

The Turing Test on the Human Stage

Stephen Abbott

Department of Mathematics, Middlebury College; abbott@middlebury.edu

Abstract

Alan Turing's empirical test for machine intelligence requires the computer to be a highly proficient improvisational actor with a rich knowledge of human culture. This curious criteria is an indication of why theater is an especially effective artistic medium for exploring Turing's life and work. Although the new generation of chatbots is calling into question the relevance of the Turing Test as an AI benchmark, playwright John Mighton shows us how Turing's imitation game provides a novel means for understanding human intelligence—and why it may not be so irrelevant after all.

The End of Turing's Test?

The following short skit was penned by none other than Alan Turing:

INTERROGATOR: In the first line of your sonnet which reads "Shall I compare thee to a summer's day," would not "a spring day" do as well or better?

WITNESS: It wouldn't scan.

INTERROGATOR: How about "a winter's day." That would scan all right.

WITNESS: Yes, but nobody wants to be compared to a winter's day.

INTERROGATOR: Would you say Mr. Pickwick reminded you of Christmas?

WITNESS: In a way.

INTERROGATOR: Yet Christmas is a winter's day, and I do not think Mr. Pickwick would mind the comparison.

WITNESS: I don't think you're serious. By a winter's day one means a typical winter's day, rather than a special one like Christmas.

This entertaining piece of theater from the founder of computer science was actually part of an extended argument in defense of his so-called "imitation game." Since its appearance in *Computing Machinery and Intelligence* in 1950, the imitation game has achieved celebrity status under its adopted stage name, the Turing Test.

That celebrity status is now in jeopardy. Over the last few months, the collective sense of wonderment generated by the release of chatbots such as ChatGPT has generated, among other speculations and consternations, a declaration that the fabled Turing Test is now a relic of a past time. (See, for instance, Metz in *NYT*.) To be fair, the Turing Test has not been a driver of mainstream AI research for many years, if it ever was at all, but it has been a long-standing anchor of philosophical debates, in classrooms and barrooms, over the possibility of machine intelligence. In his original paper, Turing predicted that it would take about 50 years until a computer would be declared intelligent, a title it would earn by passing itself off as a human in a free-flowing and unrestricted conversation with a skeptical interrogator. Turing later raised his estimate to a 100 years. The release of ChatGPT in year 72 has called the question.

ChatGPT is constructed specifically *not* to deceive users by pretending to be a human. That said, the dexterity it displays responding to prompts from across the intellectual spectrum is remarkable enough that it would seem to be only a matter of adjusting the parameters and fine-tuning the training to create a formidable Turing Test competitor. And to be clear, no one, including ChatGPT, thinks ChatGPT is doing any intelligent thinking. So was Turing wrong? Maybe, but maybe not. And anyway, it is the wrong question.

The primary weakness of the Turing Test as a benchmark for AI progress is its focus on deception. In their textbook, *Artificial Intelligence: A Modern Approach*, Stuart Russell and Peter Norvig point out that aeronautical engineers “do not define the goal of their field as ‘making machines that fly so exactly like pigeons that they can fool other pigeons.’” This reasonable objection highlights the strangeness of Turing’s criteria. He was asking the computer to venture far outside its organic domain of functions and imitate a human to the point of being indistinguishable from it. Why this particular feat? It was certainly not because Turing thought human intelligence was special in some way. For Turing, the emphasis was on *flexibility* and the process of *imitation*. From Turing’s perspective, a sentient entity of any kind—human or otherwise—could best demonstrate its intellectual acuity by transforming itself into something it was not.

The roots of Turing’s conviction on this point lie in his mathematics. The universal machine was a mathematical tool that Turing invented in the early 1930s while he was a student working on Hilbert’s decision problem. Turing machines were abstract constructs—machines of the imagination—but the dexterity of the universal machine to imitate the function of any other Turing machine eventually provided the blueprint for the physical computers that Turing and others started to build after the war. For Turing, the flexibility of the universal machine also made it a blueprint of the brain. A typical defense of the uniquely soulful nature of human consciousness is to point to our capacity for empathy. Our mind must be more than the sum of its parts, the argument goes, because of our ability to viscerally hold each other in our thoughts. Turing’s counterargument would be that empathy is just a poetic euphemism for what universal machines do by design. From Turing’s perspective, the brain’s ability to generate a model of another human being for the purpose of feeling what they would feel is actually a compelling argument for a mechanical model of the human mind. In this light, the Turing Test becomes a natural corollary. An intelligent machine, whatever else it might do or think or feel, would be capable of emulating other universal machines, including humans, with all of our passions and imperfections.

But, of course, that is distinctly not what the new generation of chatbots is doing. AI engineers effectively found a clever statistical shortcut for generating conversational language that relies on massive sets of practice data. Chatbots are highly trained stochastic parrots, not sentient entities.

So can the Turing Test be saved? And should it be? To answer the first question, one strengthening clarification to Turing’s criteria is to insist that the interrogator bring some special expertise to the task. This is the point Turing intended to illustrate with the short dialogue he wrote about the “summer’s day” sonnet. To determine if the Witness really is Shakespeare, the Interrogator is getting under the hood to figure out if the Witness not only knows how to write a good poem but whether he knows he has written it. As it currently stands, ChatGPT is only marginally competent at handling this next level of scrutiny:

INTERROGATOR: In the first line of the sonnet that begins “Shall I compare thee to a summer’s day,” would not a “spring day” do as well or better?

CHATGPT: The choice of “a summer’s day” in the first line of the sonnet is deliberate and purposeful, and changing it to “a spring day” would not necessarily do as well or better.

INTERROGATOR: How about a “winter’s day.” Do you think Shakespeare’s beloved would want to be compared to a winter’s day?

CHATGPT: It’s unlikely that Shakespeare or anyone else would want to be compared to a winter’s day, as winter is traditionally associated with coldness, darkness, and a lack of life or growth.

INTERROGATOR: Would you say Mr. Pickwick reminded you of Christmas?

CHATGPT: There are certainly some similarities between Mr. Pickwick and the holiday season, particularly in terms of their associations with kindness, generosity, and goodwill.

INTERROGATOR: Yet Christmas is a winter’s day, and I do not think Mr. Pickwick would mind the comparison.

CHATGPT: You raise a good point. While winter as a season is often associated with coldness and barrenness, there are certain cultural and social contexts in which winter and the holiday season can be associated with warmth, generosity, and good cheer.¹

¹ ChatGPT’s responses were truncated in some cases but otherwise unedited.

Being a More Human Being

For a more interesting example of how the Turing Test might look when administered by someone with the proper training, we turn to an actual human playwright. Between 1988 and 1995, John Mighton penned five full-length plays that collectively established his reputation as an intellectual writer with a gift for bringing science to bear on unsettling questions about human identity. After a lengthy hiatus from writing during which Mighton completed a doctorate in mathematics and became a prominent voice in mathematics education, he found he had another play to write.

Half Life is set in a nursing home in the present day. Donald is a cognitive scientist, and the one time we encounter him outside of the nursing home he is seated next to a closed curtain asking questions to an unseen voice that goes by the name Stanley:

DONALD: I understand you're a mathematician, Stanley.

STANLEY: Who told you that?... I think there's some mistake. I'm an artist.

DONALD: Really?

STANLEY: A painter.

DONALD: Are you sure?

STANLEY: Well, I think I would know what I am.

It's clear that Donald has done this before, and he uses Stanley's artistic inclinations as a segue to discuss wallpaper designs, tilings, and patterns in general:

DONALD: My phone number has a very unusual pattern in it.

STANLEY: What is it?

DONALD: 314-159-2653.

(Pause.)

STANLEY: I said, I'm an artist, not a mathematician.

DONALD: You don't have to be a mathematician to appreciate it.

STANLEY: I'm afraid I can't see any pattern.

Taking Stanley at his word, Donald changes course and starts to explore his new friend's family relationships. Stanley's father is deceased, as it turns out, and we hear a moving account of some early childhood memories. Then, just at the moment when the evidence for Stanley's humanity starts to tip the scales in his favor, Donald abruptly circles back to the earlier topic:

DONALD: What's my phone number?

STANLEY: 314-159-2653.

SCIENTIST: Shit.

The mathematicians in the audience all recognize the familiar song of pi's decimal expansion, but Stanley the artist was supposed to pretend that he did not. "Thank you Stanley, that'll be all," Donald says with a smile, confidently concluding that he has been talking to a machine. Poor Stanley is undone. "A machine? Are you crazy?" the computer babbles on, "I'm not a machine. I'm an artist!"

It's not clear at present how well the new large language models might hold up against a clever inquisitor like Donald, but of course, there is no real reason to run the experiment. What this scene from Mighton's play shows us is that the allure of the Turing Test has never been about confirming the possibility of machine intelligence, whatever form that might one day take. The Turing Test has kept our attention over the decades because of the novel ways it shines a light on the essence of our own humanity. In this scene from *Half Life*, for example, we are reminded of the importance to human cognition of *not* retaining

all the information we encounter. “We wouldn't survive if we remembered everything,” Donald tells his fellow researchers at the conclusion of the test.

This is the moment to remember that, in Turing's original description of the imitation game, there are *two* entities responding to the interrogator's questions: one is a computer and the other a person. Imagine the challenge of being the control in the experiment—of being tasked with the responsibility of being more human than the language generating program behind the other curtain. Defending the human race in a Turing Test against this new generation of AI technology is not about deftly switching from poetry to politics or concocting clever turns of phrase. That territory is now claimed by the chatbots. It's about making personal connections, about revealing authentic glimpses of the human journey, and about empathy, as Turing instinctively knew from the start. From this perspective, Turing's test might indeed still be a useful tool to confirm a generalized artificial intelligence, if and when there is a candidate for such a thing somewhere down the road. In the meantime, we do ourselves a favor by holding onto this curious gift from Alan Turing, whose work on computing machines continues to offer new epiphanies about the flesh-and-blood business of everyday life.

With regard to extracting human insights from Turing's mathematics, theater has proven itself to be an especially adept medium. *Breaking the Code*, by Hugh Whitmore, and *Lovesong of the Electric Bear*, by Snoo Wilson are two well-crafted plays that directly engage Turing's biography. In Mighton's *Half Life*, Turing is mentioned in passing, as are codebreaking and a few other related mathematical concepts, but the playwright is primarily focused on the theatrical possibilities of the Turing Test to help him tell his human story. Having primed his audience with the entertaining repartee between Donald and Stanley, Mighton spends the rest of the play applying what we have learned to the residents of the nursing home. An angry Agnes is tormented by childhood memories of abuse that she has only superficially managed to forget. The womanizing Patrick echoes Stanley in the way that he is so convinced of his fabrications that he may not be aware that he is lying. At the heart of Mighton's play, however, is Donald's mother Clara. Clara is showing signs of cognitive loss, which to her son the brain scientist equates to the slow loss of his mother. Assuaging Donald's grief, the reverend at the nursing home laments “the extent to which we judge adults by their function,” and makes a case for a soul that shines through when the mind starts to fail.

In the closing moment of *Half Life*, a nurse gently tucks Clara into her bed. When the nurse pulls Clara's curtain closed and departs, the scene is blocked to replicate the Turing Test setting with Donald and Stanley from earlier in the play. This time, however, it is Clara's perambulations we hear coming from behind the curtain:

CLARA: I knew a Patrick once during the war...You would be too young to remember... Apparently he has a daughter now... Of course, Dad was on the railroad...

With no one else on stage, it becomes the audience's heart-wrenching job to decide whether or not the voice they hear is coming from an intelligent source.

Acknowledgements

Portions of the discussion of *Half Life* are adopted from the forthcoming book, *The Proof Stage* (Princeton, 2023), which contains a more thorough exploration of the theater related to Turing's life and work.

References

- [1] S. Abbott. *The Proof Stage*. Princeton University Press, 2023.
- [2] C. Metz. “How Smart Are the Robots Getting?” *New York Times*, January 20, 2023.
- [3] J. Mighton. *Half Life*. Playwrights Canada Press, Toronto, Ontario, 2005.
- [4] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Pearson, 2021.
- [5] H. Whitmore. *Breaking the Code*. Samuel French, Inc. New York, 1987.
- [6] S. Wilson. *Lovesong of the Electric Bear*. Bloomsbury Methuen Drama, London, 2015.