

Artificial Intelligence:

The Turing Test and Searle's "Chinese Room"



Artificial Intelligence (AI) or “Can a Computer Think ?”

Definitions

Weak AI is the claim that computers can act in a way that makes it *look like they're thinking*.

Strong AI is the claim that computers *do understand and do have cognitive states*.


One way of understanding Strong AI is expressed:

“There is something that it is like to be that computer”

In other words, the computer is an experiencing subject

The Turing Test

In 1950 Alan Turing devised a test for determining the extent to which a computer could perform human-like conversation.



Imagine a situation in which there are three separate rooms. In two of the rooms there are humans and in one there is a machine. Human 1 can communicate with Human 2 and the machine by typing messages and reading replies. Human 1 does not know who or what is in each room. He/She sends messages and receives replies from both rooms.

The machine is said to have passed the Turing test if Human 1 cannot tell the difference between the responses that the machine gives and those of Human 2. In other words, the machine's behaviour in the conversation is indistinguishable from that of Human 2.

Some proponents of strong AI have claimed that a machine that passes the Turing test is intelligent and can therefore think.

Searle's (1980) Criticism of strong AI (The "Chinese room")

Syntactic processing – manipulation of symbols with no meaning attached to them. No understanding is involved in this process.

Semantic processing – manipulation of symbols that have meaning. Processing information this way involves understanding.

John Searle described the following situation in an attempt to demonstrate that a machine that passed the Turing test could not be viewed as "thinking" or "understanding":

Suppose that, many years from now, we have constructed a computer which behaves as if it understands Chinese. In other words, the computer takes Chinese characters as input, consults a large look-up table (as all computers can be described as doing), and then produces other Chinese characters as output. Suppose that this computer performs this task so convincingly that it easily passes the Turing test. In other words, it convinces a human Chinese speaker that it is a Chinese speaker. All the questions the human asks are responded to appropriately, such that the Chinese speaker is convinced that he or she is talking to another Chinese speaker. The conclusion proponents of strong AI would like to draw is that the computer understands Chinese, just as the person does.

Now, Searle asks us to suppose that he is sitting inside the computer. In other words, he is in a small room in which he receives Chinese characters, looks them up in the look-up table, and returns the Chinese characters that are indicated by the table. Searle notes, of course, that he doesn't understand a word of Chinese. Furthermore, he argues that his lack of understanding goes to show that computers don't understand Chinese either, because they are in the same situation as he is. They are mindless manipulators of symbols, just as he is - and they don't understand what they're 'saying', just as he doesn't.

A computer processes information syntactically, not semantically. A computer does not think.