networks,[76] information value theory,[33] Markov decision processes,[81] dynamic decision networks,[81] game theory and mechanism design[82] These tools have been especially important for planning problems.

Classifiers and statistical learning methods

The simplest AI applications can be divided into two types: classifiers ("if shiny then diamond") and controllers ("if shiny then pick up"). Controllers do however also classify conditions before inferring actions, and therefore classification forms a central part of many AI systems.

Classifiers[83] are functions that use pattern matching to determine a closest match. They can be tuned according to examples, making them very attractive for use in AI. These examples are known as observations or patterns. In supervised learning, each pattern belongs to a certain predefined class. A class can be seen as a decision that has to be made. All the observations combined with their class labels are known as a data set.

When a new observation is received, that observation is classified based on previous experience. A classifier can be trained in various ways; there are mainly statistical and machine learning approaches.

A wide range of classifiers are available, each with its strengths and weaknesses. Classifier performance depends greatly on the characteristics of the data to be classified. There is no single classifier that works best on all given problems; this is also referred to as the "no free lunch" theorem. Various empirical tests have been performed to compare classifier performance and to find the characteristics of data that determine classifier performance. Determining a suitable classifier for a given problem is however still more an art than science.

The most widely used classifiers are the neural network,[84] kernel methods such as the support vector machine,[85] k-nearest neighbor algorithm,[86] Gaussian mixture model,[87] naive Bayes classifier,[88] and decision tree.[89] The performance of these classifiers have been compared over a wide range of classification tasks[90] in order to find data characteristics that determine classifier performance.

(b)(3) - P.L. 86-36

Neural networks

The study of neural networks[84] began with cybernetics researchers, working in the decade before the field AI research was founded. In the 1960s Frank Rosenblatt developed an important early version, the perceptron.[91] Paul Werbos discovered the backpropagation algorithm in 1984,[92] which led to a renaissance in neural network research and connectionism in general in the middle 1980s. The Hopfield net, a form of attractor network, was first described by John Hopfield in 1982.

Neural networks are applied to the problem of learning, using such techniques as Hebbian learning[93] and the relatively new field of Hierarchical Temporal Memory which simulates the architecture of the neocortex.[94]

Social and emergent models

Several algorithms for learning use tools from evolutionary computation, such as genetic algorithms[95] and swarm intelligence.[96]

Specialized languages

AI researchers have developed several specialized languages for AI research:

- IPL, one of the first programming languages, developed by Alan Newell, Herbert Simon and J. C. Shaw.[97]
- Lisp[98] was developed by John McCarthy at MIT in 1958.[99] There are many dialects of Lisp in use today.
- Prolog,[100] a language based on logic programming, was invented by French researchers Alain Colmerauer and Phillipe Roussel, in collaboration with Robert Kowalski of the University of Edinburgh.[65]
- STRIPS, a planning language developed at Stanford in the 1960s.
- Planner developed at MIT around the same time.

AI applications are also often written in standard languages like C++ and languages designed for mathematics, such as Matlab and Lush.

AI Competitions and Prizes

The Loebner prize is an annual competition to determine the best Turing test competitors. The winner is the computer system that, in the judges' opinions, demonstrates the "most human" conversational behaviour (with learning AI Jabberwacky winning in 2005 and 2006, and A.L.I.C.E. before that), they have an additional prize for a system that in their opinion passes a Turing test. This second prize has not yet been awarded.

The DARPA Grand Challenge is an annual race for a $2 million prize where driverless cars must travel over a hundred miles without any communication with humans, using GPS, computers and a sophisticated array of sensors. The challenge is aimed at a congressional mandate stating that by 2015 one-third of the operational ground combat vehicles of the US Armed Forces should be unmanned.[101] In November 2007, DARPA introduced the DARPA Urban Challenge, a sixty-mile urban area race.

(b)(3) - P.L. 86-36

A popular challenge amongst AI research groups is the RoboCup and FIRA annual international robot soccer competitions. Hiroaki Kitano has formulated the International RoboCup Federation challenge: "In 2050 a team of fully autonomous humanoid robot soccer players shall win the soccer game, comply with the official rule of the FIFA, against the winner of the most recent World Cup."[102]

A lesser known challenge to promote AI research is the annual Arimaa challenge match. The challenge offers a $10,000 prize until the year 2020 to develop a program that plays the board game Arimaa and defeats a group of selected human opponents.

Applications of artificial intelligence

Robots have become common in many industries. They are often given jobs that are considered dangerous to humans. Robots have proven effective in jobs that are very repetitive which may lead to mistakes or accidents due to a lapse in concentration and other jobs which humans may find degrading. General Motors uses around 16,000 robots for tasks such as painting, welding, and assembly. Japan is the leader in using and producing robots in the world. In 1995, 700,000 robots were in use worldwide; over 500,000 of which were from Japan.[103]

Banks use artificial intelligence systems to organize operations, invest in stocks, and manage properties. In August 2001, robots beat humans in a simulated financial trading competition (BBC News, 2001).

Medical clinics use artificial intelligence systems to organize bed schedules, rotate staff, and provide medical information. Many practical applications depend on artificial neural networks. Financial institutions have long used such systems to detect charges or claims outside of the norm, flagging these for human investigation. Neural networks are also being widely deployed in Homeland Security, speech and text recognition, medical diagnosis (such as in Concept Processing technology in EMR software), Data Mining, and E-mail spam filtering.

The 1990s saw some of the first attempts to mass-produce domestically aimed types of basic Artificial Intelligence for education, or leisure. This prospered greatly with the Digital Revolution, and helped introduce people, especially children, to a life of dealing with various types of AI, specifically in the form of Tamagotchis and Giga Pets, the Internet (example: basic search engine interfaces are one simple form), and the first widely released robot, Furby. A mere year later an improved type of domestic robot was released in the form of Aibo, a robotic dog with intelligent features and autonomy.

Some other applications of AI methods:

- Pattern recognition
  - Optical character recognition
  - Handwriting recognition
  - Speech recognition
  - Face recognition
- Artificial Creativity
- Computer vision, Virtual Worlds and Image processing
- Diagnosis (artificial intelligence)
- Game theory and Strategic planning
- Game artificial intelligence and Computer game bot
- Natural language processing, Translation and Chatterbots
- Non-linear control and Robotics

(b)(3) - P.L. 86-36

External links

- AI at the Open Directory Project
- AI-Tools, the Open Source AI community homepage
- Intelligent Systems Program at the University of Pittsburgh
- Artificial Intelligence Directory, a directory of Web resources related to artificial intelligence
- The Association for the Advancement of Artificial Intelligence
- John McCarthy's frequently asked questions about AI
- Artificial Intelligence in the Computer science directory
- Mindmakers.org, an online organization for people building large scale A.I. systems
- Ray Kurzweil's website dedicated to AI including prediction of future development in AI
- Artificial intelligence library and other useful links
- International Journal of Computational Intelligence
- International Journal of Intelligent Technology
- Virtual Humans Forum and Directory

References

Major AI textbooks

-

Luger, George & Stubblefield, William (2004), Artificial Intelligence: Structures and Strategies for Complex Problem Solving (5th ed.), The Benjamin/Cummings Publishing Company, Inc., pp. 720, ISBN 0-8053-4780-1, <http://www.cs.unm.edu/~luger/ai-final/tocfull.html>

-

Nilsson, Nils (1998), Artificial Intelligence: A New Synthesis, Morgan Kaufmann Publishers, ISBN 978-1-55860-467-4

- Template:Russell Norvig 2003
-

Poole, David; Mackworth, Alan & Goebel, Randy (1998), Computational Intelligence: A Logical Approach, Oxford University Press, <http://www.cs.ubc.ca/spider/poole/ci.html>

Other sources

-

ACM, (Association of Computing Machinery) (1998), ACM Computing Classification System: Artificial intelligence, <http://www.acm.org/class/1998/I.2.html>

-

Brooks, Rodney (1990), "Elephants Don't Play Chess", Robotics and Autonomous Systems 6: 3-15,

(b)(3) - P.L. 86-36

<http://people.csail.mit.edu/brooks/papers/elephants.pdf>

- *Buchanan, Bruce G. (2005), "A (Very) Brief History of Artificial Intelligence", AI Magazine: 53-60, <http://www.aaai.org/AITopics/assets/PDF/AIMag26-04-016.pdf>*

- *Haugeland, John (1985), Artificial Intelligence: The Very Idea, Cambridge, Mass.: MIT Press, ISBN 0-262-08153-9.*

- *Hawkins, Jeff & Blakeslee, Sandra (2004), On Intelligence, New York, NY: Owl Books, ISBN 0-8050-7853-3.*

- *Kahneman, Daniel; Slovic, D. & Tversky, Amos (1982), Judgment under uncertainty: Heuristics and biases, New York: Cambridge University Press.*

- *Kurzweil, Ray (1999), The Age of Spiritual Machines, Penguin Books, ISBN 0-670-88217-8*

- *Kurzweil, Ray (2005), The Singularity is Near, Penguin Books, ISBN 0-670-03384-7*

- *Lakoff, George & Núñez, Rafael E. (2000), Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being, Basic Books, ISBN 0-465-03771-2.*

- *Lenat, Douglas (1989), Building Large Knowledge-Based Systems, Addison-Wesley*

- *Lighthill, Professor Sir James (1973), "Artificial Intelligence: A General Survey", Artificial Intelligence: a paper symposium, Science Research Council*

- *McCarthy, John; Minsky, Marvin & Rochester, Nathan et al. (1955), A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>.*

(b)(3) - P.L. 86-36

- *McCarthy, John & Hayes, P. J. (1969), "Some philosophical problems from the standpoint of artificial intelligence", Machine Intelligence **4**: 463-502, <http://www-formal.stanford.edu/jmc/mcchay69.html>*

- *McCorduck, Pamela (2004), Machines Who Think (2nd ed.), Natick, MA: A. K. Peters, Ltd., ISBN 1-56881-205-1.*

- *Minsky, Marvin (1967), Computation: Finite and Infinite Machines, Englewood Cliffs, N.J.: Prentice-Hall*

- *Minsky, Marvin (2006), The Emotion Machine, New York, NY: Simon & Schusterl, ISBN 0-7432-7663-9*

- *Moravec, Hans (1976), The Role of Raw Power in Intelligence, <http://www.frc.ri.cmu.edu/users/hpm/project.archive/general.articles/1975/Raw.Power.html>*

- *Moravec, Hans (1988), Mind Children, Harvard University Press*

- *NRC (1999), "Developments in Artificial Intelligence", Funding a Revolution: Government Support for Computing Research, National Academy Press*

- *Newell, Allen & Simon, H. A. (1963), "GPS: A Program that Simulates Human Thought", in Feigenbaum, E.A. & Feldman, J., Computers and Thought, McGraw-Hill*

- *Searle, John (1980), "Minds, Brains and Programs", Behavioral and Brain Sciences **3** (3): 417-457, <http://www.bbsonline.org/documents/a/00/00/04/84/bbs00000484-00/bbs.searle2.html>*

- *Shapiro, Stuart C. (1992), "Artificial Intelligence", in Shapiro, Stuart C., Encyclopedia of Artificial Intelligence (2nd ed.), New York: John Wiley, pp. 54-57, <http://www.cse.buffalo.edu/~shapiro/Papers/ai.ps>.*

(b)(3) - P.L. 86-36

- *Simon, H. A. (1965), The Shape of Automation for Men and Management, New York: Harper & Row*

- *Sun, R. & Boookman, L., eds. (1994), Computational Architectures: Integrating Neural and Symbolic Processes, Needham, MA: Kluwer Academic Publishers, <http://search.barnesandnoble.com/booksearch/isbninquiry.asp?r=1&ean=9780792395171>*

- *Turing, Alan (October 1950), "Computing machinery and intelligence", Mind* ***LIX*** *(236): 433-460, ISSN 0026-4423, doi:10.1093/mind/LIX.236.433, <http://loebner.net/Prizef/TuringArticle.html>*

- *Wason, P. C. (1966), "Reasoning", in Foss, B. M., New horizons in psychology, Harmondsworth: Penguin*

- *Weizenbaum, Joseph (1976), Computer Power and Human Reason, San Francisco: W.H. Freeman & Company, ISBN 0716704641*

Notes

1. ↑ Russell & Norvig 2003, pp. 5-16
2. ↑ See http://www.aaai.org/AITopics/html/applications.html AI Topics: applications]
3. ↑ Crevier 1993, pp. 64-65
4. ↑ Simon 1965, p. 96 quoted in Crevier 1993, p. 109
5. ↑ Minsky 1967, p. 2 quoted in Crevier 1993, p. 109
6. ↑ Crevier 1993, pp. 161-162,197-203 and and Russell & Norvig 2003, p. 24
7. ↑ Russell Norvig, pp. 25-26
8. ↑ McCarthy et al. 1955 See also Crevier 1993, p. 28
9. ↑ Newell & Simon 1963 and Russell & Norvig 2003, p. 18
10. ↑ Kurzweil 2005, p. 262. Also see Russell Norvig, p. 957 and Crevier 1993, pp. 271 and 279. The most extreme form of this argument (the brain replacement scenario) was put forward by Clark Glymour in the mid-70s and was touched on by Zenon Pylyshyn and John Searle in 1980. It is now associated with Hans Moravec and Ray Kurzweil.
11. ↑ This is a paraphrase of the most important implication of Gödel's theorems, according Template:Hartvtxt. See also Russell & Norvig 2003, p. 949, Gödel 1931, Church 1936, Kleene 1935, Turing 1937, Turing 1950 under "(2) The Mathematical Objection"
12. ↑ Dreyfus 1992, p. 156. See also Dreyfus & Dreyfus 1986, Russell & Norvig 2003, pp. 950-952, Crevier & 1993 120-132 and Hearn 2007, pp. 50-51
13. ↑ Searle 1980. See also Template:Harvtxt: "The assertion that machines could possibly act intelligently (or, perhaps better, act as if they were intelligent) is called the 'weak AI' hypothesis by philosophers, and the assertion that machines that do so are actually thinking (as opposed to simulating thinking) is called the 'strong AI' hypothesis," although Searle's arguments, such as the Chinese Room, apply only to physical symbol systems, not to machines in general (he would

(b)(3) - P.L. 86-36

consider the brain a machine). Also, notice that the positions as Searle states them don't make any commitment to how *much* intelligence the system has: it is one thing to say a machine can act intelligently, it is another to say it can act as intelligently as a human being.

14. ↑ Problem solving, puzzle solving, game playing and deduction: Russell & Norvig 2003, chpt. 3-9, Poole et al. chpt. 2,3,7,9, Luger & Stubblefield 2004, chpt. 3,4,6,8, Nilsson, chpt. 7-12.
15. ↑ Uncertain reasoning: Russell & Norvig 2003, pp. 452-644, Poole, Mackworth & Goebel 1998, pp. 345-395, Luger & Stubblefield 2004, pp. 333-381, Nilsson 1998, chpt. 19
16. ↑ Intractability and efficiency and the combinatorial explosion: Russell & Norvig 2003, pp. 9, 21-22
17. ↑ Several famous examples: Template:Harvtxt showed that people do poorly on completely abstract problems, but if the problem is restated to allowed the use of intuitive social intelligence, performance dramatically improves. (See Wason selection task) Template:Harvtxt have shown that people are terrible at elementary problems that involve uncertain reasoning. (See list of cognitive biases for several examples). Template:Harvtxt have controversially argued that even our skills at mathematics depend on knowledge and skills that come from "the body", i.e. sensorimotor and perceptual skills. (See Where Mathematics Comes From)
18. ↑ Knowledge representation: ACM 1998, I.2.4, Russell & Norvig 2003, pp. 320-363, Poole, Mackworth & Goebel 1998, pp. 23-46, 69-81, 169-196, 235-277, 281-298, 319-345 Luger & Stubblefield 2004, pp. 227-243, Nilsson 1998, chpt. 18
19. ↑ Knowledge engineering: Russell & Norvig 2003, pp. 260-266, Poole, Mackworth & Goebel 1998, pp. 199-233, Nilsson 1998, chpt. ~17.1-17.4
20. ↑ [20.0] [20.1] Representing categories and relations: Semantic networks, description logics, inheritance, including frames and scripts): Russell & Norvig 2003, pp. 349-354, Poole, Mackworth & Goebel 1998, pp. 174-177, Luger & Stubblefield 2004, pp. 248-258, Nilsson 1998, chpt. 18.3
21. ↑ [21.0] [21.1] Representing events and time: Situation calculus, event calculus, fluent calculus (including solving the frame problem): Russell & Norvig 2003, pp. 328-341, Poole, Mackworth & Goebel 1998, pp. 281-298, Nilsson 1998, chpt. 18.2
22. ↑ [22.0] [22.1] Causal calculus: Poole, Mackworth & Goebel 1998, pp. 335-337
23. ↑ [23.0] [23.1] Representing knowledge about knowledge: Belief calculus, modal logics: Russell & Norvig 2003, pp. 341-344, Poole, Mackworth & Goebel 1998, pp. 275-277
24. ↑ Ontology: Russell & Norvig 2003, pp. 320-328
25. ↑ McCarthy & Hayes 1969
26. ↑ [26.0] [26.1] Default reasoning and default logic, non-monotonic logics, circumscription, closed world assumption, abduction (Poole *et al.* places abduction under "default reasoning". Luger *et al.* places this under "uncertain reasoning"): Russell & Norvig 2003, pp. 354-360, Poole, Mackworth & Goebel 1998, pp. 248-256, 323-335 Luger & Stubblefield 2004, pp. 335-363, Nilsson 1998, ~18.3.3
27. ↑ Crevier 1993, pp. 113-114, Moravec 1988, p. 13, Lenat 1989 (Introduction), Russell & Norvig 2003, p. 21
28. ↑ Planning: ACM 1998, ~I.2.8, Russell & Norvig 2003, pp. 375-459, Poole, Mackworth & Goebel 1998, pp. 281-316, Luger & Stubblefield 2004, pp. 314-329, Nilsson 1998, chpt. 10.1-2, 22
29. ↑ Classical planning: Russell & Norvig 2003, pp. 375-430 Poole, Mackworth & Goebel 1998, pp. 281-309, Luger & Stubblefield 2004, pp. 314-329, Nilsson 1998, chpt. 10.1-2, 22
30. ↑ Partial order planning: Russell & Norvig 2003, pp. 387-395, Poole, Mackworth & Goebel 1998, pp. 309-315, Nilsson 1998, chpt. 22.2
31. ↑ Planning and acting in non-deterministic domains: conditional planning, execution monitoring, replanning and continuous planning: Russell & Norvig 2003, pp. 430-449
32. ↑ [32.0] [32.1] Decision theory and decision analysis: Russell & Norvig 2003, pp. 584-597, Poole, Mackworth & Goebel 1998, pp. 381-394
33. ↑ [33.0] [33.1] Information value theory: Russell & Norvig 2003, pp. 600-604

(b)(3) - P.L. 86-36

34. ↑ Multi-agent planning and emergent behavior Russell & Norvig 2003, pp. 449-455
35. ↑ Learning: ACM 1998, I.2.6, Russell & Norvig 2003, pp. 649-788, Poole, Mackworth & Goebel 1998, pp. 397-438, Luger & Stubblefield 2004, pp. 385-542 Nilsson 1998, chpt. 3.3 , 10.3, 17.5, 20
36. ↑ Reinforcement learning: Russell & Norvig 2003, pp. 763-788, Luger & Stubblefield 2004, pp. 442-449
37. ↑ Natural language processing: ACM 1998, I.2.7, Russell & Norvig 2003, pp. 790-831, Poole, Mackworth & Goebel 1998, pp. 91-104, Luger & Stubblefield 2004, pp. 591-632
38. ↑ Syntax and parsing: Russell & Norvig 2003, pp. 795-810, Luger & Stubblefield 2004, pp. 597-616
39. ↑ Semantics and disambiguation: Russell & Norvig 2003, pp. 810-821
40. ↑ Discourse understanding (coherence relations, speech acts, pragmatics): Russell & Norvig 2003, pp. 820-824
41. ↑ Applications of natural language processing, including information retrieval (or text mining) and machine translation Russell & Norvig 2003, pp. 840-857, Luger & Stubblefield 2004, pp. 623-630
42. ↑ Robotics: ACM 1998, I.2.9, Russell & Norvig 2003, pp. 901-942, Poole, Mackworth & Goebel 1998, pp. 443-460
43. ↑ Robotic mapping (localization, etc) Russell Norvig, pp. 908-915
44. ↑ 44.0 44.1 Moving and configuration space: Russell Norivg, pp. 916-932
45. ↑ Minsky 2007, Picard 1997
46. ↑ **Cite error: Invalid <ref> tag; no text was provided for refs named GI**
47. ↑ Shapiro 1992, p. 9
48. ↑ See Science at Google Books, and McCarthy's presentation at AI@50
49. ↑ **Cite error: Invalid <ref> tag; no text was provided for refs named IA**
50. ↑ Search algorithms: Russell & Norvig 2003, pp. 59-189, Poole, Mackworth & Goebel 1998, pp. 113-163, Luger & Stubblefield 2004, pp. 79-164, 193-219, Nilsson 1998, chpt. 7-12
51. ↑ Adversarial search: Russell & Norvig 2003, pp. 161-185, Luger & Stubblefield 2004, pp. 150-157, Nilsson 1998, chpt. 12
52. ↑ 52.0 52.1 Forward chaining, backward chaining, Horn clauses, and logical deduction as search: Russell & Norvig 2003, pp. 217-225, 280-294, Poole, Mackworth & Goebel 1998, pp. ~46-52, Luger & Stubblefield 2004, pp. 62-73, Nilsson 1998, chpt. 4.2, 7.2
53. ↑ Constraint satisfaction: Russell & Norvig 2003, pp. 137-156, Poole, Mackworth & Goebel 1998, pp. pp. 147-163
54. ↑ Dynamic programming: Russell & Norvig 2003, p. 293, Poole, Mackworth & Goebel 1998, pp. 145-147, Nilsson 1998, p. 178
55. ↑ State space search and planning: Russell & Norvig 2003, pp. 382-387, Poole, Mackworth & Goebel 1998, pp. 298-305, Nilsson 1998, chpt. 10.1-2
56. ↑ Graphplan: Russell & Norvig 2003, pp. 395-402
57. ↑ Hierarchical task network: Russell & Norvig 2003, pp. 422-430
58. ↑ Naive searches: Russell & Norvig 2003, pp. 59-93, Poole, Mackworth & Goebel 1998, pp. 113-132, Luger & Stubblefield 2004, pp. 79-121, Nilsson 1998, chpt. 8
59. ↑ John McCarthy writes that "the combinatorial explosion problem has been recognized in AI from the beginning" in Review of Lighthill report
60. ↑ Heuristic or informed searches: Russell & Norvig 2003, pp. 94-109, Poole, Mackworth & Goebel 1998, pp. pp. 132-147, Luger & Stubblefield 2004, pp. 133-150, Nilsson 1998, chpt. 9
61. ↑ Optimization searches: Russell & Norvig 2003, pp. 110-116,120-129, Poole, Mackworth & Goebel 1998, pp. 56-163, Luger & Stubblefield 2004, pp. 127-133
62. ↑ Genetic algorithms: Russell & Norvig 2003, pp. 116-119, Poole, Mackworth & Goebel 1998, pp. 162, Luger & Stubblefield 2004, pp. 509-530, Nilsson 1998, chpt. 4.2
63. ↑ Logic: ACM 1998, ~I.2.3, Russell & Norvig 2003, pp. 194-310, Luger & Stubblefield 2004, pp. 35-77, Nilsson 1998, chpt. 13-16
64. ↑ Resolution and unification: Russell & Norvig 2003, pp. 213-217, 275-280, 295-306, Poole,

(b)(3) - P.L. 86-36

Mackworth & Goebel 1998, pp. 56-58, Luger & Stubblefield 2004, pp. 554-575, Nilsson 1998, chpt. 14 & 16

65. ↑ [65.0] [65.1] History of logic programming: Crevier 1993, pp. 190-196. Advice Taker: McCorduck 2004, p. 51, Russell & Norvig 2003, pp. 19
66. ↑ Satplan: Russell & Norvig 2003, pp. 402-407, Poole, Mackworth & Goebel 1998, pp. 300-301, Nilsson 1998, chpt. 21
67. ↑ Explanation based learning, relevance based learning, inductive logic programming, case based reasoning: Russell & Norvig 2003, pp. 678-710, Poole, Mackworth & Goebel 1998, pp. 414-416, Luger & Stubblefield 2004, pp. ~422-442, Nilsson 1998, chpt. 10.3, 17.5
68. ↑ Propositional logic: Russell & Norvig 2003, pp. 204-233, Luger & Stubblefield 2004, pp. 45-50 Nilsson 1998, chpt. 13
69. ↑ First order logic and features such as equality: ACM 1998, ~I.2.4, Russell & Norvig 2003, pp. 240-310, Poole, Mackworth & Goebel 1998, pp. 268-275, Luger & Stubblefield 2004, pp. 50-62, Nilsson 1998, chpt. 15
70. ↑ Fuzzy logic: Russell & Norvig 2003, pp. 526-527
71. ↑ Russell & Norvig 2003, pp. 25-26 (on Judea Pearl's contribution). Stochastic methods are described in all the major AI textbooks: ACM 1998, ~I.2.3, Russell & Norvig 2003, pp. 462-644, Poole, Mackworth & Goebel 1998, pp. 345-395, Luger & Stubblefield 2004, pp. 165-191, 333-381, Nilsson 1998, chpt. 19
72. ↑ Probability: Russell & Norvig 2003, pp. 462-489, Poole, Mackworth & Goebel 1998, pp. 346-366, Luger & Stubblefield 2004, pp. ~165-182, Nilsson 1998, chpt. 19.1
73. ↑ Bayesian networks: Russell & Norvig 2003, pp. 492-523, Poole, Mackworth & Goebel 1998, pp. 361-381, Luger & Stubblefield 2004, pp. ~182-190, ~363-379, Nilsson 1998, chpt. 19.3-4
74. ↑ Bayesian inference algorithm: Russell & Norvig 2003, pp. 504-519, Poole, Mackworth & Goebel 1998, pp. 361-381, Luger & Stubblefield 2004, pp. ~363-379, Nilsson 1998, chpt. 19.4 & 7
75. ↑ Bayesian learning and the expectation-maximization algorithm Russell & Norvig 2003, pp. 712-724, Poole, Mackworth & Goebel 1998, pp. 424-433, Nilsson 1998, chpt. 20
76. ↑ [76.0] [76.1] Bayesian decision networks: Russell & Norvig 2003, pp. 597-600
77. ↑ Russell & Norvig 2003, pp. 537-581
78. ↑ Hidden Markov model: Russell & Norvig 2003, pp. 549-551
79. ↑ Kalman filter: Russell & Norvig 2003, pp. 551-557
80. ↑ Dynamic Bayesian network: Russell & Norvig 2003, pp. 551-557
81. ↑ [81.0] [81.1] Markov decision processes and dynamic decision networks: Russell & Norvig 2003, pp. 613-631
82. ↑ Game theory and mechanism design: Russell & Norvig 2003, pp. 631-643
83. ↑ Statistical learning methods and classifiers: Russell & Norvig 2003, pp. 712-754, Luger & Stubblefield 2004, pp. 453-541
84. ↑ [84.0] [84.1] Neural networks and connectionism: Russell & Norvig 2003, pp. 736-748, Poole, Mackworth & Goebel 1998, pp. 408-414, Luger & Stubblefield 2004, pp. 453-505, Nilsson 1998, chpt. 3
85. ↑ Kernel methods: Russell & Norvig 2003, pp. 749-752
86. ↑ K-nearest neighbor algorithm: Russell & Norvig 2003, pp. 733-736
87. ↑ Gaussian mixture model: Russell & Norvig 2003, pp. 725-727
88. ↑ Naive Bayes classifier: Russell & Norvig 2003, pp. 718
89. ↑ Decision tree: Russell & Norvig 2003, pp. 653-664, Poole, Mackworth & Goebel 1998, pp. 403-408, Luger & Stubblefield 2004, pp. 408-417
90. ↑ Template:Cite-web
91. ↑ Perceptrons: Russell & Norvig 2003, pp. 740-743, Luger & Stubblefield 2004, pp. 458-467
92. ↑ Backpropagation: Russell & Norvig 2003, pp. 744-748, Luger & Stubblefield 2004, pp. 467-474, Nilsson 1998, chpt. 3.3
93. ↑ Competitive learning, Hebbian coincidence learning, Hopfield networks and attractor networks:

(b)(3) - P.L. 86-36

Luger & Stubblefield 2004, pp. 474-505.
94. ↑ Hawkins & Blakeslee 2004
95. ↑ Genetic algorithms for learning: Luger & Stubblefield 2004, pp. 509-530, Nilsson 1998, chpt. 4.2
96. ↑ Artificial life and society based learning: Luger & Stubblefield 2004, pp. 530-541
97. ↑ Crevier 1993, p. 46-48
98. ↑ Lisp: Luger & Stubblefield 2004, pp. 723-821
99. ↑ Crevier 1993, pp. 59-62, Russell & Norvig 2003, p. 18
100. ↑ Prolog: Poole, Mackworth & Goebel 1998, pp. 477-491, Luger & Stubblefield 2004, pp. 641-676, 575-581
101. ↑ Congressional Mandate DARPA
102. ↑ The RoboCup2003 Presents: Humanoid Robots playing Soccer PRESS RELEASE: 2 June 2003
103. ↑ "Robot," Microsoft® Encarta® Online Encyclopedia 2006

(b)(3) - P.L. 86-36

Retrieved from [redacted]

Categories: Artificial Intelligence | Cybernetics | Science and Technology

UNCLASSIFIED

(b)(6)