

Artificial intelligence

This article is about modelling human thought with computers,. For other uses of the term AI, see [Ai](#).

Introduction

Although there is no clear definition of AI (not even of intelligence), it can be described as the attempt to build machines that think and act like humans, that are able to learn and to use their knowledge to solve problems on their own.

A 'by-product' of the intensive studies of the human brain by AI researchers is a far better understanding of how it works.

The human brain consists of 10 to 100 billion neurons, each of which is connected to between 10 and 10,000 others through synapses. The single brain cell is comparatively slow (compared to a microprocessor) and has a very simple function: building the sum of its inputs and issuing an output, if that sum exceeds a certain value. Through its highly parallel way of operation, however, the human brain achieves a performance that has not been reached by computers yet; and even at the current speed of development in that field, we still have about twenty years until the first supercomputers will be of equal power.

In the meantime, a number of different approaches are tried to build models of the brain, with different levels of success.

The only test for intelligence there is, is the [Turing Test](#). A thinking machine has yet to be built.

2. Why the current approach of AI is wrong

These following thoughts do not deal with technical problems of AI, nor am I going to prove that humans are the only intelligent species (exactly the opposite, see [section 5](#)). What I want to show, is that the whole idea of AI needs to be changed in order to lead to more than just partial results.

2.1. Reference Points

Today's AI concentrates entirely on the brain. If you look at the human body, however, it is not clear where to draw the line between which parts of the nervous system belong to the brain, and which don't. But a number of functions are performed by the spinal cord, for example, like withdrawing the hand quickly when touching something hot. It can be vitally important that this action is taken as fast as possible, in order to limit the damage. The only way of doing so is through reflexes, without the intervention of the brain. This is not an example of intelligent decisions outside the brain, but it provides an entry point to the following.

Whenever you talk to somebody, you use a huge amount of assumptions about the background of your counterpart. You usually start with assuming that other person is almost identical to yourself, and by small misunderstandings and questions of him (for the sake of easier reading I will assume that our counterpart is male) you correct that picture you have. When you know somebody, you don't have a list of all his features in your mind, but you know the differences between you and that person (or the difference between him and a third party, be it a single person or a group). I cannot prove this, but it explains why it often is so difficult to describe somebody to another person, since that other person's reference point is different from your own.

We build relations between all the things we know, and we build classes of things by putting objects with certain similar features into one such class. But the most basic difference we see is the difference between ourselves and the objects, the not-ourselves. Man's first reference point is himself, which is obvious when looking around oneself: Isn't that cat looking very nosy at you? Doesn't that monitor's face look at you? Don't that car's headlights look into a certain direction? Haven't you taught that stupid computer programme who the master was, yesterday?

People's categories are based on people, first of all. This is not an intellectual decision, but a natural necessity. How could you ever find out anything without a first reference point that you could relate it to? This also is the reason why in children's minds everything 'lives': They live themselves, so why shouldn't other things? Why shouldn't that teddybear feel hunger, exactly like I do?

The point is clear: No knowledge can be accumulated without a reference point. AI doesn't obey this. Most 'intelligent' programmes are equipped with knowledge, but none has ever had a clear picture of itself that it could relate everything else to. This is the first deficiency.

2.2. The Role of the Body

The most obvious difference between man and other animals is his mind, his ability to accumulate knowledge and pass it on to his descendants. Yet many of man's highly developed abilities can be completely switched off by the sheer terror of a single aching tooth. This also applies for other strong feelings like hatred, grief or pain in general. They can make people act against their better knowledge and their principles - these being higher developed parts of the mind. This leads to a conclusion that is obvious from looking at the ancestry of man: the vital functions rule over everything else. Man has not been built to wear digital watches (as Douglas Addams states), but he is the winner of a game that is as old as life itself: Evolution. If people were able to simply ignore hunger, they would starve to death; if they would have to control their lungs consciously, they would sooner or later suffocate. Vital functions must have priority over everything else.

Considering this, it is not surprising that many of our expressions involve basic needs, like 'being hungry for love', 'being tired of something', 'being fed up', 'having a bleeding heart', 'saving someone's skin'. This is also an example for what was said in the last paragraph: The main reference point is oneself, this is of course also true for strong emotions that are on a less 'basic' level.

Additionally, most (if not all) emotions are accompanied by physical symptoms, such as the production of hormones, shivering, gnashing of teeth or goose-pimples.

The point I want to make here is that the human body plays an important role in all intellectual processes, since they are mere subordinates of its needs in order to stay alive. It is therefore short-sighted (another example!) to try to build artificial minds not only without any body, but also without even the *concept* of a body. How should an artificial mind ever be able to understand tiredness, excitement, happiness or fear without ever having *felt* it? And by feeling, I mean the physical symptoms, and the intellectual processes that accompany the fear of injury or death, for example. A body-less mind can never understand that, and thus will never be able to understand humans, let alone act like one. This is the second deficiency.

Two additional points here: AI, being a science, acts very unscientific here. What is being done (although not on purpose) is, that while denying the existence of an immaterial soul that is independent of the body, only the soul is taken into account, but not the effects of the body. The second point is an observation: The reason why men and women often have such difficulties understanding one another, might be their different reference points, and also their different bodies, that affect the way they think. As shown above, these two points can easily lead to misunderstandings, when one assumes the counterpart to be too similar to oneself. This does not, however, mean that one of the sexes is superior, but that they are different, and that these differences have to be worked out and brought into people's minds.

2.3. Evolution

This little interlude also brought another point up: the desire to replicate. This being the most basic of all basic desires, it cannot be ignored. No experiment in the field of AI has yet specified which gender the mind would be. In a world, where there is no need for a partner to replicate (because there is a simple 'copy'-mechanism, for example), social structures would be completely different from man's. It is even the question, if different genders were necessary at all. Thus, an artificial mind would also need the 'ability' to die, as well as that to mate and replicate, otherwise the resulting being would be beyond any recognition by a human mind.

A mind that cannot die and that doesn't feel the need to replicate in a manner similar to humans, would be very different from man. Such an environment would have to be created artificially, but in a different sense: The conditions would have to be made more difficult than they needed to be, only to force the beings to act human-like. And only an evolutionary process would lead to a mind similar to a human one. This is the third deficiency.

3. The Art of Learning

Learning isn't a 'static' ability, but one which continuously changes: You must learn how to learn, and the way you learn changes. The new-born child can only learn by first-hand experience, and hardly generalize. But the older the child gets, the more he/she can learn without having experienced a corresponding situation. Indeed, most of what you know is what you were taught by others, read in books, etc. This probably is the main advantage of man over all other animals: that we can pass on knowledge, so that the next generation doesn't have to make the same mistakes again (it does anyway, but that's a different problem ...). Your knowledge includes the first-hand experiences of hundreds of thousands of people, whose knowledge and experiences were collected and put into a structured form, in order to make learning these facts easier. People nowadays aren't more intelligent than 1000 years ago, but we have more knowledge, and thus can achieve far higher goals. Like Newton said: "I am a dwarf, but I can see very far for I am standing on the shoulders of giants"

But what effects does this have on AI? The brain changes, and not only does its knowledge change, but also the way it accumulates knowledge and makes use of it. An artificial intelligence must be able to change its own programme.

4. Thinking Outside the Brain

The previous point contains another interesting thought: Whatever you do, whatever you think or say - it hasn't been thought up entirely by you, it always contains parts from other people. This is the key to developing beyond what a single generation can reach (see previous section). But it also leads us to a somewhat discomfoting question: How much of that brain is actually mine? What percentage of what I think has been thought by others already? How different am I?

Imagine a group of five people working on the solution of a complicated problem. They're sitting around a table, discussing ideas and problems, taking notes, and thinking about whether or not that last proposal was good. In the end, they will come up with a solution that will be far better than one that any of them had worked out in his/her own! The result will even be better than if each of them was assigned a fifth of the problem, and solved it alone. This team is able to achieve more than five single persons can do independently.

But what is the difference between five isolated persons and a team of five? If the team develops an idea that the single persons don't, which of the members created it? It's not a person that created the idea, but the interaction process, the discussion. An act of thinking has been done by an immaterial process, not a single person.

5. The idea of a General Concept of Intelligence

All the points made above make one thing clear: In order to build an artificial intelligence, it must be built as human-like as possible. Without basic human 'ingredients', the resulting mind might not even be recognized as such. This boils down to the feeling that the goal is to build a mere copy of the human mind.


Why on earth, one might wonder, would anybody want to build a copy of the human mind? Isn't the original working fine? Isn't it superior to everything known? Isn't one's mind the most difficult thing to be examined by itself? What would be the use of such an artificial mind, that would need even more artificial means, only to stay human-like?

The only logical solution to this is to completely separate human from artificial intelligence, in order to build something entirely new.

This naturally leads to the idea of a higher principle of Intelligence, that human intelligence is only one manifestation of (in order to distinguish between the traditional human intelligence and this new idea of a more general concept, I want

to spell the latter with a capital I: Intelligence). Another one would be artificial intelligence, and another one still the intelligence developed on a planet many lightyears from here. Again, I remind you of the points just made. Considering these, how should a mind that is the result of evolution on an entirely different planet be similar to ours in any way? There must be similarities, but on a higher level: on the level of Intelligence (note the capital I).

In that hierarchy, AI is on the same level as human intelligence, together with animal intelligence and any other kind of intelligence that one might encounter. The following figure illustrates this:



Artificial intelligence, also known as **machine intelligence**, is defined as [intelligence](#) exhibited by anything manufactured (i.e. [artificial](#)) by humans or other sentient beings or systems (should such things ever exist on [Earth](#) or [elsewhere](#)). It is usually hypothetically applied to general-purpose [computers](#). The term is also used to refer to the field of scientific investigation into the plausibility of and approaches to creating such systems.

Overview

The question of what artificial intelligence is can be reduced to two parts: "what is the nature of artifice" and "what is intelligence"? The first question is fairly easy to answer, though it does point to the question of what it is possible to manufacture (within the constraints of certain types of system, e.g. classical computational systems, of available processes of manufacturing and of possible limits on human intellect, for instance).

The second is much harder, raising questions of [consciousness](#) and [self](#), [mind](#) (including the [unconscious mind](#)) and the question of what components are involved in the only type of [intelligence](#) it is universally agreed we have available to study: that of human beings. Intelligent behavior in humans is complex and difficult to study or understand. Study of animals and artificial systems that are not just models of what exists already are also considered widely pertinent.

Several distinct types of artificial intelligence have been elucidated below. Also, the subject divisions, history, proponents and opponents and applications of research in the subject are described. Finally, references to fictional and non-fictional descriptions of AI are provided.

Strong AI and weak AI

One popular and early definition of artificial intelligence research, put forth by [John McCarthy](#) at the [Dartmouth Conference](#) in 1956, is "*making a machine behave in ways that would be called intelligent if a human were so behaving.*", repeating the claim put forth by [Alan Turing](#) in "[Computing machinery and intelligence](#)" (Mind, October 1950). However this definition seems to ignore the possibility of strong AI (see below). Another definition of **artificial intelligence** is [intelligence](#) arising from an artificial device. Most definitions could be categorized as concerning either systems that *think like humans*, systems that *act like humans*, systems that *think rationally* or systems that *act rationally*.

Strong artificial intelligence

Strong artificial intelligence research deals with the creation of some form of computer-based artificial intelligence that can truly [reason](#) and [solve problems](#); a strong form of AI is said to be [sentient](#), or self-aware. In theory, there are two types of strong AI:

- Human-like AI, in which the computer program thinks and reasons much like a [human mind](#).
- Non-human-like AI, in which the computer program develops a totally non-human sentience, and a non-human way of thinking and reasoning.

Weak artificial intelligence

Weak artificial intelligence research deals with the creation of some form of computer-based artificial intelligence that can reason and solve problems only in a limited domain; such a machine would, in some ways, act *as if* it were intelligent, but it would not possess true intelligence or sentience. The classical test for such abilities is the [Turing test](#).

There are several fields of weak AI, one of which is [natural language](#). Many weak AI fields have specialised software or programming languages created for them. For example, the 'most-human' natural language [chatbot A.L.I.C.E.](#) uses a programming language [AIML](#) that is specific to its program, and the various clones, named [Alicebots](#).

To date, much of the work in this field has been done with [computer](#) simulations of intelligence based on predefined sets of rules. Very little progress has been made in strong AI. Depending on how one defines one's goals, a moderate amount of progress has been made in weak AI.

Philosophical criticism and support of strong AI

The term "Strong AI" was originally coined by [John Searle](#) and was applied to digital [computers](#) and other [information processing](#) machines. Searle defined strong AI:

"according to strong AI, the computer is not merely a tool in the study of the mind; rather, the appropriately programmed computer really is a mind" (J Searle in Minds Brains and Programs. The Behavioral and Brain Sciences, vol. 3, 1980).

Searle and most others involved in this debate are addressing the problem of whether a machine that works solely through the transformation of encoded data could be a mind, not the wider issue of [Monism](#) versus [Dualism](#) (ie: whether a machine of any type, including biological machines, could contain a mind).

Searle points out in his [Chinese Room](#) Argument that information processors carry encoded data which describe other things. The encoded data itself is meaningless without a cross reference to the things it describes. This leads Searle to point out that there is no meaning or understanding in an information processor itself. As a result Searle claims to demonstrate that even a machine that passed the [Turing test](#) would not necessarily be [conscious](#) in the human sense.

Some philosophers hold that if Weak AI is accepted as possible then Strong AI must also be possible. Daniel C. Dennett argues in Consciousness Explained that if

there is no magic spark or soul, then Man is just a machine, and he asks why the Man-machine should have a privileged position over all other possible machines when it comes to intelligence or 'mind'. [Simon Blackburn](#) in his introduction to philosophy, Think, points out that you might appear intelligent but there is no way of telling if that intelligence is real (ie: a 'mind'). However, if the discussion is limited to strong AI rather than [artificial consciousness](#) it may be possible to identify features of human minds that do not occur in information processing computers.

Strong AI seems to involve the following assumptions about the mind and brain:

1. the mind is software, a [finite state machine](#) so the [Church-Turing thesis](#) applies to it
2. [presentism](#) describes the mind
3. the brain is purely hardware (i.e. only follows the rules of a classical computer)

The first assumption is particularly problematic because of the old adage that any computer is just a glorified abacus. It is indeed possible to construct any type of information processor out of balls and wood, although such a device would be very slow and prone to failure it would be able to do anything that a modern computer can do. This means that the proposition that information processors can be minds is equivalent to proposing that minds can exist as devices made of rolling balls in wooden channels.

Some (including [Roger Penrose](#)) attack the applicability of the Church-Turing thesis directly by drawing attention to the [halting problem](#) in which certain types of computation cannot be performed by information systems yet seem to be performed by human minds.

Ultimately the truth of Strong AI depends upon whether information processing machines can include all the properties of minds such as [Consciousness](#). However, Weak AI is independent of the Strong AI problem and there can be no doubt that many of the features of modern computers such as multiplication or database searching might have been considered 'intelligent' only a century ago.

History

Development of AI theory

Much of the (original) focus of artificial intelligence research draws from an experimental approach to [psychology](#), and emphasizes what may be called linguistic intelligence (best exemplified in the [Turing test](#)).

Approaches to artificial intelligence that do not focus on linguistic intelligence include [robotics](#) and [collective intelligence](#) approaches, which focus on active manipulation of an environment, or [consensus decision making](#), and draw from [biology](#) and [political science](#) when seeking models of how "intelligent" behavior is organized.

Artificial intelligence [theory](#) also draws from animal studies, in particular with insects, which are easier to emulate as robots (see [artificial life](#)), as well as animals with more complex cognition, including [apes](#), who resemble humans in many ways but have less developed capacities for planning and cognition. AI researchers argue that animals, which are simpler than humans, ought to be considerably easier to mimic. But satisfactory computational models for animal intelligence are not available.

Seminal papers advancing the concept of machine intelligence include *A Logical Calculus of the Ideas Immanent in Nervous Activity* (1943), by [Warren McCulloch](#) and [Walter Pitts](#), and *On Computing Machinery and Intelligence* (1950), by [Alan Turing](#), and *Man-Computer Symbiosis* by J.C.R. Licklider. See [cybernetics](#) and [Turing test](#) for further discussion.

There were also early papers which denied the possibility of machine intelligence on [logical](#) or [philosophical](#) grounds such as *Minds, Machines and Gödel* (1961) by [John Lucas](#) [1].

With the development of practical techniques based on AI research, advocates of AI have argued that opponents of AI have repeatedly changed their position on tasks such as [computer chess](#) or [speech recognition](#) that were previously regarded as "intelligent" in order to deny the accomplishments of AI. They point out that this moving of the goalposts effectively defines "intelligence" as "whatever humans can do that machines cannot".

[John von Neumann](#) (quoted by [E.T. Jaynes](#)) anticipated this in 1948 by saying, in response to a comment at a lecture that it was impossible for a machine to think: "You insist that there is something a machine cannot do. If you will tell me *precisely* what it is that a machine cannot do, then I can always make a machine

which will do just that!". Von Neumann was presumably alluding to the [Church-Turing thesis](#) which states that any effective procedure can be simulated by a (generalized) computer.

In [1969](#) McCarthy and Hayes started the discussion about the [frame problem](#) with their essay, "Some Philosophical Problems from the Standpoint of Artificial Intelligence".
Experimental AI research

Artificial intelligence began as an experimental field in the [1950s](#) with such pioneers as [Allen Newell](#) and [Herbert Simon](#), who founded the first artificial intelligence laboratory at [Carnegie-Mellon University](#), and McCarthy and [Marvin Minsky](#), who founded the [MIT AI Lab](#) in [1959](#). They all attended the aforementioned [Dartmouth College](#) summer AI conference in [1956](#), which was organized by McCarthy, Minsky, Nathan Rochester of [IBM](#) and [Claude Shannon](#).

Historically, there are two broad styles of AI research - the "neats" and "scruffies". "Neat", *classical* or [symbolic](#) AI research, in general, involves symbolic manipulation of abstract concepts, and is the methodology used in most expert systems. Parallel to this are the "scruffy", or "connectionist", approaches, of which [neural networks](#) are the best-known example, which try to "evolve" intelligence through building systems and then improving them through some automatic process rather than systematically designing something to complete the task. Both approaches appeared very early in AI history. Throughout the [1960s](#) and [1970s](#) scruffy approaches were pushed to the background, but interest was regained in the [1980s](#) when the limitations of the "neat" approaches of the time became clearer. However, it has become clear that contemporary methods using *both* broad approaches have severe limitations.

Artificial intelligence research was very heavily funded in the [1980s](#) by the [Defense Advanced Research Projects Agency](#) in the [United States](#) and by the [fifth generation computer systems project](#) in [Japan](#). The failure of the work funded at the time to produce immediate results, despite the grandiose promises of some AI practitioners, led to correspondingly large cutbacks in funding by government agencies in the late 1980s, leading to a general downturn in activity in the field known as AI winter. Over the following decade, many AI researchers moved into related areas with more modest goals such as [machine learning](#), [robotics](#), and [computer vision](#), though research in pure AI continued at reduced levels.

Brief History of Artificial Intelligence

The intellectual roots of AI, and the concept of intelligent machines, may be found in Greek mythology. Intelligent artifacts appear in literature since then, with real (and fraudulent) mechanical devices actually demonstrated to behave with some degree of intelligence. Some of these conceptual achievements are listed below under "[Ancient History](#)."

After modern computers became available, following World War II, it has become possible to create programs that perform difficult intellectual tasks. From these programs, general tools are constructed which have applications in a wide variety of everyday problems. Some of these computational milestones are listed below under "[Modern History](#)."

ANCIENT HISTORY

Greek myths of Hephaestus and Pygmalion incorporate the idea of intelligent robots. Many other myths in antiquity involve human-like artifacts. Many mechanical toys and models were actually constructed, e.g., by Hero, Daedalus and other real persons.

5th century B.C.

Aristotle invented syllogistic logic, the first formal deductive reasoning system.

13th century

Talking heads were said to have been created, Roger Bacon and Albert the Great reputedly among the owners.

Ramon Llull, Spanish theologian, invented machines for discovering nonmathematical truths through combinatorics.

15th century

Invention of printing using moveable type. Gutenberg Bible printed (1456).

15th-16th century

Clocks, the first modern measuring machines, were first produced using lathes.

16th century

Clockmakers extended their craft to creating mechanical animals and other novelties.

Rabbi Loew of Prague is said to have invented the Golem, a clay man brought to life (1580).

17th century

Early in the century, Descartes proposed that bodies of animals are nothing more than complex machines. Many other 17th century thinkers offered variations and elaborations of Cartesian mechanism.

Hobbes published *The Leviathan*, containing a material and combinatorial theory of thinking.

Pascal created the first mechanical digital calculating machine (1642).

Leibniz improved Pascal's machine to do multiplication & division (1673) and envisioned a universal calculus of reasoning by which arguments could be decided mechanically.

18th century

The 18th century saw a profusion of mechanical toys, including the celebrated mechanical duck of Vaucanson and von Kempelen's phony mechanical chess player, The Turk (1769).

19th century

Luddites (led by Ned Ludd) destroyed machinery in England (1811-1816).

Mary Shelley published the story of Frankenstein's monster (1818).

George Boole developed a binary algebra representing (some) "laws of thought."

Charles Babbage & Ada Byron (Lady Lovelace) worked on programmable mechanical calculating machines.

20th century - First Half

Bertrand Russell and Alfred North Whitehead published *Principia Mathematica*, which revolutionized formal logic. Russell, Ludwig Wittgenstein, and Rudolf Carnap lead philosophy into logical analysis of knowledge.

Karel Capek's play "*R.U.R.*" (*Rossum's Universal Robots*) opens in London (1923). - First use of the word 'robot' in English.

Warren McCulloch & Walter Pitts publish "A Logical Calculus of the Ideas Immanent in Nervous Activity" (1943), laying foundations for neural networks.

Arturo Rosenblueth, Norbert Wiener & Julian Bigelow coin the term "cybernetics" in a 1943 paper. Wiener's popular book by that name published in 1948.

Vannevar Bush published *As We May Think* (Atlantic Monthly, July 1945) a prescient vision of the future in which computers assist humans in many activities.

A.M. Turing published "Computing Machinery and Intelligence" (1950). - Introduction of Turing's Test as a way of operationalizing a test of intelligent behavior.

Claude Shannon published detailed analysis of chess playing as search (1950).

Isaac Asimov published his three laws of robotics (1950).

MODERN HISTORY

1956	John McCarthy coined the term "artificial intelligence" as the topic of the Dartmouth Conference, the first conference devoted to the subject.
	Demonstration of the first running AI program, the Logic Theorist (LT) written by Allen Newell, J.C. Shaw and Herbert Simon (Carnegie Institute of Technology, now Carnegie Mellon University).
1957	The General Problem Solver (GPS) demonstrated by Newell, Shaw & Simon.
1952-62	Arthur Samuel (IBM) wrote the first game-playing program, for checkers, to achieve sufficient skill to challenge a world champion. Samuel's machine learning programs were responsible for the high performance of the checkers player.
1958	John McCarthy (MIT) invented the Lisp language.
	Herb Gelernter & Nathan Rochester (IBM) described a theorem prover in geometry that exploits a semantic model of the domain in the form of diagrams of "typical" cases.
	Teddington Conference on the Mechanization of Thought Processes was held in the UK and among the papers presented were John McCarthy's "Programs with Common Sense," Oliver Selfridge's "Pandemonium," and Marvin Minsky's "Some Methods of Heuristic Programming and Artificial Intelligence."
Late 50's & Early 60's	Margaret Masterman & colleagues at Cambridge design semantic nets for machine translation.
1961	James Slagle (PhD dissertation, MIT) wrote (in Lisp) the first symbolic integration program, SAINT, which solved calculus problems at the college freshman level.
1962	First industrial robot company, Unimation, founded.
1963	Thomas Evans' program, ANALOGY, written as part of his PhD work at MIT, demonstrated that computers can solve the same analogy problems as are given on IQ tests.
	Ivan Sutherland's MIT dissertation on Sketchpad introduced the idea of interactive graphics into computing.

	Edward A. Feigenbaum & Julian Feldman published <i>Computers and Thought</i> , the first collection of articles about artificial intelligence.
1964	Danny Bobrow's dissertation at MIT (tech.report #1 from MIT's AI group, Project MAC), shows that computers can understand natural language well enough to solve algebra word problems correctly.
	Bert Raphael's MIT dissertation on the SIR program demonstrates the power of a logical representation of knowledge for question-answering systems
1965	J. Alan Robinson invented a mechanical proof procedure, the Resolution Method, which allowed programs to work efficiently with formal logic as a representation language.
	Joseph Weizenbaum (MIT) built ELIZA, an interactive program that carries on a dialogue in English on any topic. It was a popular toy at AI centers on the ARPA-net when a version that "simulated" the dialogue of a psychotherapist was programmed.
1966	Ross Quillian (PhD dissertation, Carnegie Inst. of Technology; now CMU) demonstrated semantic nets.
	First Machine Intelligence workshop at Edinburgh - the first of an influential annual series organized by Donald Michie and others.
	Negative report on machine translation kills much work in Natural Language Processing (NLP) for many years.
1967	Dendral program (Edward Feigenbaum, Joshua Lederberg, Bruce Buchanan, Georgia Sutherland at Stanford) demonstrated to interpret mass spectra on organic chemical compounds. First successful knowledge-based program for scientific reasoning.
	Joel Moses (PhD work at MIT) demonstrated the power of symbolic reasoning for integration problems in the Macsyma program. First successful knowledge-based program in mathematics.
	Richard Greenblatt at MIT built a knowledge-based chess-playing program, MacHack, that was good enough to achieve a class-C rating in tournament play.
Late 60s	Doug Engelbart invented the mouse at SRI.
1968	Marvin Minsky & Seymour Papert publish Perceptrons, demonstrating

	limits of simple neural nets.
1969	SRI robot, Shakey, demonstrated combining locomotion, perception and problem solving.
	Roger Schank (Stanford) defined conceptual dependency model for natural language understanding. Later developed (in PhD dissertations at Yale) for use in story understanding by Robert Wilensky and Wendy Lehnert, and for use in understanding memory by Janet Kolodner.
	First International Joint Conference on Artificial Intelligence (IJCAI) held in Washington, D.C.
1970	Jaime Carbonell (Sr.) developed SCHOLAR, an interactive program for computer-aided instruction based on semantic nets as the representation of knowledge.
	Bill Woods described Augmented Transition Networks (ATN's) as a representation for natural language understanding.
	Patrick Winston's PhD program, ARCH, at MIT learned concepts from examples in the world of children's blocks.
Early 70's	Jane Robinson & Don Walker established influential Natural Language Processing group at SRI.
1971	Terry Winograd's PhD thesis (MIT) demonstrated the ability of computers to understand English sentences in a restricted world of children's blocks, in a coupling of his language understanding program, SHRDLU, with a robot arm that carried out instructions typed in English.
1972	Prolog developed by Alain Colmerauer.
1973	The Assembly Robotics group at Edinburgh University builds Freddy, the Famous Scottish Robot, capable of using vision to locate and assemble models.
1974	Ted Shortliffe's PhD dissertation on MYCIN (Stanford) demonstrated the power of rule-based systems for knowledge representation and inference in the domain of medical diagnosis and therapy. Sometimes called the first expert system.
	Earl Sacerdoti developed one of the first planning programs, ABSTRIPS, and developed techniques of hierarchical planning.
1975	Marvin Minsky published his widely-read and influential article on

	Frames as a representation of knowledge, in which many ideas about schemas and semantic links are brought together.
	The Meta-Dendral learning program produced new results in chemistry (some rules of mass spectrometry) the first scientific discoveries by a computer to be published in a refereed journal.
Mid 70's	Barbara Grosz (SRI) established limits to traditional AI approaches to discourse modeling. Subsequent work by Grosz, Bonnie Webber and Candace Sidner developed the notion of "centering", used in establishing focus of discourse and anaphoric references in NLP.
	Alan Kay and Adele Goldberg (Xerox PARC) developed the Smalltalk language, establishing the power of object-oriented programming and of icon-oriented interfaces.
	David Marr and MIT colleagues describe the "primal sketch" and its role in visual perception.
1976	Doug Lenat's AM program (Stanford PhD dissertation) demonstrated the discovery model (loosely-guided search for interesting conjectures).
	Randall Davis demonstrated the power of meta-level reasoning in his PhD dissertation at Stanford.
Late 70's	Stanford's SUMEX-AIM resource, headed by Ed Feigenbaum and Joshua Lederberg, demonstrates the power of the ARPAnet for scientific collaboration.
1978	Tom Mitchell, at Stanford, invented the concept of Version Spaces for describing the search space of a concept formation program.
	Herb Simon wins the Nobel Prize in Economics for his theory of bounded rationality, one of the cornerstones of AI known as "satisficing".
	The MOLGEN program, written at Stanford by Mark Stefik and Peter Friedland, demonstrated that an object-oriented representation of knowledge can be used to plan gene-cloning experiments.
1979	Bill VanMelle's PhD dissertation at Stanford demonstrated the generality of MYCIN's representation of knowledge and style of reasoning in his EMYCIN program, the model for many commercial expert system "shells".
	Jack Myers and Harry Pople at University of Pittsburgh developed INTERNIST, a knowledge-based medical diagnosis program based on

	Dr. Myers' clinical knowledge.
	Cordell Green, David Barstow, Elaine Kant and others at Stanford demonstrated the CHI system for automatic programming.
	The Stanford Cart, built by Hans Moravec, becomes the first computer-controlled, autonomous vehicle when it successfully traverses a chair-filled room and circumnavigates the Stanford AI Lab.
	Drew McDermott & Jon Doyle at MIT, and John McCarthy at Stanford begin publishing work on non-monotonic logics and formal aspects of truth maintenance.
1980's	Lisp Machines developed and marketed.
	First expert system shells and commercial applications.
1980	Lee Erman, Rick Hayes-Roth, Victor Lesser and Raj Reddy published the first description of the blackboard model, as the framework for the HEARSAY-II speech understanding system.
	First National Conference of the American Association of Artificial Intelligence (AAAI) held at Stanford.
1981	Danny Hillis designs the connection machine, a massively parallel architecture that brings new power to AI, and to computation in general. (Later founds Thinking Machines, Inc.)
1983	John Laird & Paul Rosenbloom, working with Allen Newell, complete CMU dissertations on SOAR.
	James Allen invents the Interval Calculus, the first widely used formalization of temporal events.
Mid 80's	Neural Networks become widely used with the Backpropagation algorithm (first described by Werbos in 1974).
1985	The autonomous drawing program, Aaron, created by Harold Cohen, is demonstrated at the AAAI National Conference (based on more than a decade of work, and with subsequent work showing major developments).
1987	Marvin Minsky publishes <i>The Society of Mind</i> , a theoretical description of the mind as a collection of cooperating agents.

1989	Dean Pomerleau at CMU creates ALVINN (An Autonomous Land Vehicle in a Neural Network), which grew into the system that drove a car coast-to-coast under computer control for all but about 50 of the 2850 miles.
1990's	Major advances in all areas of AI, with significant demonstrations in machine learning, intelligent tutoring, case-based reasoning, multi-agent planning, scheduling, uncertain reasoning, data mining, natural language understanding and translation, vision, virtual reality, games, and other topics.
	Rod Brooks' COG Project at MIT, with numerous collaborators, makes significant progress in building a humanoid robot
Early 90's	TD-Gammon, a backgammon program written by Gerry Tesauro, demonstrates that reinforcement learning is powerful enough to create a championship-level game-playing program by competing favorably with world-class players.
1997	The Deep Blue chess program beats the current world chess champion, Garry Kasparov, in a widely followed match.
	First official Robo-Cup soccer match featuring table-top matches with 40 teams of interacting robots and over 5000 spectators.
Late 90's	Web crawlers and other AI-based information extraction programs become essential in widespread use of the world-wide-web.
	Demonstration of an Intelligent Room and Emotional Agents at MIT's AI Lab. Initiation of work on the Oxygen Architecture, which connects mobile and stationary computers in an adaptive network.
2000	Interactive robot pets (a.k.a. "smart toys") become commercially available, realizing the vision of the 18th cen. novelty toy makers.
	Cynthia Breazeal at MIT publishes her dissertation on Sociable Machines, describing KISMET, a robot with a face that expresses emotions.
	The Nomad robot explores remote regions of Antarctica looking for meteorite samples.

Practical applications of AI techniques

Whilst progress towards the ultimate goal of human-like intelligence has been slow, many spinoffs have come in the process. Notable examples include the languages [LISP](#) and [Prolog](#), which were invented for AI research but are now used for non-AI tasks. [Hacker](#) culture first sprang from AI laboratories, in particular the [MIT AI Lab](#), home at various times to such luminaries as McCarthy, Minsky, [Seymour Papert](#) (who developed [Logo](#) there), [Terry Winograd](#) (who abandoned AI after developing [SHRDLU](#)).

Many other useful systems have been built using technologies that at least once were active areas of AI research. Some examples include:

- [Chinook](#) was declared the Man-Machine World Champion in [checkers \(draughts\)](#) in 1994.
- [Deep Blue](#), a chess-playing computer, beat [Garry Kasparov](#) in a famous match in 1997.
- InfoTame, a text analysis search engine developed by the KGB for automatically sorting millions of pages of communications intercepts.
- [Fuzzy logic](#), a technique for reasoning under uncertainty, has been widely used in industrial control systems.
- [Expert systems](#) are being used to some extent industrially.
- [Machine translation](#) systems such as [SYSTRAN](#) are widely used, although results are not yet comparable with human translators.
- [Neural networks](#) have been used for a wide variety of tasks, from [intrusion detection systems](#) to computer games.
- [Optical character recognition](#) systems can translate arbitrary typewritten European script into text.
- [Handwriting recognition](#) is used in millions of [personal digital assistants](#).
- [Speech recognition](#) is commercially available and is widely deployed.
- [Computer algebra systems](#), such as [Mathematica](#) and [Macsyma](#), are commonplace.
- [Machine vision](#) systems are used in many industrial applications ranging from hardware verification to [security systems](#).
- AI Planning methods were used to automatically plan the deployment of US forces during Gulf War I. This task would have cost months of time and millions of dollars to perform manually, and DARPA stated that the money saved on this single application was more than their total expenditure on AI research over the last 50 years.

The vision of artificial intelligence replacing human professional judgment has arisen many times in the history of the field, in [science fiction](#) and today in some specialized areas where "[expert systems](#)" are used to augment or to replace professional judgment in some areas of engineering and of medicine.

Hypothetical consequences of AI

Some observers foresee the development of systems that are far more intelligent and complex than anything currently known. One name for these hypothetical systems is **artilects**. With the introduction of artificially intelligent non-deterministic systems, many [ethical](#) issues will arise. Many of these issues have never been encountered by [humanity](#).

Over time, debates have tended to focus less and less on "possibility" and more on "desirability", as emphasized in the "[Cosmist](#)" (versus "[Terran](#)") debates initiated by [Hugo de Garis](#) and [Kevin Warwick](#). A Cosmist, according to de Garis, is actually seeking to build more intelligent successors to the human species. The emergence of this debate suggests that desirability questions may also have influenced some of the early thinkers "against".

Some issues that bring up interesting ethical questions are:

- Determining the sentience of a system we create.
 - [Turing test](#)
 - [Cognition](#)
 - Why do we have a need to categorize these systems at all?
- Can AI be defined in a graded sense?
- Freedoms and rights for these systems
- Can AIs be "smarter" than humans in the same way that we are "smarter" than other animals?
- Designing systems that are far more intelligent than any one human
- Deciding how many safe-guards to design into these systems
- Seeing how much learning capability a system needs to replicate human thought, or how well it could do tasks without it (e.g. [expert systems](#))
- [The Singularity](#)
- Effect on careers and jobs. The problems may resemble problems seen under [free trade](#)

Sub-fields of AI research

- [Combinatorial search](#)
- [Computer vision](#)
- [Expert system](#)
- [Genetic programming](#)
- [Genetic algorithm](#)
- [Knowledge representation](#)
- [Machine learning](#)
- Machine planning
- [Neural network](#)
- [Natural language processing](#)
- [Program synthesis](#)
- [Robotics](#)
- [Artificial life](#)
- [Artificial being](#)
- [Distributed artificial intelligence](#)
- [Swarm Intelligence](#)