

Do Machines Understand? Could Machines Ever Understand?



Douglas Hofstadter
Cognitive Science and Comparative Literature
Indiana University, Bloomington

In 1950, Alan Turing proposed the “Imitation Game”, today known as the Turing Test, as a way of detecting the presence of genuine thinking in a computer. His idea was that if an unseen computer, participating in a reasonably long-lasting no-holds-barred dialogue with a careful human interrogator, could someday trick that person into the mistaken belief that it was a human, then it would be utterly reasonable to declare that the computer was actually thinking.

Today we are at a stage where many people believe that the Turing Test has already been passed, or is at least on the verge of being passed. But is that really so? Indeed, do today’s computers really understand any language at all? What does “understand language” mean?

Consider poetry translation — a quintessential linguistic task — as a challenge for computers. Asking for a translation of a poem (or at least a decent attempt) would be fair game in the Turing Test. How far are we from achieving a human level of performance in that activity? If that’s still hugely far off, then how do computers stack up against humans in translating plain, ordinary prose? What does the astonishing rapidity of Google Translate tell us about computers understanding human language?

What does the stunning recent success of IBM’s “Watson” in playing Jeopardy tell us about computers understanding human language today? What do question-and-answer sessions with Apple’s Siri tell us about computers understanding human language today? What do the performances of “chatbots” such as “Eugene Goostman” in the annual Loebner Competition — supposedly a true realization of the Turing Test — tell us about computers understanding human language today?

What is it that is lacking in today’s computational approaches to understanding language? If we just look attentively at what real human thinking is and what mechanized imitations of it are capable of doing, we will see that the performances of Google Translate, Siri, Watson, and Eugene Goostman, though very impressive on their surface, are really just glittery window-dressing, and that machines still have a very long pathway to follow before they will genuinely *understand* anything at all — and this is a sobering fact that, rather than saddening us in the slightest, should instead make us happy, and should instill in us the deepest respect for the subtlety of our very own and very human minds.

