Jacob Devasier

1/30/2021

Computing Machinery and Intelligence

In this paper, Alan Turing proposes the question “Can machines think?” and provides arguments for and counters arguments against the idea of “machines” “thinking”. But what does it mean for a machine to think? Due to the ambiguous nature of the question, especially in its time, Turing proposes that perhaps the better way to answer the question is with a different problem that comes to the same conclusion as the original. This problem is called The Imitation Game.

The objective of The Imitation Game is for an interrogator to determine, between two participants, who the male participant is and who the female one is. Now suppose one of the participants, i.e. the male one, is replaced by a machine; will the interrogator be able to determine the male is a computer? Turing claims that this problem is, in essence, the same problem as the original question. In addition to the introduction of The Imitation Game, Turing also introduces the idea of the digital computer whereby he describes what we now know as computers.

With this idea comes many critiques and arguments against it which will be touched on further. Many of the objections to Turing’s position hold very well against common arguments provided. Specifically, the first three counterarguments easily hold against their respective arguments. One argument which causes concern and deserves to be expanded on is the fourth, i.e. the argument from consciousness. This argument makes many points but nicely boils down

to “how can one know if anyone is thinking without being that person?” While this works as a decent argument for thinking machines, we can see from current technology that this is not quite enough to strongly refute this argument.

Current architectures for what we consider to be computers “learning” are based in statistics and mathematics. They learn in such a way that is similar to the Chinese room problem in which an English speaker tells a translator to translate a message to Chinese for a different Chinese speaker. The translator is placed in a room and equipped with a book containing every possible combination of words and their perfect translation. When we look at it from outside of the translator demonstrates perfect ability to translate from English to Chinese implying said translator has perfect understanding of the task and the languages. However, when we look inside, we can see in fact that the translator is only looking up translations in a book.

Much like the translator, the thinking machine with modern architecture will perform some deterministic actions by which the lack of consciousness will be exposed. Consider the scenario where we toss a coin in the air and it lands on heads. Did the coin make the conscious decision to land on heads? Any sane person would argue that the coin did not make the decision to land on heads and in fact it was some combination of a rotational force exerted onto it and the gravitational pull of earth causing it to land as heads, but according to this argument, because I am not the coin, I cannot state that the coin has no consciousness. Likewise, despite knowing exactly how a computer would choose the words to deceive the interrogator, this argument forces us to accept that the machine is conscious. Perhaps with more technological advances, this will be also be seen as obviously wrong.

The arguments that follow this are provided equally as strong refutations and do not require extended consideration. Overall, this paper can understandably be pointed to as the paper to introduce the artificial intelligence field and will forever be one of the most influential papers of the 20th century.