② Can The Turing Test Determine Whether Computers Can Think?

The History and Status of the Debate — Map 2 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

2 Alan Turing, 1950
The Turing test is an adequate test of thinking. The Turing test is a test to determine whether a machine can think. If a computer can persuade judges that it is human via teletyped conversation, then it passes the test and is deemed able to think.
Notes:
- Shown here is a standard interpretation of Turing's test. Turing's original test, and Turing's original version of the Turing test takes the form of an imitation game (see "The Imitation Game," Box 4).
- Authors differ in whether they take the Turing test to establish "thinking" or "intelligence."

The Turing Test

1 **In one room** ... a machine answers questions posed by the interrogator.

In another room ... a human control answers question posed by the interrogator.

3 **In a third room** ... the interrogator engages in teletyped (computer) conversation with the contestants in judges' rooms. If a machine can trick the interrogator and the judges into thinking it is a human, then that machine has passed the Turing test.

Machine contestant · Judges · Interrogator · Referee · Human control

3 Alan Turing, 1950
The question "can machines think?" is too vague. The question of whether machines can think is too vague. It cannot be decided by asking what these words commonly mean, because that reduces the question of machine intelligence to a statistical survey like the Gallup Poll. Hence, the question "can machines think?" should be replaced by the more precise question of whether they can pass the Turing test.

This is the Gallup Poll. Do you think that machines can think?

can't *test*: A comparison case for which a theory predicts that test results will be independent of the test procedure or some other factor. In the Turing test the responses of human controls are compared with those of machines. Without such controls we would never know whether observed conversational ability was a result of actual machine intelligence or just something else, such as the interrogator's bias, gullibility, ...

Players in the Turing Test

Machine contestant: the machine that is attempting to pass the Turing test.

Human contestant, or control: a human whose conversation is compared with that of the machine. The control is sometimes also called a "confederate" or "player" or "foil."

Judge: a person who interprets the conversations between interrogator and contestant. Often it is assumed that an interrogator is at the same time a judge.

Referee: a person who watches to make sure contestants ask no unfair questions.

Interrogator: a person who asks questions of the contestants.

Does the imitation game determine whether computers can think?

4 Alan Turing, 1950
The imitation game. Turing's original formulation of his test takes the form of an imitation game, which takes place in 2 stages.

Part I **In one room** ... a man (X) and a woman (Y) each try to convince an interrogator that he or she is a woman.

In another room ... the interrogator tries to determine whether X or Y is the woman, based on their responses.

Interrogator: Will X please tell me the length of his or her hair?
X: My hair is shingled and the longest strands are about 9 inches long.
Y: I am the woman, don't listen to him!

Part II To see whether a machine can think, replace the man (X) in the imitation game with a machine. If the machine can successfully imitate the person, then we say the machine can think.

5 Kenneth Colby, F. Hilf, S. Weber, and H. Kraemer, 1972
The imitation game is flawed. The imitation game is a weak test for several reasons.
- The concept of "womanlikeness" is too vague.
- Control subjects may be inexperienced at deception.
- The test fails to specify what to do when a computer fails to imitate a man imitating a woman. Is the computer a successful imitation of a man failing to imitate a woman? Or is the computer a poor imitation of man? Or what?

6 Jack Copeland, 1993
A man simulating a woman is not a woman. So, if a machine fails the imitation game, he would only have succeeded in simulating a woman. He would not have shown he is a woman. Therefore, the imitation game is not an adequate test.

7 Keith Gunderson, 1964
A box of rocks could pass the toe-stepping game. There is an analogy between the imitation game and a toe-stepping game, in which an interrogator puts his foot through an opening in the wall and tries to determine whether a man or a woman has stepped on his foot. A box of rocks could be programmed to win which a box of rocks hooked up to an electric eye could win the toe-stepping game by producing the same effect as a human foot stepping on an interrogator's toe. However, a box of rocks cannot duplicate real human toe stepping. Similarly, the imitation game just shows that a system can produce the outward effects of conversation, not that it can duplicate real human intelligence.

8 Keith Gunderson, 1964
One example cannot explain thinking. The imitation game attempts to explain thinking in terms of one example: a man's ability to imitate a woman. But this is just one of many different examples of mental activity, such as wondering, reflecting, deliberating, writing sonnets, and so forth. Thinking is "all purpose," whereas playing the imitation game uses just one ability.

Keith Gunderson

9 John Stevenson, 1976
The imitation game is an all-purpose test. Playing the imitation game is a "second order" ability that presupposes many other abilities. Hence, it is misleading to characterize the ability to pass the imitation game as just one example or property of thinking.

10 Keith Gunderson, 1964
The vacuum salesman example. If a vacuum-cleaner salesman demonstrated an all-purpose Swish 000 by sucking up everything in sight, the customer would be unsatisfied. He or she would want to see it suck up mud, straw, paper, and cut hairs from crockets, tight corners, and so forth. One example is not enough to demonstrate the all-purpose Swish 000, just as one example is not enough to demonstrate thinking.

11 James Moor, 1987
The customer could infer that the vacuum was all purpose. The buyer would not have to actually watch the vacuum cleaner do everything in its repertoire to infer that it was a good vacuum. Even if the salesman only demonstrated the vacuum cleaner picking up dust, if the buyer saw that it came with various attachments, he or she could reasonably infer that they work, too.
Supported by
"The Inductive Evidence Interpretation," Box 108.

If a simulated intelligence passes, is it intelligent?

21 *Anticipated by* Jack Copeland, 1993
Simulated intelligence is not real intelligence. Just as simulated diamonds are not real diamonds, a machine that simulates intelligence to pass the Turing test is not really intelligent. In general, a simulated X is not a real X.

22 Jack Copeland, 1993
Some simulations are duplications. Two kinds of simulation can be distinguished: (1) Some lack essential features of the simulated thing, and (2) others duplicate all death lacks essential physiological features of real death. Such simulations are not duplications. (2) A second type captures the essential features of the object simulated despite being produced in a nonstandard way, for example artificial coal. Simulations of this kind are duplications.

I'm dying!!

Simulation of death · Duplication of coal

23 John Searle, 1980b
Simulations are not duplications. A simulation of a fire will not burn. A simulation of metabolism in the body will not metabolize. A simulation of an automobile engine will not get anywhere. Similarly, a simulation of understanding is not actually understanding.

24 Lawrence Carleton, 1984
Simulations are duplications if the inputs and outputs are of the right kind. If a simulation uses the same kinds of inputs and outputs as the simulated phenomenon uses, then it is also a duplication. The Chinese Room, for example, uses the same kinds of inputs and outputs (i.e., symbol strings) to model speaking as humans use when they speak, hence it is a duplication as well as a simulation. Computer simulations of fire, digestion, etc., are different. Because they don't use the right kinds of input and output, they are mere simulations. (See Box 55.)
"The Chinese Room Is More Than A Simulation," Box 55, Map 4.

Other Turing test arguments

25 *Anticipated by* Alan Turing, 1950
ESP would confound the Turing test. Extrasensory perception could invalidate the Turing test in a variety of ways. In a competitor had ESP she could "read in" on the judges's minds and thoughts. If a judge had ESP he could easily discern human from machine by clairvoyance. As Turing put it, "With ESP anything may happen."

26 Alan Turing, 1950
Perform the test in telepathy-proof rooms. To prevent ESP-related problems, the test would have to be carried out in "telepathy-proof" rooms.

27 Michael Agnes, 1971
Turing assumes that the brain is a machine. Where Turing claims that machines can be taught in much the same way that children are, he assumes the brain is a machine. But whether or not the brain is a machine is part of what's at issue, because if brains were machines then obviously some machines (brains) would be able to think. So, Turing begs the question of whether machines can think—he assumes part of what he's trying to prove.

The Turing test doesn't assume that!

28 James Moor, 1976
The Turing test does not depend on the assumption that the brain is a machine. If a machine were to pass the Turing test, it wouldn't tell us anything about the relation between brains and machines. It would only show that some machines are capable of thinking, whether these machines are like brains or not.

29 B. McInnes, 1971
The Turing test misleads AI research. The Turing test motivates AI researchers to accomplish a machine that can fool people into thinking it is an intelligence involved in carrying on a normal conversation. This kind of intelligence is arbitrary in that it is relative to particular cultures. AI researchers should focus their efforts on specific intellectual tasks that can be aggregated to form more generally intelligent systems, such as the ability to recognize patterns, solve differential equations, etc.

Turing test proficiency · Specific intellectual tasks · AI

30 Bonnie Shanon, 1989
The Turing test assumes a representationalist theory of mind. By confining itself to teletyped conversations, the Turing test assumes a representational theory of mind, according to which thinking involves computational operations on symbolic representations. But critics have pointed out various phenomena that cannot be accounted for in this framework. Hence, the Turing test begs the question of whether machines can think because its very presuppositions are inhospitable to the possibility that machines can think because its very presuppositions are inhospitable to the possibility that machines can think because its very presuppositions are inhospitable to the possibility that