What is this?

This info-mural is one of seven "argumentation maps" in a series that explores Turing's question: "Can computers think and/or will they ever be able to?" Argumentation mapping is a method that provides:

- a method for portraying major philosophical, political, and pragmatic debates
- a summary of an ongoing, major philosophical debate of the 20th century
- a new way of doing intellectual history.

What does it contain?

Altogether the seven maps:

- summarize over 800 major moves in the debates threaded into claims, rebuttals, and counterrebuttals
- 97-130 arguments and rebuttals per map
- 70 issue areas in the 7 maps
- 32 sidebars history and further background

The argumentation maps:

- arrange debate so that the current stopping point of each debate thread is easily seen
- identify original arguments by over 380 protagonists worldwide over 40 years
- make the current frontier of debate easily identifiable
- provide summaries of eleven major philosophical camps of the protagonists (or schools of thought).

How do I get a printed copy?

You can order artist/researcher signed copies of all seven maps from www.macrovu.com for $500.00 plus shipping and handling.

1 Can Computers Think?

The History and Status of the Debate — Map 1 of 7

An Issue Map™ Publication

Start Here

1 Alan Turing, 1950
Yes, machines can (or will be able to) think. A computational system can possess all important elements of human thinking or understanding.

I believe that at the end of the century ... one will be able to speak of machines thinking without expecting to be contradicted.

Alan Turing

Can computers have free will?

3 Humans also lack free will. Whether or not machines have free will is irrelevant to the issue of whether machines can think. Humans think and they don't have free will. People are just highly complex machines are. So machines may yet be able to think.

supporting | supported by

4 Nozick, Steven, 1964.
Humans are programmed. If you accept determinism, then you accept that nature has programmed you to behave in certain ways in certain situations. Nature has thus been programming or rather than the programming a computer to do so.

5 Free will is an illusion of experience. We may think we are free, but that is an illusion created by our experience. Actually, we are determined to do what we do by our underlying neural machinery.

According to the modern scientific view, there is simply no room at all for "freedom of the human will" (1986, p. 306).

[H]uman beings are slaves of brute matter, subject to the same deterministic laws of biochemical and neuronal factors. What we see is the illusory nature of free will (1985, p. 189).

Marvin Minsky | Geoff Simons

6 Philip Johnson-Laird, 1988a.
Free will results from a multilevel representational structure. A multilevel representational structure is capable of producing free will. The system must have levels that:
- representing alternatives for, e.g., go to dinner, read, take a walk;
- representing the mental process by which option is take (e.g., choose the one that makes me happy, choose by flipping a coin);
- representing a method for deciding which decision-making process to use (e.g., follow the most "rational" method, follow the fastest method).

Computers that have been programmed with such multilevel structures can exhibit free will.

Options for action | Decision-making process | Grounds for choosing a decision-making process

Philip Johnson-Laird

7 Geoff Simons, 1985
Free will is a decision-making process. Free will is a decision-making process characterized by selection of options, discrimination between clusters of data, and choice between alternatives. Because computers already make such choices, they possess free will.

8 Geoff Simons, 1985
Conditional jumps constitute free will. The ability of a system to perform conditional jumps when confronted with changing information gives it the potential to make free decisions. For example, a computer may or may not "jump" when it interprets the instruction "proceed to address 9739 if the contents of register 3 are negative." The ability to make a decision making that results from this ability frees the machine from being a mere puppet of the programmer.

10 Jack Copeland, 1993
Free will arises from random selection of alternatives in nil preference situations. When an otherwise deterministic system makes a random choice in a nil preference situation, that system is free with respect to that choice. Human beings are free when they make a choice between a variety of equally preferred alternatives (for example, whether to eat one orange or another from a bag of equally tasty oranges). Such choices are free because they are not deterministic factors, but "when the dice roll," the choice is made freely.

9 Alan Turing, 1951
Machines can exhibit free will by making random selections. Free will can be produced in a machine that makes random selections, for example, by sampling random noise.

11 Randomization sacrifices responsibility. Machines that make decisions based on random choices have no responsibility for their actions, because it is then a matter of chance that they act one way rather another. Agents with responsibility is necessary for free will, such machines lack free will.

12 J. L. Jarvis, 1954
Free will is necessary for moral responsibility. Randomness and moral responsibility are incompatible. We cannot be morally responsible for what happens by chance. Neither can we be responsible for what is predetermined. Because an adequate account of moral responsibility should be grounded in the notion of free will, randomness cannot adequately ground moral responsibility.

13 Jack Copeland, 1993
Random choice and responsibility are compatible. Agents are free (and morally responsible) in a nil preference situation (one in which all choices are equally preferred) when they make a random choice. For example, one randomly chooses to kill 1 of 7 hostages. He chooses at random, but he is still responsible for the death of the hostage he picks, because he was responsible for taking the people hostage in the first place. Agents are free and have responsibility if the choice is between alternatives of differing ethical value.

14 The helplessness argument. When agents (human or machine) make random choices, they are not free will, because their choices are then beyond their control. As Skaller (1990) puts it, the agent is "at the helpless mercy of those emotions within him which control his behavior."

15 Jack Copeland, 1993
The Turing randomizer is only a tiebreaker. The randomizer in a Turing machine is only a tiebreaker, a device that renders precise control all decision making — for example, the decision of whether to wait at the curb or jump on in front of an oncoming truck. All the Turing randomizer does is determine what a machine will do in those situations in which options are equally preferred.

2 Computers can't have free will. Computers do what they have been designed or programmed to do. They lack the free will is necessary for thought. Therefore, computers can't think.

free will: The ability to make voluntary, reasonable decisions. Freely made decisions are independent of external influences (such as factors as environment and conditioning) constraint.

16 Turing's "Computing Machinery and Intelligence" (1950) is a free will problem. Turing's classic paper is a free will problem in disguise. Turing's "imitation game" is a test for free will in a machine. Turing's "Learning machines" section is devoted to the free will issue.

17 Computers only exhibit the free will of their programmers. Computers can't think because they cannot act except as they are determined to by their designers and programmers.

18 Geoff Simons, 1985.
Some computers can program themselves. Automatic programming systems can write programs by themselves. Learning some of the same heuristics that human programmers use. They specify the task that the program is to perform, choose a language to write the program in, write the program, test the program to see if it is applied to, and make use of information about various programming strategies. Programs written by such APs are not written by humans, and so computers that run these programs do not just mirror the free will of humans.

19 Fred Zill, 1999
Programmed humans and programmed robots have psychological states. Computer programmed robots have no psychological states of their own. They may act as if they have such states, but only because their programmers have psychological states and have programmed the robots to act accordingly.

20 Noonan Smart, 1964
Programmed humans have psychological states. Human beings are programmed by nature and yet have psychological states. Thus, if determinism is true, we have a counterexample to the claim that programmed entities can't have psychological states. [Smart, "Human Are Programmed," Box 4.

21 Fred Zill, 1999
The record player argument. A robot "plays" its memory like a record player plays a record. It is pre-programmed in a fixed set of ways. For example, "When we laugh at the joke of a robot, we are really appreciating the wit of a human programmer, and not the wit of the robot" (Putnam, 1964, p. 678).

22 Hilary Putnam, 1964
The robot learning response. A robot could be programmed to learn and to modify its behaviors by learning in the same way people do. For example, a program that learned to tell jokes could learn simply by repeating jokes the programmer liked the most. In memory, but would be inventing jokes in the same way humans do.

23 Fred Zill, 1999
The reprogramming argument. Humans can't be programmed in the arbitrary way that robots can be. For instance, a robot can be programmed to act total no matter what its physical state is, but humans can't. Humans are conditioned by only above some kind of emotion. The actions of the robot depend entirely on the whims of the programmer, whereas human behavior is self-determined.

24 Hilary Putnam, 1964
Reprogramming is consistent with free will. The reprogramming argument fails to show that robots lack free will. It shows only that:
- Humans can be reprogrammed without affecting their free will. For example, a criminal might be reprogrammed into a good citizen via a brain operation, but he could still make free decisions after the operation. The operation would not affect his free will.
- Robots cannot always be arbitrarily reprogrammed in the way that the reprogramming argument suggests. For instance, if a robot is psychologically isomorphic to a human, it will be as difficult to reprogram as a human.
- Even if robots can be arbitrarily reprogrammed, this does not exclude them from having free will. Such a robot may still produce spontaneous and unpredictable behavior.

That robot's been reprogrammed but it still acts spontaneously and unpredictably.

25 L. Jonathan Cohen, 1953
Computers do not choose their own rules. We refer to people as "having in mind of their own" the rules or principles, or the rules or commands of others. Computers are in a similar sense, "programmed with rules and follow commands without conscious choice." Therefore, computers lack free will.

Put it over there.

He has no mind of his own, right now. He's acting like a computer.

26 Joseph Rychlak, 1991
Computers can't do otherwise. An agent's actions are free if the agent can do otherwise and yet choose to act as he does. Humans can do otherwise, but robots can't. Humans dialectical reasoning allows an agent to change its goals and thereby act freely. Because machines are not capable of that kind of thinking they are not free.
Note: Also, see the "Can physical symbol systems think dialectically?" arguments on Map 3.

27 Selmer Bringsjord, 1992
Computers can't include that finite machines can't reproduce. Unlike deterministic machines (e.g., Turing machines) persons can be in an infinite number of states in a finite period of time. That infinite capacity allows persons to make decisions that machines could never make.
Note: Bringsjord's argument is flawed; see the "Can automata think?" arguments on Map7. Also, see the "Can computers be persons?" arguments on this map.

Can computers have emotions?

28 Machines can't have emotions. Machines can never be emotional because they can never be angry, happy, depressed, etc.). Emotions are necessary for free will. Therefore, computers can't think.

29 Paul Zill, 1999
The concept of feeling only applies to living things. Because machines are mechanistic artifacts, not organisms, they cannot have feelings.

30 J. J. C. Smart, 1964
Having feelings does not logically imply the possession of sensations. Although we haven't yet come across any nonliving entities with feelings, perhaps there could be such. There is no logical contradiction in the idea of a nonliving being that has feelings.

Rocks: nonliving and no feelings.
Clouds: Conditioning and no feelings.
Stephen ... shock!
Ouch! Ouch!
Well, no nonliving things with feelings ... yet.

31 J. J. C. Smart, 1964
We can imagine artifacts that have feelings. Several causes show that artifacts could have feelings. (1) We can imagine a robot with a pain system. Over time, that human would be both living creatures and artifacts created by God. (2) We could imagine self-replicating mechanisms whose offspring would be living creatures, mechanisms allowing them to evolve. Such mechanisms would be living and at the same time artifacts.

"They are my children. God as well as living creatures."

32 Hilary Putnam, 1964
"Alive" is not definitionally based on structure. Because the definition of "alive" is not based on structure, it is conceivable that creatures made of metal, plastic, vacuum tubes and transistors instead of neurons and blood vessels might have feelings because they might actually be alive.

33 Gunnar Rye, 1989
Machines lack the physiological components of emotion. Machines lack the human physiology that is essential to emotions. For example, the ability to secrete hormones and neurotransmitters. Because machines don't have the physical components of emotion, they can't possess emotions.

34 Aaron Sloman, 1987
Emotions are not essential to intelligence. Human emotions can be implemented on a computer system. Such emotions would be modeled (the emotion's interaction with cognitive states, for example) but the physiological aspects of emotion (which include biochemistry, etc.) would be absent. Such emotional semantics, they are not essential.

That's just a simulation of emotion. It's not essential to intelligence.

35 Joseph F. Rychlak, 1991
Machines can't think dialectically, and dialectical thinking is necessary for emotions. Dialectical thinking is a complicated dialectical circumstances, which require the ability to make judgments about others and gauge oppositions. Machines can't reason in that way, so machines can't experience emotions.
Supported by:
Toronto Systems Cannot Think Dialectically: Map 3, Box 25

Joseph Rychlak

36 Emotions are necessary for thought. Only systems that can be in emotional states can be said to think.

37 Geoffrey Jefferson, 1949
Emotional experience is necessary for thought. The only entities that can possess human abilities are entities that can act on the basis felt emotions. Machines can't think because they can't feel.
Note: Also, see "Mechanisms Can't Possess Human Consciousness," Map 6, Box 10.

38 David Gelernter, 1994
Computers must be capable of emotional association to think. In order to think, a computer must be capable of a full spectrum of thought. Computers may be able to think in a narrow sense, but they lack the full range of human thought. In order to think as humans do they must also be capable of low-end thinking, which is diffuse, analogical, associative. For example, a flower and a flowered dress might be associated in low-end thought by a diffuse set of emotionally charged feelings.

39 Tom Stonier, 1992
Emotional machines need limbic systems. Human intelligence is a function of the interaction of the human limbic system. The limbic system produces emotional states, frames drives, and motivates behavior. Such a system is essential to human intelligence and which guides the activities of all higher animals. Through the development of artificial limbic systems, emotional machines will be attainable in 20-30 years.

40 Hans Moravec, 1988
Artificial minds should mimic animal evolution. The fastest progress in AI research will be made by those who follow the evolutionary path of animals, starting near the bottom of the phylogenetic scale and working upward toward animals with more complex nervous systems.

Should we pretend computers will never be able to think?

19 Anti-*quoted* by Alan Turing, 1950
The theological objection. The consequences of machine thought are too dreadful to accept. We should "stick our heads in the sand" and hope that machines will never be able to think or have souls.

60 Alan Turing, 1950
The heads-in-the-sand objection. The consequences of machine thought are too dreadful to bear. Thinking machines must be inferior beings.

Consolation: This is not the strong objection. It is not even an argument. It is only wishful thinking. It is too trivial to deserve a response. Perhaps it is a matter of conscience to suppose that machines are souls, souls passed from humans to machines during transmigration. The theological doctrine of the transmigration of souls.

Does God prohibit computers from thinking?

61 Anti-*quoted* by Alan Turing, 1950
The theological objection. Thinking is a function of man's immortal soul. God has given souls to humans, but not to machines. Therefore, humans can think, and computers can't.

Here, God is speaking to man.

62 Alan Turing, 1950
The theological objection is weak. This argument is not very compelling, which is as unprovided and arbitrary as the view that men have souls because only men are male. It is like a young thinking theologian has to say he serving God's ends by providing dwellings for souls he creates.

Can computers understand arithmetic?

63 Stanley L. Jaki, 1969; Fred Dretske, 1990; Jerry Fodor and Ernie Lepore, 1992
Brains, not machines, can understand symbols. Even when they perform the operations corresponding to addition, they are merely shuffling symbols that are meaningless to them. They don't understand what numbers are, and thus humans interpret them.
Note: For earlier versions of this claim see Map 3, Box 24, "Computers Can't Think Dialectically" and Map 6, Box 10, "Mechanisms Can't Possess Human Consciousness." Also see "Mechanisms Can't Possess Human Consciousness," Map 6, Box 10.

64 William Rapaport, 1988
Computers can learn arithmetic. Humans possess innate numerical capacities, and yet they can learn arithmetic. Thus, while they do not have an innate ability to add, they can learn.

65 Fred Dretske, 1990
The marijuana-sniffing dog. Computers can't have an infinite number of beliefs because they can't have the symbols being added about. Dogs can't have an infinite number of beliefs because they don't have any meaning because they don't have the symbols that underlie that meaning. Irrational dog, for example, will wag its tail when it sniffs marijuana, but it doesn't know that it is marijuana that it is sniffing. It just wagged its tail, not because the meaning of the smell causes it to wag its tail.

Can computers draw analogies?

66 Computers can't understand analogies. Computers cannot understand analogical arguments. For example, a machine could not understand the sentence, "Sue was the fairy the wand."
Note: Analogy arguments are also discussed by George Lakoff in the "Symbolic Data" arguments on Map 7.

67 Computers have understood analogy. Running models have understood analogies.

Implemented Model

68 Brian Falkenhainer, K. Forbus, and D. Gentner, 1990
SME. SME (Structure-Mapping Engine) finds analogies between domains by a set of match rules. The analogies that match are judged according to the criteria of clarity, richness, abstractness, structural consistency, and direct applicability. Examples include heat and water flow, solar systems and atoms, and so other domains.

Implemented Model

72 Keith Holyoak and Paul Thagard, 1989
ACME. ACME is a connectionist system that discovers cross-domain ecological mappings. The ACME algorithm uses three constraints to seek out those mappings.

69 David Chalmers, Robert French, and Douglas Hofstadter, 1991
Computers can't understand analogies. Computers can't understand analogies because they lack high level representations that are structured with those specific properties that are required for such understanding. Such representations are not delivered hand-tailored to the model. They are the product of interactions between low level and high level components.

70 David Chalmers, Robert French, and Douglas Hofstadter, 1991
Computers can't make rigidly distinguished SME. In order for its analogical mappings to work, SME assumes a rigid structure that distinguishes between objects and relations, and between high level and low level terms, and to make such a rigid distinction. For example, we sometimes conceptually wealth as an object that flows between people, but at other times we conceptualize wealth as a fluid that changes with each transaction we make.

71 David Chalmers, Robert French, and Douglas Hofstadter, 1991
Computers can't make distinctions is too rigid. In SME, relations are treated as if place predicates that can only be mapped to other place predicates. For example, the relation "x orbits y" is treated as if the various contexts and analogies it might be used in. Such a representation would be too bulky for efficient processing.

73 David Chalmers, Robert French, and Douglas Hofstadter, 1991
All-encompassing representations could not be processed. The all-encompassing representations that such models would provide to a computer model would have to encode a vast amount of information about the world in order to cope with the various contexts and analogies it might be used in. Such a representation would be too bulky for efficient processing.

74 David Chalmers, Robert French, and Douglas Hofstadter, 1991
Perception depends on analogy. The ability to see something as something is what high-level analogical processes we lose. For example, Saddam Hussein will be seen as a Hitler by many people because he is seen (or is selected by advertisers to be seen) as an analogous to Adolf Hitler (a ruthless aggressor) or to Robin Hood (a generous crusader).

75 David Chalmers, Robert French, and Douglas Hofstadter, 1991
Perception requires high level representations. Perception requires high level representations that are structured with those specific properties that are required for such understanding. Such representations are developed by humans as they develop a sensory-motor relationship with the world. They are not delivered hand-tailored to the model.

76 David Chalmers, Robert French, and Douglas Hofstadter, 1991
Perception depends on analogy. The ability to see something as something is what high-level analogical processes we lose. For example, Saddam Hussein will be seen as a Hitler by many people because he is seen (or is selected by advertisers to be seen) as an analogous to Adolf Hitler (a ruthless aggressor) or to Robin Hood (a generous crusader).

Sensory-data | Sensory and motor data | Computer model

Are computers inherently disabled?

87 Anti-*quoted* by Alan Turing, 1950
The argument from disability. Machines can never do X, where X is any of a variety of activities that are regarded as distinctly human, for example, being kind, resourceful, beautiful, friendly, having initiative, enjoying strawberries and cream, or making someone fall in love.
Note: A great deal of the debate represented on these maps are devoted to evaluating variations of arguments that machines can't be creative, can't use analogies, can't have free will, and so on. Turing's model also he thought of as supports for this claim.

88 Alan Turing, 1950
Disability arguments derive from our limited experience with machines. We have seen machines write music and play chess, reply mechanically, and so forth, we assume that a machine could never do X, where X is any of a variety of activities. But our experience with machines is limited, and our inductions from a limited base of experience.

89 Anti-*quoted* by Alan Turing, 1950
Computers can't enjoy strawberries and cream. Computers never could enjoy strawberries and cream.

90 Alan Turing, 1950
Machines could be made to enjoy strawberries and cream. Computers might be made that will enjoy strawberries and cream. They might have sensors that would be the pleasure other tastes, such as the possibility of friendship between man and machine.

91 Anti-*quoted* by Alan Turing, 1950
Computers can't make mistakes. Computers differ from humans in that humans can make mistakes, whereas computers can't. They are infallible. Because of this fact, because humans would frequently make mistakes in interacting with humans, computers never do.
Note: For more on the Turing test, see Map 2.

92 Alan Turing, 1950
Computers can make certain kinds of mistakes. Computers can't make mistakes of a certain kind, namely, mistakes that result from the performance of an inadequate program. Such mistakes are the mistakes of a machine that cannot errors of functioning if they are correctly programmed. But machines can commit errors of conclusion, for example, by making faulty inferences based on a lack of adequate information.

93 Anti-*quoted* by Alan Turing, 1950
Computers can't think about themselves. Computers cannot be the subject of their own thoughts.

94 Alan Turing, 1950
Computers can be the subject of their own thoughts. When a computer solves equations, the equations can not be the subject of its thoughts. Similarly, when a computer is used to predict its own behavior or the behavior of other programs, we can say that it is the object of its own thoughts.

95 Anti-*quoted* by Alan Turing, 1950
Computers can't have diversity of behavior. Humans have a great deal more diversity of behavior than machines ever will.

96 Alan Turing, 1950
Diversity of behavior depends only on storage capacity. Great diversity of behavior is possible for machines if they have large enough storage capacities. Because of the immense storage capacities of our brains it is not possible for a machine to have much storage capacity.

Is the brain a computer?

79 John Searle, 1992
Brains are not digitally a digital computer. The synaptic structures that underlie human thought are intrinsic to physics, they are ascribed to physical systems in a way that they are not ascribed to digital computers[1] in general. Therefore, the brain is not a digital computer.

That will can be interpreted as a digital computer.

80 Jack Copeland, 1993
Brains are not universally digitable. Even if it is true that during any interval of time a pattern of molecular movements in a human brain will be such that, for example, the formal pattern of the WordStar computer program, that will not

That will direct support the counterfactuals necessary.

81 John Searle, 1992
Universal realizability is not essential to the arguments. Even without universal realizability, it is true that syntax is observer relative. And this is a crucial point because, nothing, including the brain, is intrinsically a digital computer.

82
Formal programs can be realized in multiple physical media. The same formal program can be realized in a human brain, in a human brain, in brain cells and toilet paper, or in any number of physical implementations. The program is independent of any of the physical implementations. The program is independent of any of the physical implementation is irrelevant.
Note: For more on multiple realizability arguments, see the "Can functional states generate consciousness?" arguments on Map 6 and sidebar, "Formal Systems: An Overview," on Map 7.

83
The biological assumption. The brain is a machine that can think. Its neurological processes are similar or similar enough with the information processes of a computer.
Note: For more on the biological assumption argument see implemented on Map 3 and on Map 5.

is similar to

84 Roger Penrose, 1990
Low-level quantum effects are uncomputable. The biological phenomena that underlie consciousness operate at a level at which quantum effects could exert a significant influence. Quantum effects are not computable, the brain and consciousness are not noncomputational and nonalgorithmic.

I'm not computable.

85 Robin Cooper, 1990
And then a miracle happens ... Penrose does not explain how quantum effects in the brain might affect consciousness. He simply assumes that some miracle happens and the brain are miraculously related.

That's an explanation of a miracle!

86 Herbert Simon, 1995
Quantum effects are irrelevant to symbolic processes. Quantum effects are irrelevant to the study of artificial intelligence because they occur at a low level of organization and are averaged out before they can affect higher level processes.

Can computers be creative?

97
Computers can never be creative. Computers only do what they are programmed to do; they have no creative potential.
Note: Similar debates play out in the "Free Will" arguments.

98 Anti-*quoted* by Alan Turing, 1950
Creativity is not entirely predictable. The belief that computers are entirely predictable arises from the false conception (widespread in philosophy and in mathematics) that humans can be entirely predicted in a formal sense, given sufficient information. But humans have new things in part through the working out of deductive consequences. Similarly, humans don't know everything they know about a subject, and they can be surprised by consequences, and so are things in part by watching them perform their calculations.

That's what you think!

99 Alan Turing, 1950
Machines frequently take us by surprise. Computer users and even experts are often surprised by what the computer does.

Wow! What a surprise!

The machine isn't creative, but the human is creatively surprised.

100 Alan Turing, 1950
Machines frequently take us by surprise. Computer users and even experts are often surprised by what a machine does; that reaction does not mean that the machine is creative or intelligent in itself. It just means that human minds a creative production about what the computer would do, and was then surprised when the computer acted differently.

101 Anti-*quoted* by Alan Turing, 1950
Surprise is a result of human creativity. Surprise is interpreted by what a machine does; that reaction does not mean that the machine is creative or intelligent in itself. It just means that human minds a creative production about what the computer would do, and was then surprised when the computer acted differently.

102 Alan Turing, 1950
The argument from human creativity is a special case of surprise. You could always say that being surprised by a machine's behavior is a special case, rather than from anything original on the other person's or machine's part, of being surprised by a human surprise you tell a joke; here you could argue that the machine's surprise is a kind of interpretation of the joke rather than anything creative on the joke teller's part.

103 Countess of Lovelace, 1842
The analytical engine can't originate anything. The analytical engine was called "The Analytical Engine") could only do what we know how to order it to perform. It can't originate anything, because it can only do what it has been programmed to do.

PHOTO # 32 Charles Babbage's analytical engine. Courtesy of Lovelace

104 Alan Turing, 1950
The analytical engine may have been able to think for itself. Ada Lovelace was perplexed in denying that her machine could think for itself. Lovelace was perplexed because she had no evidence that it was creative. But because the machine was a digital and universal digital computer, it may have had far greater capabilities than she knew about. Because of its storage capacity the analytical engine may have been able to think for itself.

The Analytical Engine

Invented by Charles Babbage circa 1860, the analytical engine was a mechanical computer powered by steam. It could store, which could be programmed by punch cards. In principle, Babbage's analytical engine could do all the calculations a modern electronic computer can, but due to construction and design flaws the machine never worked. Ada (during Babbage's lifetime (several have been constructed since).

Charles Babbage

105 Computers have already been created. Computer models that exhibit creativity or at least models that exhibit creativity have already been developed.

106 Douglas Hofstadter, 1995
The ELIZA effect. The ELIZA effect is a phenomenon that people attribute more intelligence than is warranted by their underlying code. For example, the computerized psychotherapist ELIZA (Weizenbaum, 1965, Map 3, Box 34) gives apparently sympathetic responses to human concerns, but in fact is only utilizing a set of programmed responses.
Note: The ELIZA effect was recognized and described by ELIZA's creator, Joseph Weizenbaum, though he didn't give it that title.

Douglas Hofstadter

Implemented Model

107 U. Gelernter, 1963
The geometry program. The geometry program can prove theorems in high school level geometry theorems, searching for their proofs by means-and analysis. This program uses a set of axioms and a hierarchy of goals and subgoals. To avoid impossible searches the program uses diagrams to select the most promising search paths.

Implemented Model

108 Philip Johnson-Laird, 1988a
The jazz generator. The jazz generator produces chord sequences that conform to improvisse sheets, lines like melodies, and rhythms.

Implemented Model

109 Marguerite Manteiga, 1991
Genetic algorithms can be used to evolve artificial fish. Artificial fish can be born (given ingredients provided it adequately models human rationality, which is the suitable structure necessary for personhood.

Implemented Model

110 Jim Meehan, 1975
TALE-SPIN. TALE-SPIN generates stories with characters that have goals and subgoals dependent on their personalities. It then interprets encounters in such other's plans and can form competitive relationships through interactions with humans. The program also uses "Thematic Units" to make communications between its characters.

I George Ant was very lonely. I looked around the forest for someone to talk to. I walked from his patch of ground across the meadow through the valley to a river.

All white in the bushes I flash snow peaks in the sunlight. Along the Sun has fogged.

Implemented Model

111 Harold Cohen, B. Cohen, **AARON.** AARON produces novel art by selecting a random starting point on a random walk through lists from paint using a complex set of if-then rules.

Implemented Model

112 Margaret Boden, 1990
Computer models of creativity exhibit creativity. Connectionist networks can produce novel patterns without being specifically programmed to do so.
Note: Also, see Map 4.

The computer recognizes that letter A without having been programmed to do so.

Implemented Model

113 Sheldon Klein, 1973
Computers can write novels. Novels can be generates 2,100-word novels from a 36-sentence preliminary plot based on the conflicting motivations of characters and file the model of a mystery story based on the murder mystery at the end.

Lady Buxley was near Jones. Jones disliked Lady Buxley. Jones was Lady Buxley's lover.

114 Margaret Boden, 1977
The book generator is inadequate. The book generator is inadequate for creativity for the following reasons: (1) The stories are disjoint and rambling. (2) The specific motivational patterns are not sufficiently complex and subtle. (3) The identification of the murderer comes as a surprise rather than as a discovery.

Can computers reason scientifically?

115 Computers can't reason scientifically. Computers are unable to think and reason as human scientists do.

116 Harry Collins, 1994
Scientific reasoning requires social agreement. Computers cannot reason scientifically because they are not members of society. Scientific laws are not derived from the application of an algorithm, but are the result of a sociopolitical process of negotiation.

I throwed trash in the wastebasket.

Harry Collins

I throw the trash in the wastebasket.

117 Harry Collins, 1994
The socialization test. The importance of socialization is demonstrated by the "socialization test." Scientific reasoning and discoveries are both given a passage of "stacked up" English text. If a person (or machine) has been socialized into the ways of a society, he/she/it can interpret the text into normal English. If a judge cannot tell which text was artificially generated by machines and which by the human control subject, then the machines have passed the socialization test.
Note: For more on the Turing test, see Map 2.

118 Carl Hempel, 1985
Computers can't introduce new terms or explanatory principles. A computer cannot reason because it cannot introduce new theoretical terms or principles. The meaning of theoretical terms can be expressed using the proper's observational terms and apparatus. Human discovery, by contrast, involves the introduction of new terms and explanatory principles that are not explanatory adequacy principles to be expressed in terms of the vocabulary of "casual models" — a class of mathematical theories poorly understood.

It's vocabulary is fixed by any new terms.

119 Richard Scherrer, 1988
Computers can introduce new terms. Computers can introduce new terms using automated principles of explanatory adequacy. This is a class of computer program that uses explanatory adequacy principles to generate new terms in the domain of "casual models" — a class of mathematical theories poorly understood.

120 Carl Hempel, 1985
Computers can't adequately evaluate hypotheses. A computer model of scientific reasoning would have to use a criterion of preference to choose between hypotheses that account for available data. But such a criterion of preference would be as imprecise and idiosyncratic, or it is unlikely that such a criterion could be implemented on a computer.

121 Computers have already reasoned scientifically. Computer models have reasoned as scientists do, proposing explanatory hypotheses and choosing among them.

Implemented Model

122 Pat Langley, Herbert Simon, Gary Bradshaw, and Jan Zytkow, 1987
BACON. A program for discovering laws from data. BACON's heuristics, which are based on Kepler's law of planetary motion, Galileo's law of uniform acceleration, and Ohm's law of electrical resistance.
Note: The history of BACON program is compact and extends back into the 1960s.

123 Harry Collins, 1994
BACON only works when humans filter its data. Bacon only works when humans filter its data and with scientists who filter the data and thereby predetermine its success. Bacon's success is, in part, an interpretation of its data, and it is doubtful that BACON would produce any interesting results.
Supported by:
"The Front End Assumption Is Dubious," Box 15.

124 R. G. Buchanan, D. H. Smith, W. C. White, R. Gritter, E. A. Feigenbaum, J. Lederberg, and C. Djerassi, 1976
DENDRAL. DENDRAL is an expert system that analyzes and identifies chemical compounds by using spectrometer readings and chemical data. Alan DENDRAL, a component of DENDRAL, has discovered how to synthesize previously unknown chemical compounds. It has been used in the synthesis of chemical analysis. It even has a publication to its credit.

Can computers be persons?

125 Computers can't be persons. Machines can never be persons. They lack ethical status and moral standing and they never can. At best they can display personlike behavior.
Note: Many other arguments about computers have been implemented on this map. Most are placed in other regions to emphasize what special aspect of machinehood or personhood is in question.

126 Julian Pollock, 1989
Computers can be built. An artificial person can be built from physical ingredients provided it adequately models human rationality, which is the suitable structure necessary for personhood.

127 Selmer Bringsjord, 1992
Robots can do intelligent things but will never be persons. Robots can do intelligent things but will never be persons because they lack free will and free will is a necessary condition for personhood.

Personhood: Historical Background

Many contemporary and historical debates have dealt with the question of personhood. The debates have dealt with the status of the fetus as a person. Theologians and lawyers are whether slaves are persons. Theologians have debated whether Indians are persons. Theologians have debated whether women are people and not property.

The question of whether robots are persons has been asked since at least the release of Karel Capek's play *R.U.R.* (Rossum's Universal Robots) in the 1920s. This play — which ends with the robots "killing off" destroy — is about the struggle of intelligent robots to gain their civil liberties.

In the debate over artificial intelligence, personhood again becomes an issue, because if computers are able to think, then their ethical status may have to be reconsidered. Moreover, many artificial intelligence researchers hold the dream of creating artificial persons. Some think that a sufficiently complex computer program will be conscious and already related to the concept of personhood. Some think that a thinking computer would be straight off a person, because we know that thinking is (to some degree) the defining characteristic of a person. Some think that a computer can be a genuine person, because otherwise there would be no "one" doing any thinking.

128 Dwight Van De Vate Jr., 1971
A machine isn't a person because it cannot pass the Turing test. A machine or an individual is not a person until society collectively decides that it is a person. Having a genius, a fresh and blood body, the ability to feel pain, and so forth. If a machine lacks any of these—if, for example, it cannot feel pain—it won't be recognized as or treated as a person.

That ain't a person like me.

That's just some economy machine.

129 Anti-*quoted* by Dwight Van De Vate Jr., 1971
Machines can behave like persons in the imitation game. Machines can behave in ways that are like a person and be treated like a person in an imitation game.
Note: Dwight Van De Vate Jr. argues that "Can the imitation game determine whether computers can think?" arguments on Map 2.

You say? I'm a human, just like you! I mean you may guess I fooled you many times, but I'm really a human!

130 Dwight Van De Vate Jr., 1971
Laboratory experiments isn't enough for full reciprocity of social behavior. A machine in a lab playing the imitation game is not a genuine person because it is really being treated like one. We treated like an artifact in an experiment, but we are supposed and ignore as we see fit.

Let's go to lunch.

131 Dwight Van De Vate Jr., 1971
Reciprocity of social behavior is required for personhood. To be capable of treating others like persons in a genuine way, one must be treated like a person by members of society in a variety of contexts.

you to you to you to you

One of 7 in this Issue Mapping™ series—Get the rest!

The remaining 6 maps in this Issue Mapping™ series can be ordered with MasterCard, VISA, check, or money order. Order by phone (206-780-3612), by fax (206-842-0296), or through the mail (Box 566, 321 High School Rd. NE, Bainbridge Island, WA 98110).

© 1998 R. E. Horn. All rights reserved.
Version 1.0

Legend

The arguments on these maps are organized by links that carry a range of meanings:

supporting | supported by — Arguments that uphold or defend another claim. Examples include: supporting evidence, further arguments, supporting experiments, extensions or qualifications, and implemented models.

responding to | responded to — A charge made against another claim. Examples include: logical negations, counterexamples, attacks on an argument's emphasis, potential dangers an argument might raise, thought experiments, and implemented models.

generalizing to | generalized by — A distinctive reconfiguration of an earlier claim.

Unsupported Territory — Additional arguments — This icon indicates areas of argument that lie on or near the boundaries of the central issue areas mapped on these maps. It marks regions of potential interest for future mapmakers and explorers.

Focus Box: The lowest-numbered box in each issue area is an introductory focus box. The focus box introduces and summarizes the core elements of each issue area and serves as an assumption and sometimes as a general claim with no particular author.

Arguments With No Authors: Arguments that are not attributable to a particular source (e.g., general philosophical positions, broad concepts, common tests in artificial intelligence) are listed with no accompanying author.

Citations: Complete bibliographic citations can be found in the booklet that accompanies this map.

Methodology: A further discussion of argumentation analysis methodology can be found in the booklet that accompanies this map.

Anticipated by — Where this phrase appears in a box, it identifies a potential attack on a previous argument that is raised by the author so that it can be disputed.

As articulated by — Where this phrase appears in a box, it identifies a reformulation of another author's argument. The reformulation is different enough from the original author's wording to warrant the use of the tag. This phrase is also used when the original argument is impossible to locate other than in its articulation by a later author (e.g., word of mouth), or to denote a general philosophical position that is given a special articulation by the protagonist.

The Issue Mapping™ series is published by MacroVU Press, a division of MacroVU, Inc. MacroVU is a registered trademark of MacroVU, Inc. Issue Map and Issue Mapping are trademarks of Robert E. Horn.