

II. CAN MACHINES THINK?

Functionalism seems to imply that machines can have a full mental life. Can they really, just by virtue of their functional organization? If so, they presumably have some kind of moral status; we can't shove them around as we would a lawnmower. But it's not clear how mere functional organization can give a thing the right not to be shoved around.

Turing thinks the question is too vague to be answerable. What counts as thinking? Switch to a precise question. Can a digital computer be competitive at the Imitation Game:

An interrogator uses text messaging to question a pair of "contestants" in another room: a human and a computer program. Her job is to distinguish the human from the computer program using only their answers. She may ask whatever questions she wants. The contestants' job is to convince the interrogator that they are human. (Turing's original version used gender as the variable; this might be important.)

If the computer convinces the interrogator as often as the human, it passes the Turing Test (none have so far: the Loebner prize). Claim: There's nothing more to (or no better proof of) machine thinking than passing this test.

Searle's "Chinese room" argument directed at Strong AI: an "appropriately programmed computer is a mind, in the sense that computers given the right programs can be literally said to understand ..." A thought experiment:

Suppose that I'm locked in a room and given a large batch of Chinese writing ... together with some [English] instructions ... [that] instruct me how to give back certain Chinese symbols with certain sorts of shapes in response to certain sorts of shapes given me ... I get so good at following the instructions for manipulating the Chinese symbols that ... my answers to the questions are absolutely indistinguishable from those of native Chinese speakers ... it seems to me quite obvious that ... I do not understand a word of the Chinese stories (331-2)

1. Systems Reply: It's not Searle who understands Chinese but the system of which he is a part, including the rules and the data bank and the incoming and outgoing paper.

Searle: He can memorize the rules and the data bank. Now all that's added is bits of paper. How could an individual plus paper understand when the individual doesn't?

2. Robot Reply: Tack on "perceptual" and "motor" capacities—a camera and robotic arms—and the system will understand.

Searle: Not enough, for the perceptual information and motor instructions could also be in Chinese; he still doesn't understand.

3. Brain Simulator Reply: Suppose the system's "rules" are based on the way human brains operate. A functional duplicate of a Chinese-speaker's brain surely understands!

Searle: (i) Why do we need to know how the brain works, if functionalism allows multiple realizability? (ii) A complicated brain-like system still wouldn't understand (imagine it made of

toilet paper rolls and marbles and water pipes; or, to take Block's example, the population of China).

4. Purpose Reply (Lycan): Intelligence is the ability to balance a wide variety of goals under a wide variety of outside circumstances. Computers lack intelligence because (i) they rely on humans to feed them information, which greatly limits their flexibility (see Robot Reply), and (ii) they can't develop their own goals.

Searle?: Give the data bank a bunch of goal-specifications, also in Chinese. These play a role in determining the Chinese motor instructions. Goals are chosen on the basis of perceptual input, including input about one's robotic "body." An onsite random number generator takes up the slack when the goal-determining rules yield a unique verdict.

Can a machine think, in Searle's view? "My own view is that only a machine could think, and indeed only very special kinds of machines, namely brains and machines that had the same causal powers as brains ... Whatever else intentionality is, it is a biological phenomenon, and it is as likely to be causally dependent on the specific biochemistry of its origins as lactation, photosynthesis, or any other biological phenomena." But why?

Intentionality (aboutness) is part of mentality but not all. There's also consciousness. Turing quotes from "Jefferson's Lister Oration for 1949":

Not until a machine can write a sonnet or compose a concerto because of thoughts and emotions felt, and not by the chance fall of symbols, could we agree that mind equals brain ... No mechanism could feel ... pleasure at its successes, grief when its valves fuse, be warmed by flattery, be made miserable by its mistakes, be charmed by sex, be angry or depressed when it can't get what it wants.

This seems to mix two kinds of consciousness together: self-awareness = awareness of what is going on in your own mind, phenomenal consciousness = "what it's like" to be in a certain mental state. The first might not be such a problem: the computer can be equipped with cameras directed at its own mechanisms. Some say a state is phenomenally conscious just to the extent we are aware of being in it. ("Higher-order thought" theory of consciousness.) If so phenomenal consciousness reduces to self-awareness.

Return of the Knowledge Argument. Consider a variant of Mary:

- (1) Before her release, Mary knows all the physical and functional facts.
- (2) Then Mary learns a new fact: what yellow experience is like.
- (3) Facts about what experience is like are not physical or functional facts.

Does this show that functionalism is unable to account for the "what it's like" of experience?

Functionalists can agree that to know what yellow experience is like "from the inside," you have to actually have that sort of experience yourself. That's because knowing from the inside is not knowing a special sort of information, it's rather a special sort of knowing: knowledge by acquaintance.

"[I]dentify the qualitative nature of your sensations-of-yellow with that physical feature of the brain state that instantiates to which your mechanisms of introspective discrimination are in fact responding when you judge you have a sensation of yellow" (Churchland). Until Mary introspectively discriminates that brain state, she has only book knowledge of the experience of yellow; afterwards she is acquainted with yellow.

MIT OpenCourseWare
<http://ocw.mit.edu>

24.00 Problems in Philosophy
Fall 2010

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.