



There Can Be No Turing-Test--Passing Memorizing Machines

Citation

Shieber, Stuart M. 2014. There can be no Turing-Test--passing memorizing machines. *Philosophers' Imprint* 14(16): 1-13.

Published Version

<http://hdl.handle.net/2027/spo.3521354.0014.016>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:11684156>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

THERE CAN BE NO TURING-TEST-PASSING MEMORIZING MACHINES

Stuart M. Shieber

Harvard University

© 2014, Stuart M. Shieber

*This work is licensed under a Creative Commons Attribution 4.0
International License.*

<www.philosophersimprint.org/014016/>

Abstract

Anti-behaviorist arguments against the validity of the Turing Test as a sufficient condition for attributing intelligence are based on a *memorizing machine*, which has recorded within it responses to every possible Turing Test interaction of up to a fixed length. The mere possibility of such a machine is claimed to be enough to invalidate the Turing Test.

I consider the nomological possibility of memorizing machines, and how long a Turing Test they can pass. I replicate my previous analysis of this *critical Turing Test length* based on the age of the universe, show how considerations of communication time shorten that estimate and allow eliminating the sole remaining contingent assumption, and argue that the bound is so short that it is incompatible with the very notion of the Turing Test. I conclude that the memorizing machine objection to the Turing Test as a sufficient condition for attributing intelligence is invalid.

1. Introduction

Is the Turing Test valid as a sufficient condition for attributing thinking? Only half a dozen years after Turing's *Mind* paper (1950) proposing his test for thinking, Shannon and McCarthy (1956, page vi) foreshadowed a potential anti-behaviorist argument against its validity:

A disadvantage of the Turing definition of thinking is that it is possible, in principle, to design a machine with a complete set of arbitrarily chosen responses to all possible input stimuli. . . . With a suitable dictionary such a machine would surely satisfy Turing's definition but does not reflect our usual intuitive concept of thinking.

Block (1981) established this argument and explored it in detail with his imagined "Aunt Bertha machine", which has recorded within it responses to every possible Turing Test interaction of up to a fixed length. I will call such devices *memorizing machines*. By hypothesis, a memorizing machine can pass a Turing Test of up to the length it is

designed for, while having “the intelligence of a toaster” (Block, 1981, page 21). Block argues that the mere logical possibility of a memorizing machine refutes the Turing Test as a criterion for intelligence.

In this paper, I refine and extend an earlier argument of mine (Shieber, 2007) to show how weak an assumption beyond logical possibility is necessary to refute the memorizing machine argument.¹ I provide a new technique for calculating an upper bound on how long a Turing Test a memorizing machine can pass. This length limit is the *critical Turing Test length* (CTTL), and it depends on what space of possibility we assume. If we entertain any logically possible memorizing machine, the CTTL is infinite of course, since it is logically possible to furnish a memorizing machine with arbitrary computational resources. I show, however, that the CTTL for nomologically possible memorizing machines is well defined, calculable, and in fact extremely short—about 27 seconds. This length is so short that it is hard to imagine carrying out something that could be reasonably called a Turing Test within that time restriction. I conclude that *any* functional Turing Test would be beyond the ability of a nomologically possible memorizing machine. Thus, the memorizing machine argument does not impeach the Turing Test as an empirical condition sufficient for attributing intelligence to a machine.

In the following sections, I clarify two roles for the memorizing machine thought experiment, a conceptual role and an evidentiary role (Section 2). I argue that nomological possibility is an important context within which to view the evidentiary role of the Turing Test (Section 3). I then replicate my previous analysis of the critical Turing Test length based on the age of the universe (Section 4), show how considerations of communication time shorten that estimate and allow eliminating the sole remaining contingent assumption (Section 5), and argue that

1. I assume some familiarity with my earlier argument as presented in that work (Shieber, 2007). In particular, I do not recapitulate the argument that performance of the Turing Test on particular occasions can demonstrate a general capacity to perform similarly. Leveraging the idea of “interactive proof” to that end was the primary contribution of the earlier work.

the bound is so short that it is incompatible with the very notion of the Turing Test (Section 6). I conclude (Section 7) that the memorizing machine objection to the Turing Test as a sufficient condition for attributing intelligence is invalid.

2. The role of the memorizing machine

The memorizing machine thought experiment is important because of its potential impact on two distinct roles of the Turing Test—a conceptual role and an evidentiary role.

2.1 The conceptual role

The memorizing machine argument’s impact on the conceptual role of the Test is most familiar; it undermines a conceptual analysis of intelligence² as a capacity to pass a Turing Test. This is the role Block (1981) proposes the memorizing machine thought experiment for. He addresses the question of whether passing a Turing Test (more generally, the capacity to do so) is *definitional* of intelligence, that is, whether the capacity to pass a Turing Test is a necessary and sufficient condition for intelligence. Since definition is a logical notion, the mere logical possibility of a Turing-Test-passing memorizing machine counterexamples the Turing Test as a definition of intelligence.

Turing himself thought of passing the Test not as definitional of intelligence but as a sufficient condition for attributing intelligence to a subject undergoing the test³—that is, a condition of the form “if the subject can pass a Turing Test, it is intelligent.”

2. As is traditional in the Turing Test literature, as early as Turing’s own usage (Turing, 1950), I use the terms ‘thinking’ and ‘intelligence’ as synonyms.

3. Hereafter, I will use the term ‘sufficient for intelligence’ as a shorthand for “a sufficient condition for attributing intelligence to a subject undergoing the test”.

The game may perhaps be criticised on the ground that the odds are weighted too heavily against the machine. . . . This objection is a very strong one, but at least we can say that if, nevertheless, a machine can be constructed to play the imitation game satisfactorily, we need not be troubled by this objection. (Turing, 1950, page 435)

Block also appreciates this distinction between a definition of intelligence and a sufficient condition.

The defect I have just pointed out in the case against the behaviorist view of intelligence is a moderately serious one, since behaviorists have tended to focus on giving sufficient conditions for the application of mental terms. . . . Turing, for example, was willing to settle for a “sufficient condition” formulation of his behaviorist definition of intelligence. One of my purposes in this paper is to remedy this defect in the standard objections to behaviorism by showing that no behavioral disposition is sufficient for intelligence. (Block, 1981, page 16)

Here Block’s purpose is to *strengthen* the behaviorist position (by clarifying “defects in the standard objections to behaviorism”), so as to make his eventual counterargument against the behaviorist view of intelligence that much stronger. However, Block still addresses (and disposes of) the Turing Test as a *conceptually* sufficient condition.

Memorizing machines are, of course, impractical, because they are subject to a combinatorial explosion (vide Block, 1981, pages 38ff.). But this cuts no ice: Appeal to contingent considerations (such as the impracticality of building such a machine) is irrelevant from the point of view of a conceptual analysis, whether a definition or a conceptually sufficient condition. In short, all that is required for this first role of the memorizing machine thought experiment is that the machine be logically possible and clearly unintelligent, both properties that seem clearly true of the memorizing machine.

By virtue of having to store responses for every potential Turing

Test situation, a memorizing machine can handle Turing Tests of only a limited length, which depends on the amount of storage it uses. That length might be an hour, say, or a month, or a hundred years, but not indefinitely. This length limitation might seem in and of itself problematic for the memorizing machine, but again, there is no ice-cutting from the fact that memorizing machines cannot pass tests of arbitrary length, as it is not part of our conception of intelligence that its possessors be able to behave in ways that exhibit that property indefinitely. We humans ourselves wouldn’t satisfy such a criterion. Nonetheless, the issue of how long a Turing Test a memorizing machine can pass will be crucial in the discussion below. It is a useful exercise at this point to ask yourself how long a Turing Test you think a memorizing machine could handle based on its having some fixed information storage capacity—the storage capacity of all the computers on earth, say, or (as we will develop later) the entire information storage capacity of the universe.

2.2 *The evidentiary role*

The potential power of the memorizing machine argument is not exhausted by its impact on the conceptual role of the Turing Test. It is important as well in consideration of a second role for the Test, not as an operational *conceptual analysis* of intelligence, but as a mechanism for acquiring *robust evidence* in service of coming to a conclusion that attributing intelligence is appropriate. This empirical role for the Turing Test is highlighted by James Moor:

I believe that the significance of the Turing test is that it provides one good format for gathering inductive evidence such that if the Turing test was passed, then one would certainly have very adequate grounds for inductively inferring that the computer could think on the level of a normal, living, adult human being. (Moor, 1976)

Essentially, Moor is arguing for the Turing Test's role in an *abductive* argument for intelligence (Shieber, 2004). Abduction is "inference to the best explanation", and Moor argues that the best explanation for a machine passing a Turing Test is that it is intelligent. In that sense, the Turing Test is an evidentially or empirically sufficient condition for intelligence, even if not a conceptually sufficient condition.

When examining the strength of Turing-Test-passing as evidence—deciding whether the best explanation for the Turing-Test-passing behavior is the machine's intelligence—we need to examine what other explanations might be in play. Certainly, one explanation for the intelligent behavior is that the Turing-Test-passing machine is intelligent. But there may be alternative explanations, explanations consistent with it being an unintelligent machine that just happens to have the capacity to pass Turing Tests. Are there any proposals for such machines? Yes, the memorizing machine. In fact, in discussions of the empirical sufficiency of the Turing Test for intelligence, the memorizing machine holds a special place, because it is the *only* concrete alternative explanation that we have of that sort.⁴ Thus the memorizing machine might vitiate this role of the Turing Test as well. *To the extent that a memorizing machine is possible*, and only to that extent, the Turing Test as a sufficient condition in this evidentiary sense is weakened. So it behooves us to know under what conditions the memorizing machine counterexample is in play, in order to understand in what contexts Turing-Test-passing might continue to serve as a sufficient condition for attributing intelligence.⁵

4. Readers may propose Searle's "Chinese Room" (Searle, 1980) as another potential alternative. However, the Chinese Room is not a concretely specified machine but an abstractly described proceeding. In particular, it is not described at a level at which its nomological possibility is even a well-defined question, and therefore falls outside the present discussion.

5. It is unclear whether Turing himself would have thought of the Test as a conceptually sufficient condition or an evidentially sufficient one. In fact, on its face, his *Mind* paper attempts to avoid any relation whatsoever between the Test and an intelligence criterion, famously deeming the matter to be "too meaningless to deserve discussion" (Turing, 1950, page 442). But as Moor (1976) notes, glib dismissal of the relationship isn't possible. "[I]f Turing intends that

3. Logical versus nomological possibility

Before turning to the primary question that I consider—whether a Turing-Test-passing memorizing machine is nomologically possible—I address an important framing issue. Is nomological possibility a useful context for evaluating the scope of Turing Tests? That is, does it make sense to restrict Turing Tests to nomologically possible subjects?

Once we move from the conceptual realm of definition to the evidentiary realm of sufficient conditions, there is no fact of the matter as to the right set of background assumptions—the space of possibilities—with respect to which we evaluate the sufficiency of the conditions. The question becomes "Sufficient under what assumptions?" The Test might be sufficient for intelligence in certain contexts but not others. In fact, at the moment, passing a Turing Test is transparently sufficient for intelligence, for the simple reason that no machines at present can pass a Turing Test (or even, frankly, come close). It is thus vacuously true that (at present) all Turing-Test-passing machines are intelligent. There can be no false positives given current technology, because there are no positives at all. But this notion of the evidentiary sufficiency of the Turing Test is certainly not what Turing had in mind in proposing the Test:

The short answer is that we are not asking whether all digital computers would do well in the game nor whether the computers at present available would do well, but whether there are imaginable computers which would do well. (Turing, 1950, page 436)

the question of the success of the machine at the imitation game replace the question about machines thinking, then it is difficult to understand how we are to judge the propriety and adequacy of the replacement if the question being replaced is too meaningless to deserve discussion. Our potential interest in the imitation game is aroused not by the fact that a computer might learn to play yet another game, but that in some way this test reveals a connection between possible computer activities and our ordinary concept of thinking." Certainly, since publication of the 1950 *Mind* paper, both views of the relationship have been active areas of discussion.

He was interested not only in a *contingently* accurate test, but in something closer to a characterization of intelligence, even if not a definition. Being human is also a temporally contingently sufficient condition for attributing intelligence, but as a property it begs the very question that Turing was interested in, whether machines can be intelligent.

Turing entertains the idea that a Turing Test might be an appropriate sufficient condition for attributing intelligence not just in current practice, that is contingently, but *robustly*, allowing for a broader range of possibilities than mere contingent possibility. The memorizing machine argument already shows that entertaining all logically possible machines rules out the Turing Test as sufficient for intelligence. But that alone does not end the discussion; there may well be expansive notions of subjunctive possibility for which memorizing machines do not defeat the Turing Test as sufficient for intelligence. How expansive a notion will still allow use of the Turing Test is thus pertinent to our understanding of the Turing Test as sufficient for intelligence.

What possibilities must we be including for a robust test? Certainly, the particularities of current technology are not important. We want to entertain the use of Turing Tests for testing *future* technologies as well. Turing himself made occasional prognostications of future Turing Test performances, for instance, in his predictions in the *Mind* paper (Turing, 1950, page 442):

I believe that in about fifty years' time it will be possible to programme computers, with a storage capacity of about 10^9 , to make them play the imitation game so well that an average interrogator will not have more than 70 per cent. chance of making the right identification after five minutes of questioning.

and a BBC interview (Newman et al., 1952):

Newman: I should like to be there when your match between a man and a machine takes place, and perhaps to try my hand at making up some of the questions. But that will be a long time from now, if the machine is to stand any chance with no

questions barred?

Turing: Oh yes, at least 100 years, I should say.

We thus do not want to predicate the appropriateness of a Turing Test on mere contingent facts; that would be hopelessly weak. At the other extreme of robustness would be a test applicable in any context consistent with the laws of physics. A maximally robust sufficient condition (short of conceptual sufficiency) would be one sufficient for intelligence for any subjects that are nomologically possible. Then, it behooves us to understand *whether a Turing-Test-passing memorizing machine is nomologically possible*. If so, the Turing Test is not sufficient for intelligence under this maximally robust view; the view that the Turing Test is a sound sufficient condition for attributing intelligence must be narrowed to incorporate some further restrictions on possible testing contexts. But if not (as I argue in the following sections), the Turing Test may well be sufficient for intelligence in all nomologically possible worlds, even though passing a Turing Test is inappropriate as a conceptual definition of intelligence.⁶

To date, there has been no proof that a Turing-Test-passing memorizing machine is nomologically impossible. The following sections provide such a proof.

6. I am not claiming that the Turing Test *is* a sound sufficient condition for attributing intelligence in all nomologically possible worlds, but only that it still *may be* so far as the memorizing machine argument goes. Arguments other than the memorizing machine argument might refute the Turing Test as a sufficient condition in this broad context, though I know of no such arguments.

4. Replicating the prior critical length estimate

The notion of a critical Turing Test length comes from my previous work (Shieber, 2007), which explored how weak an assumption beyond logical possibility is necessary to refute the memorizing machine argument.

The force of the memorizing machine argument comes from the fact that the memorizing machine has a general capacity by construction to pass Turing Tests of up to a fixed length. It does so by storing responses to every Test scenario that might be confronted during the Test period. The number of such scenarios is exponential in the length of the Test. Any memorizing machine must therefore possess storage capacity that is exponential in the Turing Test length. So a simple expedient to protect the Turing Test criterion of intelligence would be to require that the machine *not* use storage exponential in the length of the Test. Equivalently, we can require that the Test be at least logarithmic in the storage capacity of the machine.

Let s be the machine's storage capacity and t be the Test length. Requiring that the storage be less than exponential in the Test length amounts to requiring $s < 2^t$. Then, $\log s < t$, that is, the Test length must be longer than the logarithm of the storage. The threshold T , the longest Turing Test that a machine using exponential storage could pass, would then be at the boundary of this inequality, where $T = \log s$. This threshold value is just the critical Turing Test length alluded to above.⁷

But bounding the storage of the machine s appears to require investigation of an internal property of the machine's construction, not something that the behavioral Turing Test could provide. There is no logical constraint on the storage capacity of the machine. It seems that some further constraint is necessary. Perhaps surprisingly, adding the further constraint of mere *physical existence in the current universe* is sufficient to provide a strict limit on the storage capacity of the machine

7. The inequations here are up to a multiplicative constant. The development below, by sticking to consistent units, accounts for this detail.

and hence how long a Turing Test it could pass. This follows from the fact that the universe itself has a finite capacity to store information.

In the calculations below, I fix a standard set of units: bits for information, seconds for time, and light-seconds for distance. I use two conversions between them. First, the spatial density of information in bits per cubic light-second, specified as the information capacity α of a unit sphere of radius 1 light-second, is approximately $4.7633 \times 10^{131} \frac{\text{bits}}{\text{light-second}^3}$. (I base the conversion on a unit sphere rather than a unit cube to simplify the formulas below.) Second, the temporal density of spoken information β , measured in bits per second, is about $16.7 \frac{\text{bits}}{\text{second}}$. Appendix A provides the physical basis of these estimates.

Now, imagine a machine whose size is bounded by a sphere of radius r . (We can think of the Turing Test proper, the interaction between the interrogator and the subject, going on at the center of this sphere.) We can correspondingly bound the information capacity of the machine as $s = \alpha r^3$. Again, as argued above (and in detail previously [Shieber, 2007]), that information capacity can support a memorized Turing Test of logarithmic length, that is, of

$$T_{\text{bits}} = \log \alpha r^3$$

bits.⁸ We can convert the length to seconds by dividing by β and define the supported Turing Test length in seconds to be

$$T_{\text{sec}} = \frac{\log \alpha r^3}{\beta} \quad . \quad (1)$$

We can already provide a bound on the critical Turing Test length, by setting r to be the effective radius of the universe. The age of the universe is 13.798×10^9 years (as currently estimated by Planck Collaboration et al. [2013, page 36]), so the diameter of the universe is

8. The log is taken in base 2 so that information is measured in bits. I later use the notation $\ln(\cdot)$ for natural (base e) log.

bounded by a light cone of that many light-years.⁹ We can thus take the radius of the universe to be half that, 6.899×10^9 light-years. Using this value for r in Equation 1 yields an estimate of the critical Turing Test length of about 37 seconds.¹⁰

This upper bound on the critical Turing Test length is extremely short, and the assumption of physical existence in the current universe, though stronger than nomological possibility, is quite weak. One would be hard-pressed to say that it would not be revealed by the behavior of the machine, since it would be revealed by the mere existence of the machine independent of its behavior.

Nonetheless, there are two problems with this estimate:

1. The estimate is dependent on a contingent fact about the universe, in particular, its age. As time progresses, the upper bound on the size of the universe increases without limit, and so does the critical Turing Test length as estimated in this way. Admittedly, the estimate is quite insensitive to the value of this parameter. Even if the universe were ten times as old, the critical Turing Test length estimated in this way increases only by about half a second. Nonetheless, reliance on a contingent fact of this sort is not only inelegant but violates our intuitions about the time-invariance of any adequate test for intelligence. For instance, had humans evolved at some later point in the future of the universe when the CTTL was, say, hours long, this argument would be quite weakened.

9. Interestingly, Turing makes this same point in a single sentence on a postcard to Robin Gandy from March 1954 on which he aphorizes that “The Universe is the interior of the Light Cone of the Creation” (Turing, 1954).

10. This estimate of 37 seconds is a bit shorter than the one given in the earlier work (Shieber, 2007) (some 43 seconds, described as “less than a minute”), since the numerical parameters are here used more tightly. The prior paper overestimated the physical parameters, often by many orders of magnitude, so as to make the point that even under such wildly optimistic estimates, the critical length is still quite short.

2. The estimate ignores communication time. At the physical scales we are discussing, the time it takes to get information stored at the periphery of the machine to its center is well beyond the length of the Turing Test itself. (Indeed, it is ludicrously long, as long as the age of the universe itself.)

I solve both problems in the next section, by folding communication time into the estimation, leading to a refined and shorter estimate of the critical Turing Test length. This revised estimate also eliminates any contingent assumption, requiring only nomological possibility.

5. Considering communication time

We have one constraint on the Turing Test length, namely that it is bounded by a function of the volume devoted to the machine’s storage, as given by Equation 1. But the distance r must also be sufficiently short that appropriate communication can occur across that distance *within the length of the Turing Test itself*. If we imagine the interaction happening at the center of the bounding sphere of the machine, then to get information from the periphery of the machine to its center requires a time given by r itself. (I have purposefully chosen to work in units of light-seconds and seconds so that the value of r can be considered as a distance or a time via trivial conversion.) If r is longer than the length T of the Turing Test, then information stored at the periphery cannot be used during the Test. Indeed, no information stored beyond distance T can be used. Thus, if the critical Turing Test length were 37 seconds (as calculated above), no storage beyond 37 light-seconds away could be used by the machine, shrinking the effective storage capacity of the machine from a sphere of radius 7×10^9 light-years to one of 37 light-seconds. This tremendous reduction in storage capacity affects our estimate of the critical length, lowering it, as it turns out, from 37 seconds to 27.17 seconds.¹¹

11. It may seem surprising that restricting the size of the machine from the entire universe to 37 light-seconds (less than a tenth the distance from the Sun to Earth) yields only a 10 second reduction in the critical Turing Test length.

Now, with a Turing Test length of 27.17 seconds, we have an even tighter constraint on the radius, and therefore a further reduction in storage capacity, leading to a yet shorter estimate, and so forth.

We seem to have an infinite regress. Fortunately, as in Zeno's paradox, there exists a fixed point to this serial process. Indeed, we've almost reached it. The next estimate in the series is 27.10 seconds, and the one after that is identical up to four significant digits. So taking into account the requirement of communicating from the periphery to the center of the machine shortens the critical Turing Test length to about 27 seconds.

One might imagine that additional constraints of this sort would substantially shorten the critical Turing Test length further. For example, merely communicating information from the periphery to the center of the machine, a distance of r , is not sufficient, for how does the machine notify the pertinent portion of the periphery to provide the information? The signaling from the center to the periphery itself takes communication time, another r . At the very least, the machine would have to send a signal out to the periphery of the sphere and receive information back, a distance of $2r$ within T_{mins} . More generally, we may want to require that k traversals across the machine, a distance of kr , be able to be carried out within the Turing Test length. This provides a further constraint on r , namely,

$$kr \leq \frac{\log \alpha r^3}{\beta} \quad . \quad (2)$$

This is a result of the tremendous growth of exponentials. The space required to store all conversations of a certain length grows exponentially in the length; the extra 10 seconds thus take a phenomenal amount of storage relative to that required by the first 27 seconds. The mathematical reflex of that is the log in the formula.

Appendix B considers the direct solution of this inequation in more detail, rather than the iterative approach used above. Solving Equation 2 when $k = 1$ yields a maximum value for r of 27.10 seconds, corroborating the iterative solution above.

Perhaps surprisingly, increasing the number of round trips does not substantially decrease the estimated CTTL. For instance, requiring a single round trip ($k = 2$) leads to a maximum value of r of 13.46 seconds. A memorizing machine of this radius could, based on Equation 1, support a Turing Test length of up to 26.9 seconds, which is only .2 seconds shorter than our previous estimate. Requiring two round trips ($k = 4$) shortens the CTTL by about another .2 seconds. Even requiring the ability to perform 100 round trips only shortens the CTTL to 25.7 seconds. The estimate decreases exponentially slowly in k , and is thus extraordinarily insensitive to this parameter. In essence, the value of the critical Turing Test length in the neighborhood of about 27 seconds is close to a physical invariant.

This provides us with the answer to the question posed. It would be nomologically impossible to realize a memorizing machine for a Turing Test longer than about 27 seconds; any such machine would be so large that communication across the machine could not be completed within the Turing Test time allotted.

This estimate of the critical Turing Test length is somewhat shorter than that based on the current size of the universe. More importantly, however, this estimate involves even fewer assumptions than the prior estimate. In particular, it is time-invariant. It does not depend on the age of the universe, and therefore does not change as the universe ages.¹² The *only* constraint imposed is nomological possibility.

As an aside, Equation 1 gives us a simple way of estimating the CTTL under various other assumptions as well, by varying α , β , and r .

12. The time-invariance of the estimate isn't strictly true, but in a benign way. In the period before the universe was 27 seconds old, this bound on the critical Turing Test length would have been shorter than 27 seconds. It stabilized to 27 seconds from that point on. Of course, there were no humans back then.

Scenario	Radius	Info. density	CTTL (seconds)
Upper bound	13.46 light-seconds	*	26.92
Real time	1 light-second	*	26.24
Earth size	0.02 light-seconds	*	25.23
City size	1 mile	*	23.09
City size	1 mile	†	5.62

Table 1: Examples of critical Turing Test lengths under varying assumptions. (*) nomological limit on information density; (†) a more plausible information density of 1 terabit per cubic centimeter.

Plugging in other values gives us other estimates, which must be modified to satisfy the constraint on r imposed by Equation 2. I provide a few examples in Table 1, restricting the size of the machine to allowing a round of communication in real time (2 seconds, a radius of 1 light-second), or to roughly the size of the earth, or to the size of a city, and further modifying the information density from the nomological limit of 1 bit per Planck volume to a more plausible (yet still currently unattainable) limit of one terabit per cubic centimeter.

6. The critical Turing Test length is too short

The most important ramification of this new limit on the critical Turing Test length under just the assumption of nomological possibility is this: A Turing Test of 27 seconds is *no Turing Test at all*. Turing cites as one of the principal advantages of the Turing Test method that it “seems to be suitable for introducing almost any one of the fields of human endeavour that we wish to include.” Making use of this potentiality requires more than 27 seconds. A 27-second Turing Test under our assumptions is less than 90 words long. Turing (1950, pages 433–4, 434–5, 446) provides three short snippets from sample Turing Tests; they are 25, 66, and 102 words in length, respectively. These snippets

are thus individually on the order of the CTTL bound, and together are more than twice the bound. Similarly, Turing’s famous prediction of even a partial success on a time-limited Turing Test, quoted above, still envisions a test of five minutes, some ten times longer than the new estimate of the critical Turing Test length. Clearly, Turing did not conceive of a 27-second Turing Test as a valid sufficient condition for attributing intelligence to a machine.

Block also discusses the importance of the length of a Turing Test. He describes the Turing Test as involving “a machine in one room, and a person in another, each responding by teletype to remarks made by a human judge in a third room for some fixed period of time, e.g., *an hour*” (Block, 1981, page 7, emphasis added). He counters the worry that a Turing Test of only an hour might not be sufficient by allowing for longer tests (Block, 1981, page 21): “Note also that [the Aunt Bertha machine’s] limitation to Turing Tests of an hour’s length is not essential. For a Turing Test of *any* given length, the machine could in principle be programmed in just the same way to pass a Turing Test of that length.” He devotes his Objections 7, 7a, and 7b to discussion of the length limitation of memorizing machines. He grants that memorizing machines have a length limit, and that the length of the test may be important, but that an Aunt Bertha machine can be constructed to satisfy any particular length requirement. But a memorizing machine for a Turing Test of an hour, much less any particular length, is not nomologically possible. Our upper bound shows that a memorizing machine could at best pass a Turing Test of 27 seconds.

Of course, a 27-second Turing Test may provide *some* information—at least from an information-theoretic standpoint—about the subject under test, and perhaps even about its human/machine status. But that is not what is at issue for the purpose of viewing the Turing Test as evidentially sufficient for attributing intelligence. Turing Tests need to be defined in such a way that we can conclude from the statistical indistinguishability of human and machine in repeated Tests that the machine is intelligent, not merely that we can conclude some information or other about the machine or obtain some evidence one way

or the other on the human/machine status. Statistical indistinguishability between humans and machines in repeated 27-second tests is not the kind of condition that would lead one to conclude (with Dennett [1985]) that “any computer that actually passes the unrestricted Turing test will be, in every theoretically interesting sense, a thinking thing.” In summary, the convincingness of a Turing Test is not a linear, or even monotonically increasing, function of its length. It is for that reason that no length limit was prescribed for the Test.

The claim that a 27-second Turing Test is no Turing Test at all might be viewed as itself predicated on contingent facts about humans, namely, that we can communicate only at certain rates through verbal interaction. But this contingency is intrinsic. The Turing Test, after all, is about how *people* are to verify that a machine can think via verbal behavior as defined by people. This contingent aspect is an assumption even in the memorizing machine counterargument to the appropriateness of the Turing Test as a sufficient condition for attributing thinking. Shannon and McCarthy’s and Block’s worries about memorizing machines don’t rest on claims that the Turing Test shouldn’t be contingent on human communication and its properties, but rather that *even in that context*, the Turing Test can fail by virtue of a machine that operates within the contingent rules of the Turing Test itself. I have shown to the contrary that this is not true; such a machine is not nomologically possible.

7. Conclusion

Returning to the topic of Section 3, nothing in our argument indicts the logical possibility of a memorizing machine. One could maintain the view that memorizing machines that can pass functional Turing Tests are logically possible, though they are nomologically impossible, and that this is sufficient to invalidate the Turing Test as a test for intelligence. If one thinks of the capacity to pass a Turing Test as definitional of intelligence, then a logically possible counterexample ought to be sufficient to invalidate the definition, since defining a concept is a logical notion.

There Can Be No Turing-Test-Passing Memorizing Machines

But the Turing Test was not proposed as part of a definition of intelligence. Some have promoted the Test as an evidentially sufficient, but not necessary, condition; perhaps Turing himself thought of the Test in this way. In that context, logically possible but nomologically impossible putative counterexamples are not persuasive. I conclude that the memorizing machine objection to the Turing Test as an evidentially sufficient condition for attributing intelligence is invalid.

Appendix A. Estimating conversion factors

We can estimate the conversion constants α and β from Section 4 based on current best estimates of various physical constants.¹³

For α , we follow our earlier methodology (Shieber, 2007) in assuming the ability to store one bit per Planck volume as the limit of nomological possibility.¹⁴ The size of a unit sphere is $\frac{4\pi}{3}r^3$, where we need r to be expressed in light-seconds. To convert Planck lengths to light-seconds, we make use of two measurement standards, the Planck length in meters (1.616×10^{35}), as specified by NIST (Mohr, Taylor, and Newell, 2012, page 1587), and the definition of a meter according to the International System of Units (“the length of the path travelled by light

13. All of the technical results in the paper and the appendices are derived in detail in a supplemental Mathematica document available at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:11684156>.

14. This estimate is itself quite generous.

Current results in quantum gravity yield even smaller estimates of the information-storage capacity of the universe. Work on the so-called *holographic principle* (regarding which see the survey by Bousso [2002] for a review) limits the information stored in a volume based on its surface area rather than volume. ... An important property of this result is that (unlike the estimate of ... one bit per Planck volume) it does not depend on any assumptions about the fine structure of physical theory. It is a pure principle of physics, like relativity; regardless of future discoveries of more and more finely differentiated particles, say, this limit on information content will hold. (Shieber, 2007)

Taking this further limitation into account reduces the CCTL by a factor of 2/3 (since volume grows as the cube of the radius, but surface area only as the square), to about 18 seconds.

in vacuum during a time interval of $1/299\,792\,458$ of a second" [Taylor and Thompson, 2008, page 18]). Putting these together, we have

$$\begin{aligned}\alpha &= \frac{4\pi}{3} \times \frac{1 \text{ bit}}{\text{Planck length}^3} \times \left(\frac{1.616 \times 10^{35} \text{ Planck length}}{m} \right)^3 \\ &\quad \times \left(\frac{2.998 \times 10^8 \text{ m}}{\text{light-second}} \right)^3 \\ &\approx 4.7633 \times 10^{131} \frac{\text{bit}}{\text{light-second}^3} .\end{aligned}$$

For β , we similarly follow our earlier methodology (Shieber, 2007) in assuming a conservative estimate of the entropy of English of about 5 bits per word. Brown et al. (1992) provide an upper bound on the entropy of printed English of 1.75 bits per character. An estimate of 1 bit per character is thus a conservative value for entropy. With an average word length of 5 characters per word, this gives an entropy of 5 bits per word. We take a (fast, and therefore again conservative) estimate of speaking rate to be 200 words per minute.

$$\begin{aligned}\beta &= \frac{5 \text{ bit}}{\text{word}} \times \frac{200 \text{ word}}{\text{minute}} \times \frac{\text{minute}}{60 \text{ second}} \\ &\approx 16.7 \frac{\text{bit}}{\text{second}}\end{aligned}$$

Appendix B. Solving the critical constraint

In Section 5, we showed that the inequation

$$kr \leq \frac{\log \alpha r^3}{\beta}$$

characterizes radii for which the time for k traversals of communication falls within the Turing Test length storable in the corresponding space.

For what values of r is this constraint solvable? That depends, of course, on the values of the various constants. We can get some intuition from further examining the formula.

We simplify to

$$k\beta r \leq \log \alpha r^3 .$$

Exponentiating yields

$$2^{k\beta r} \leq \alpha r^3$$

so

$$r^3 2^{-k\beta r} \geq \alpha^{-1} .$$


Now, this function

$$f(r) \equiv r^3 2^{-k\beta r}$$

is unimodal, as depicted in Figure 1 for $k = 2$. To calculate the maximum, we set the derivative to zero

$$f'(r) = 3r^2 2^{-k\beta r} - k\beta r^3 2^{-k\beta r} \ln 2 = 0$$

and solve for r (ignoring the solution at $r = 0$, corresponding to a

Figure 1: The function $r^3 2^{-2\beta r}$

minimum):

$$\frac{3r^2 2^{-k\beta r}}{k\beta r^3 2^{-k\beta r} \ln 2} = 1$$

$$r = \frac{3}{k\beta \ln 2}$$

The maximum value is thus

$$f_{\max} = f\left(\frac{3}{k\beta \ln 2}\right) = \left(\frac{3}{k\beta \ln 2}\right)^3 2^{-\frac{3}{\ln 2}}$$

and there are no solutions for r if

$$f_{\max} \leq \alpha^{-1}.$$

There Can Be No Turing-Test-Passing Memorizing Machines




Figure 2: The function $r^3 2^{-2\beta r}$, shown on a log-log scale. The upper dashed line is at the maximum value of the function. The lower is at α^{-1} . The region in blue shows the range of r values where a single round trip communication to the periphery of the storage region is possible within the Turing Test length.

For the physically reasonable values of α and β from Section 4 and $k = 2$, f_{\max} is about 10^{-4} and α^{-1} is about 8×10^{-132} . Thus α^{-1} is much smaller than f_{\max} , as depicted in Figure 2, giving rise to a range of solutions for r shown by the blue-filled region. The minimum and maximum values for r that bound the region are just the roots of the equation $f(r) = \alpha^{-1}$, which we solve numerically to yield the extremal values for r of 1.28×10^{-44} and 13.46 light-seconds. For that radius, the critical Turing Test length T given by Equation 1 is 26.9 seconds. (The same method allows calculating r and T for other values of k , α , and β .) Thus a machine with any radius between those two values supports at least a single round of communication within the Turing Test length that that storage provides for. Machines of larger size are *communication-bounded*—too large to communicate with their periphery—and machines of smaller size are *storage-bounded*—too small to store the data required for a Test of length given by the round trip time to their periphery.

Acknowledgements

I thank Michael Brenner for help with the analysis of Appendix B and Ned Block, Bernhard Nickel, and Jason Stanley for helpful discussions on earlier drafts of the paper. They of course bear no responsibility for any remaining errors.

References

- Block, Ned (1981). "Psychologism and Behaviorism". In: *The Philosophical Review* XC.1, pp. 5–43. DOI: 10.2307/2184371.
- Bousoo, Raphael (2002). "The Holographic Principle". In: *Reviews of Modern Physics* 74.3, pp. 825–874. DOI: 10.1103/RevModPhys.74.825. arXiv: hep-th/0203101.
- Brown, Peter F., Stephen A. Della Pietra, Vincent J. Della Pietra, Jennifer C. Lai, and Robert L. Mercer (1992). "An Estimate of an Upper Bound for the Entropy of English". In: *Computational Linguistics* 18.1, pp. 31–40. URL: <http://acl.ldc.upenn.edu/J/J92/J92-1002.pdf>.
- Dennett, Daniel (1985). "Can Machines Think?" In: *How We Know*. Ed. by Michael Shafto. San Francisco, CA: Harper & Row, pp. 121–145.
- Mohr, Peter J., Barry N. Taylor, and David B. Newell (2012). "CODATA Recommended Values of the Fundamental Physical Constants: 2010". In: *Reviews of Modern Physics* 84.4, pp. 1527–1605. DOI: 10.1103/RevModPhys.84.1527. URL: <http://physics.nist.gov/cuu/pdf/RevModPhysCODATA2010.pdf>.
- Moor, James H. (1976). "An Analysis of the Turing Test". In: *Philosophical Studies* 30.4, pp. 249–257. DOI: 10.1007/BF00372497.
- Newman, M. H. A., Alan M. Turing, Sir Geoffrey Jefferson, and R. B. Braithwaite (1952). "Can Automatic Calculating Machines Be Said to Think?" Radio interview, recorded 10 January 1952 and broadcast 14 and 23 January 1952. Turing Archives reference number B.6.
- Planck Collaboration et al. (2013). *Planck 2013 Results. I. Overview of Products and Scientific Results*. arXiv: 1303.5062 [astro-ph.CO].
- Searle, John R. (1980). "Minds, Brains, and Programs". In: *Behavioral and Brain Sciences* 3.03, pp. 417–424. DOI: 10.1017/S0140525X00005756.
- Shannon, Claude E. and John McCarthy, eds. (1956). *Automata Studies*. Princeton, NJ: Princeton University Press.
- Shieber, Stuart M. (2004). "The Turing Test's Evidentiary Value". In: *The Turing Test: Verbal Behavior as the Hallmark of Intelligence*. Ed. by Stuart M. Shieber. Cambridge, MA: MIT Press, pp. 293–296.
- Shieber, Stuart M. (2007). "The Turing Test as Interactive Proof". In: *Noûs* 41.4, pp. 686–713. DOI: 10.1111/j.1468-0068.2007.00636.x.
- Taylor, Barry N. and Ambler Thompson (2008). *The International System of Units (SI)*. NIST Special Publication 330. National Institute of Standards and Technology. URL: <http://physics.nist.gov/Pubs/SP330/sp330.pdf>.
- Turing, Alan M. (1950). "Computing Machinery and Intelligence". In: *Mind* LIX.236, pp. 433–460. DOI: 10.1093/mind/LIX.236.433.
- Turing, Alan M. (1954). *Postcard to Robin Gandy*. Turing Archives AMT D/4/14. URL: <http://www.turingarchive.org/viewer/?id=154&title=14>.