


being a universal machine. In defending this point of view, Turing referred to what would now be called chaotic effects in the brain and argued that these did not prevent computer simulation. Notably, at this time Turing was also founding a new branch of mathematical biology: He was applying the insights of an applied mathematician who was also one of the first to use a computer for simulating physical systems.

In 1951, however, Turing gave a radio talk with a different take on this question, suggesting that the nature of quantum mechanics might make simulation of the physical brain impossible. This consideration can be traced back in Turing's thought to 1932, when he first studied the axioms of quantum mechanics [see (6)]. Turing then took up renewed interest in quantum theory and noted a problem about the observation of quantum systems (now known as the quantum Zeno effect). With his death, this train of thought was lost, but the serious question of relating computation to fundamental physics has remained.

Since the 1980s, quantum computing has given a practical technological arena in which computation and quantum physics interact excitingly, but it has not yet changed Turing's picture of what is computable. There are also many thought-experiment models that explore what it would mean to go beyond the limits of the computable. Some rather trivi-



ally require that machine components could operate with boundless speed or allow unlimited accuracy of measurement. Others probe more deeply into the nature of the physical world. Perhaps the best-known body of ideas is that of Roger Penrose (7). These draw strongly on the very thing that motivated Turing's early work—the relationship of mental operations to the physical brain. They imply that uncomputable physics is actually fundamental to physical law and oblige a radical reformulation of quantum mechanics.

Superficially, any such theory contradicts the line that Turing put forward after 1945. But more deeply, anything that brings together the fundamentals of logical and physical description is part of Turing's legacy. He was most unusual in disregarding lines between mathematics, physics, biology, technology, and philosophy. In 1945, it was of immediate practical concern to him that physical media could be found to embody the 0-or-1 logical states needed for the practical construction of a computer. But his work always pointed to the more abstract problem of how those discrete states are embodied in the continuous world. The problem remains: Does computation with discrete symbols give a complete account of the physical world? If it does, how can we make this connection manifest? If it does not, where does computation fail, and what would this tell us about fundamental science?

#### References

1. A. M. Turing, *Proc. London Math. Soc.* **52-42**, 230 (1937).
2. M. Davis, *The Universal Computer: The Road from Leibniz to Turing* (Taylor & Francis, Boca Raton, FL, Turing Centenary edition, 2012).
3. A. M. Turing, *Proc. London Math. Soc.* **52-45**, 161 (1939).
4. A. M. Turing, *Mind* **49**, 433 (1950).
5. R. M. French, *Science* **336**, 164 (2012).
6. A. Hodges, *Alan Turing: The Enigma* (Princeton Univ. Press, Princeton, NJ, Turing Centenary edition, 2012).
7. R. Penrose, *The Emperor's New Mind* (Oxford Univ. Press, Oxford, 1989).

10.1126/science.1218417

## COMPUTER SCIENCE

# Dusting Off the Turing Test

Robert M. French

**H**old up both hands and spread your fingers apart. Now put your palms together and fold your two middle fingers down till the knuckles on both fingers touch each other. While holding this position, one after the other, open and close each pair of opposing fingers by an inch or so. Notice anything? Of course you did. But could a computer without a body and without human experiences ever answer that question or a million others like it? And even if recent revolutionary advances in collecting, storing, retrieving, and analyzing data lead to such a computer, would this machine qualify as “intelligent”?

Just over 60 years ago, Alan Turing published a paper on a simple, operational test for

machine intelligence that became one of the most highly cited papers ever written (1). Turing, whose 100th birthday is celebrated this year, made seminal contributions to the mathematics of automated computing, helped the Allies win World War II by breaking top-secret German codes, and built a forerunner of the modern computer (2). His test, today called the Turing test, was the first operational definition of machine intelligence. It posits putting a computer and a human in separate rooms and connecting them by teletype to an external interrogator, who is free to ask any imaginable questions of either entity. The computer aims to fool the interrogator into believing it is the human; the human must convince the interrogator that he/she is the human. If the interrogator cannot determine which is the real human, the computer will be judged to be intelligent.

Revolutionary advances in data capture, storage, retrieval, and analysis revive questions raised by the Turing test.

In the early days of artificial intelligence (AI), the Turing test was held up by many as the true litmus test for computational intelligence (3, 4). However, workers in AI gradually came to realize that human cognition emerges from a web of explicit, knowledge-based processes and automatic, intuitive, “subcognitive” processes (5), the latter deriving largely from humans' direct interaction with the world. It was argued, therefore, that by tapping into this subcognitive substrate—something a disembodied computer did not have—a clever interrogator could unfailingly distinguish a computer from a person (6). By 1995, most serious researchers in AI had stopped talking about machines passing Turing's original, teletype-based test (7), let alone harder versions involving testing visual, auditory, and object-manipulation abilities (8). The Turing

LEAD-CNRS, Université de Bourgogne, Dijon, France.  
E-mail: robert.french@u-bourgogne.fr


test had been, as one researcher put it, “consigned to history” (9).

However, two revolutionary advances in information technology may bring the Turing test out of retirement. The first is the ready availability of vast amounts of raw data—from video feeds to complete sound environments, and from casual conversations to technical documents on every conceivable subject. The second is the advent of sophisticated techniques for collecting, organizing, and processing this rich collection of data. Two deep questions for AI arise from this new technology. The first is whether this wealth of data, appropriately processed, could be used by a machine to pass an unrestricted Turing test.

The second question, first asked by Turing, is whether a machine that had passed the Turing test using this technology would necessarily be intelligent.

Suppose, for a moment, that all the words you have ever spoken, heard, written, or read, as well as all the visual scenes and all the sounds you have ever experienced, were recorded and accessible, along with similar data for hundreds of thousands, even millions, of other people. Ultimately, tactile, and olfactory sensors could also be added to complete this record of sensory experience over time. Researchers at the cutting edge of today’s computer industry think that this kind of life-experience recording will become commonplace in the not-too-distant future (10). Recently, a home fully equipped with cameras and audio equipment continuously recorded the life of an infant from birth to age three, amounting to ~200,000 hours of audio and video recordings, representing 85% of the child’s waking experience (11, 12).

Assume also that the software exists to catalog, analyze, correlate, and cross-link everything in this sea of data. These data and the capacity to analyze them appropriately could allow a machine to answer heretofore computer-unanswerable questions that tap into facts derived from our embodiment or from our subcognitive associative networks, like the finger experiment that began this article or like asking native English speakers whether the neologism “Flugblogs” would be a better name for a start-up computer company or for air-filled bags that you tie on your feet for walking across swamps (6). Someone, somewhere has almost certainly done the finger experi-



ment and may well have posted their observations about it to the Internet—or will do so after reading this article—and this information would be accessible to a data-gathering Web crawler. By extension, if a complete record of the sensory input that produced your own subcognitive network over your lifetime were available to a machine, is it so far-fetched to think that the machine might be able to use that data to construct a cognitive and subcognitive network similar to your own? Similar enough, that is, to pass the Turing test.

Computers are already extremely good at collecting and analyzing data from 8 billion (and counting) Web pages, document databases, TV programs, Twitter feeds, etc. (13). In early 2011, IBM’s Watson (14), a 2880-processor, 80-teraflop (i.e., 80 trillion operations/s) computing behemoth with 15 terabytes of RAM, won a *Jeopardy* challenge against two of the best *Jeopardy* players in history. Watson’s success was attributable, at least in part, to its meticulous study of *Jeopardy*-like answers and questions, but its performance was nevertheless astounding (15). How much would be required to retool Watson for a no-holds-barred Turing test?

The real challenge is not to store countless petabytes (1 million gigabytes) of information, but to selectively retrieve and analyze that information in real time. The human brain processes data in a highly efficient manner, requiring little energy and relying on a densely interconnected network of ~100 billion relatively slow and imprecise neurons. It is still not known to what extent the mechanisms of neuronal firing and the patterns of neuronal interconnectivity are optimal for the analysis of the data stored in the

brain. IBM is betting that it just might be. The company recently unveiled a new generation of experimental “neurosynaptic” computer chips, based on principles that underlie neurons, with which they hope to design cognitive computers that will “emulate the brain’s abilities for perception, action and cognition” (16).

Yes, you say, but data-crunching computers will never be able to think about their own thoughts, which in the final analysis is what makes us human. But there is nothing stopping the computer’s data-analysis processes, themselves, from also being data for the machine. Programs already exist that self-monitor their own data processing (17).

All of this brings us squarely back to the question first posed by Turing at the dawn of the computer age, one that has generated a flood of philosophical and scientific commentary ever since. No one would argue that computer-simulated chess playing, regardless of how it is achieved, is not chess playing. Is there something fundamentally different about computer-simulated intelligence?

#### References and Notes

1. A. Turing, *Mind* **59**, 433 (1950).
2. A. Hodges, *Science* **336**, 163 (2012).
3. H. Dreyfus, *What Computers Still Can’t Do* (MIT Press, Cambridge, MA, 1992).
4. J. Haugeland, *Artificial Intelligence, the Very Idea* (MIT Press, Cambridge, MA, 1985).
5. D. R. Hofstadter, *Metamagical Themas* (Basic Books, New York, 1985), pp. 631–665.
6. R. M. French, *Mind* **99**, 53 (1990).
7. R. M. French, *Trends Cogn. Sci.* **4**, 115 (2000).
8. S. Harnad, *Minds Mach.* **1**, 43 (1991).
9. B. Whitby, in *Machines and Thought: The Legacy of Alan Turing*, P. Millican, A. Clark, Eds. (Oxford Univ. Press, Oxford, 1996), pp. 53–63.
10. G. Bell, J. Gemmell, *Total Recall: How the E-Memory Revolution Will Change Everything* (Dutton, New York, 2009).
11. D. Roy et al., in *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, R. Sun, N. Miyake, Eds. (Erlbaum, Mahwah, NJ, 2006), pp. 2059–2064.
12. [www.media.mit.edu/research/groups/1446/human-speechome-project](http://www.media.mit.edu/research/groups/1446/human-speechome-project)
13. D. Talbot, “A Social-media decoder,” *Technol. Rev.* (Nov/Dec. 2011); [www.technologyreview.com/computing/38910/](http://www.technologyreview.com/computing/38910/).
14. D. Ferrucci et al., *AI Mag.* **31**, 59 (2010).
15. R. Kurzweil, “Why IBM’s *Jeopardy* victory matters,” *PC Mag.* (2011); [www.pcmag.com/article2/0,2817,2376035,00.asp](http://www.pcmag.com/article2/0,2817,2376035,00.asp).
16. See [www-03.ibm.com/press/us/en/pressrelease/35251.wss](http://www-03.ibm.com/press/us/en/pressrelease/35251.wss), posted on 18 August 2011.
17. J. Marshall, *J. Exp. Theor. Artif. Intell.* **18**, 267 (2006).
18. This work was supported in part by ANR grant 10-065-GETPIMA. Thanks to D. Dennett, M. Weaver, and especially M. Mitchell for comments on an early draft of this article.

10.1126/science.1218350