

formance capacity, and that happens to call for all of our robotic performance capacities too: the Total Turing Test (TTT). And, as a bonus, the robotic capacities can be used to *ground* the pen-pal (symbolic) capacities, thereby solving the “symbol grounding problem” [3], which afflicts the pen-pal version of the TT, but not the robotic TTT.¹

In fact, one of the reasons no computer has yet passed the TT may be that even successful TT capacity has to draw upon robotic capacity. A TT computer pen-pal alone could not even tell you the color of the flower you had enclosed with its birthday letter—or indeed that you had enclosed a flower at all, unless you mention it in your letter. An infinity of possible interactions with the real world, interactions of which each of us is capable, is completely missing from the TT (and again, “tricks” have nothing to do with it).

Is the Total Turing Test Total Enough?

Note that all talk about “percentages” in judging TT performance is just numerology. Designing a machine to exhibit 100% Turing indistinguishable performance capacity is an empirical goal, like designing a plane with the capacity to fly. Nothing short of the TTT or “total” flight, respectively, meets the goal. For once we recognize that Turing-indistinguishable performance capacity is our mandate, the Totality criterion comes with the territory. Subtotal “toy” efforts are interesting only insofar as they contain the means to scale up to life-size. A “plane” that can only fall, jump, or taxi on the ground is no plane at all; and gliding is pertinent only if it can scale up to autonomous flight.

The Loebner Prize Competition is accordingly trivial from a scientific standpoint. The scientific point is not to fool some judges, some of the time, but to design a candidate that *really* has indistinguishable performance capacities (respectively, pen-pal performance [TT] or pen-pal + robotic performance [TTT]); indistinguishable to any judge, and for a lifetime, just as yours and mine are. No tricks! The real thing!

The only open questions are (1) whether there is more than one way to design a candidate to pass the TTT, and if so, (2) do we then need a stronger test, the TTTT (neuromolecular indistinguishability), to pick out the one with the mind? My guess is that the constraints on the TTT are tight enough, being roughly the same ones that guided the Blind Watchmaker who designed us (evolutionary adaptations—survival and reproduction—are largely performance matters; Darwinian selection can no more read minds than we can).

Let me close with the suggestion that the problem under discussion is not one of definition. You don’t have to be able to define intelligence (knowledge, understanding) in order to see that people have it and today’s machines don’t. Nor do you need a defini-

1. In a nutshell, the symbol grounding problem can be stated as follows: Computers manipulate meaningless symbols that are systematically *interpretable* as meaning something. The problem is that the interpretations are not intrinsic to the symbol manipulating system; they are made by the mind of the external interpreter (as when I interpret the letters from my TT pen-pal as meaningful messages). This leads to an infinite regress if we try to assume that what goes on in *my* mind is just symbol manipulation too, because the thoughts in my mind do not mean what they mean merely because they are interpretable by someone *else’s* mind: Their meanings are intrinsic. One possible solution would be to ground the meanings of a system’s symbols in the system’s capacity to discriminate, identify, and manipulate the objects that the symbols are interpretable as standing for [1], in other words, to ground its symbolic capacities in its robotic capacities. Grounding symbol-manipulating capacities in object-manipulating capacities is not just a matter of attaching the latest transducer/effector technologies to a computer, however. Hybrid systems may need to make extensive use of analog components and perhaps also neural nets, in order to connect symbols to their objects [5, 6].

tion to see that once you can no longer tell them apart, you will no longer have any basis for denying of one what you affirm of the other.

References

- [1] Harnad, S., (ed.) *Categorical Perception: The Groundwork of Cognition*. New York: Cambridge University Press. 1987.
- [2] Harnad, S., “Minds, Machines and Searle.” *Journal of Theoretical and Experimental Artificial Intelligence* 1: 5–25. 1989.
- [3] Harnad, S., “The Symbol Grounding Problem.” *Physica D* 42: 335–346. 1990.
- [4] Harnad, S., “Other bodies, Other minds: A machine incarnation of an old philosophical problem.” *Minds and Machines* 1: 43–54. 1991.
- [5] Harnad, S., Hanson, S.J. & Lubin, J., “Categorical Perception and the Evolution of Supervised Learning in Neural Nets.” In: *Working Papers of the AAAI Spring Symposium on Machine Learning of Natural Language and Ontology* (DW Powers & L Reeker, Eds.) pp. 65–74. Presented at Symposium on Symbol Grounding: Problems and Practice, Stanford University, March 1991; also reprinted as Document D91-09, Deutsches Forschungszentrum für Künstliche Intelligenz GmbH, Kaiserslautern FRG. 1991.
- [6] Harnad, S., “Connecting Object to Symbol in Modeling Cognition.” In: A. Clarke and R. Lutz (Eds) *Connectionism in Context*, Springer Verlag. 1992.

The Turing Test and The Economist

Stuart C. Shapiro
Department of Computer Science
and Center for Cognitive Science
SUNY at Buffalo
226 Bell Hall
Buffalo, NY 14260 U.S.A.
(716) 645-3935; FAX (716) 645-3464
shapiro@cs.buffalo.edu

The Turing Test

I have always thought the Turing Test a perfectly good test for AI, and I want to say briefly why.

First of all, I am assuming the version of the Turing Test in which the Interrogator interacts with entities A and B, one of which is a human, and the other of which is a computer, and the Interrogator must guess which is which. Furthermore, I have always assumed that the Interrogator knows these facts, an assumption met neither in the cases where Eliza nor where Parry were said to have passed the Turing Test. I have also assumed a reasonable amount of time—not a mere five minutes—and that it would not be significant to the outcome that the computer intentionally be programmed to make mistakes so that it not look “too smart.” (If interrogators guess correctly, and attribute their success to the computer’s making fewer mistakes than the human, then certainly the computer has demonstrated thinking, intelligence, etc.)

I believe the above assumptions to be at least in the spirit of what Turing specified for his imitation game. Perhaps they are so different from what Turing said explicitly that you will conclude that this is no longer the Turing Test, but a replacement, but I don’t think so.

What I may have missed in my thinking about the Turing Test is that sufficiently unsophisticated Interrogators might provide such a weak test that “obviously unintelligent” programs might

pass. I assumed that good interrogators would cover a reasonably large number of fields of knowledge and experience, and would require the contestants to answer questions based on information given by the Interrogator during the test. This last point is also missed by many people, including, I believe, most discussants of Searle's Chinese Room Puzzle. If the Chinese Room is to behave exactly like a Chinese-understanding person, it will have to not only write some output for each input, but maintain a data/knowledge base of information contained in inputs—equivalent to updating the book of rules (program). It is this posing of puzzles in arbitrary domains that is the hardest part of the Turing Test, and a part that no program has yet passed.

I think that Harnad's Symbol Grounding Problem is irrelevant to the adequacy of the Turing Test, though it is an empirical question whether it is relevant to producing an AI. If a computer passes the Turing Test, as I have been describing it, we must credit it with thought and intelligence, regardless of how it came to have its abilities. Remember, the Interrogator may ask questions whose answers seem to require direct sensory experience (e.g., "What is the difference in taste between a McDonald's hamburger and a Wendy's hamburger?" "Under what circumstances would you or wouldn't you be willing to go on an amusement park ride that got you soaking wet?").

Finally, I don't think that Turing's proposal nor AI's success rests on the Turing Test's being passed by a serial digital computer. Early in my computer science studies, I learned the proofs that serial and parallel computers are computationally equivalent, as are digital and analog computers. When a computer scientist talks about a computer program passing the Turing Test, it should be clear that any Church-Turing equivalent computing device will do. The question is, does a human count among those devices.

The Economist Editorial

This discussion was mainly triggered by an editorial in *The Economist* [*The Economist* Vol 324, No. 770 (August 1st 1992) p. 14],

which I will describe for those of you who haven't read it. I believe it is very pro-AI, as is the accompanying article, "White-collar computers" [*ibid.*, pp. 57-58]. The most negative aspects of the editorial are the title, "Artificial Stupidity," and the subtitle, "Creating machines that think like people is a great challenge, but a bad idea." The negative message is two-part: it seems like the Turing Test might be passed by trickery, such as making typing errors (see my above remarks about this); machine intelligence that duplicates human intelligence is not economical since "people are in plentiful supply [and] should a shortage arise, there are proven and popular methods for making more." On the contrary, "the point of using machines ought to be that they perform differently from people, and preferably better ... gradually, this is happening ... machines have learnt a lot ... already machines can match, or better, human performance on many problems ... watching such machines at work, nobody could mistake them for humans—or deny their intelligence" [italics added]. Even on the subject of human-like intelligent machines, the editorial recognizes the research motivation, even if the writer denies an economical motivation, "One day researchers may use the precision and power of computers to re-create human reasoning. In the process they may unravel many mysteries—including, possibly, the roots of human intelligence."

The final paragraph is so stirring, I quote it in full:

The real challenge, then is not to recreate people but to recognise the uniqueness of machine intelligence, and learn to work with it. Surrendering the human monopoly on intelligence will be confusing and painful. But there will be large consolations. Working together, man and machine should be able to do things that neither can do separately. And as they share intelligence, humans may come to a deeper understanding of themselves. Perhaps nothing other than human intelligence—constantly struggling to recreate itself despite crumbling memories and helter-skelter reasoning—could even conceive of something as illogical and wonderful as machines that think, let alone build them and learn to live with them.