

Is the Turing test good enough? The fallacy of resource-unbounded intelligence.

Virginia Savova
Johns Hopkins University
savova@jhu.edu

Leonid Peshkin
Harvard University
pesha@hms.harvard.edu

Abstract

This goal of this paper is to defend the plausibility of the argument that passing the Turing test is a sufficient condition for the presence of intelligence. To this effect, we put forth new objections to two famous counter-arguments: Searle's "Chinese Room" and Block's "Aunt Bertha." We take Searle's argument to consist of two points: 1) intelligence is not merely an ability to manipulate formal symbols; it is also the ability of relating those symbols to a multi-sensory real-world experience; and 2) intelligence presupposes an internal capacity for generalization. On the first point, while we concede that multi-sensory real-world experience is not captured by the test, we show that intuitions about the relevance of this experience to intelligence are not clear-cut. Therefore, it is not obvious that the Turing test should be dismissed on this basis alone. On the second point, we strongly disagree with the notion that the test cannot distinguish a machine with internal capacity for generalization from a machine which has no such capacity. This view is best captured by Ned Block, who argues that a sufficiently large look-up table is capable of passing any Turing test of finite length. We claim that, contrary to Block's assumption, it is impossible to construct such a table, and show that it is possible to ensure that a machine relying solely on such table will fail an appropriately constructed Turing test.

1 Introduction

The poet James Whitcomb Riley (1849-1916) is often remembered for his formulation of the "duck criterion": "When I see a bird that walks like a duck and swims like a duck and quacks like a duck, I call that bird a duck."

With a similarly practical attitude, many AI researchers would say a machine that talks intelligently and behaves intelligently should be called intelligent. This assumption was formalized by Alan Turing who proposed a behavioral test for determining whether a machine can think [Turing, 1950; Turing *et al.*, 1952]. In its original formulation, the Turing test is an imitation game, in which a machine does its best to imitate a human participant in free-flowing conversation with

a judge. At the end of the game, the judge is asked to point out the human participant. If under repeated sessions of the game, the judge is at chance, the machine should be considered intelligent.

It is fairly obvious that passing the Turing test is not a necessary condition of intelligence. For one, most humans have never been subjected to such a test. Furthermore, an intelligent agent might fail the test due to temporary impairments, such as being sleepy, hungry or on drugs. Finally, an intelligent machine might fail the Turing test because it may believe it is in its best interest not to show how intelligent it is.

Thus, the Turing test debate in the literature has centered around the question of whether the test represents a sufficient condition for intelligence. Turing himself is fairly unconvinced on this issue. His skepticism derives from the conclusion that the question of machine intelligence is very much alike the question of other minds. It is impossible, the solipsist argument goes, to know for sure whether there exist any other conscious beings apart from one's self, since one only has direct access to one's own consciousness. Thus, rather than directly arguing in favor of his test as a sufficient condition, he claims that the test is simply a good approximation to some true characteristic of intelligence, which is both currently unknown and impossible to observe.

But despite Turing's reluctance to adopt a strong stance, philosophers have taken this position seriously in both defending and disproving it, and for a good reason. In the absence of this strong assertion, the power of the Turing test as a diagnostic is severely limited, and its status is reduced to a dubious replacement for the true definition of intelligence. In other words, why would we take the Turing test more seriously than any other arbitrarily formulated criterion? Its only advantage appears to be that it would pick out all humans, who constitute- by assumption- the pool of known intelligent entities. However, we can achieve similar success by replacing the definition of intelligence with the definition of featherless biped. Countless other criteria would do the same, while having little to do with our intuition of what intelligence is as a trait. That is why it is important to figure out the extent to which Turing's criterion can be deemed sufficient for ascertaining intelligence.

Just like Riley's duck criterion, Turing's intelligence criterion is attractive because the decision is based on purely observable, external characteristics. Thus, without having to

know how intelligence arises or what it is made of, we can ascertain its presence purely by observing its interactions with the environment. This gives rise to the foundational claim of AI and cognitive science — that intelligence can be understood and reproduced in an abstract computational device.

But the Turing test's most attractive aspect is precisely the one open for most criticism. Powerful objections have been raised to the effect that the definition of intelligence requires access to the inner workings of the computational mechanism used by the machine, rather than its behavior per se. The objections are illustrated by two famous thought experiments: Searle's Chinese Room and Block's Aunt Bertha.

2 The Chinese Room

Searle (1980) imagines a monolingual English speaker locked in a room with Chinese symbols and a large set of instructions in English about mapping sequences of these symbols to other sequences. His interrogators outside come up with a Chinese story along with a set of questions about the story. They hand the story and the questions to the person inside the room, who then looks up the sequence of symbols in the instruction manual and produces what his interrogators might view as a reply. Despite performing adequately on the verbal Chinese test, the person inside the room does not understand Chinese. Instead, s/he merely appears to understand Chinese, i.e. simulates understanding of Chinese. Similarly, Searle argues, a machine can fake its way out of the Turing Test without satisfying our deep intuition of what it means to be truly intelligent.

2.1 Previous objections to the Chinese Room

There have been a number of objections to the Chinese Room over the years. We will briefly review here the ones most relevant to our current discussion.

The "systems reply" states that even if the person inside the Chinese Room does not understand Chinese, the room with all its content, does. In effect, the person is to the room like the CPU is to the computer [Rey, 1986]. Therefore the correct level of analysis is the room as a whole.

While the "systems reply" by itself may be insufficiently effective, it is an important step. Intelligence is perceived by modern science to be an emergent property of a system, rather than a property of an essential component. Therefore, the correct level of analysis of the Chinese Room is, indeed, the room.

Since it is generally very difficult to have the right intuition about the perceived intelligence of rooms, it is to our advantage to reformulate the setup. Searle's counter-reply allows us a way out: assume that the person was not locked inside a room, but instead had memorized the rules and symbols completely inside his head. Despite this change, we would still be reluctant to say that s/he understands Chinese. We will call this version of the experiment "the Chinese impostor".

The "English reply" states that even if the person does not understand Chinese, s/he understands something - the rule-book, for example. However, Searle claims that this fact is irrelevant since the Chinese room test is about understanding of Chinese in particular, rather than general understanding [Chrisley, 1995].

Finally, the "robot reply" concedes that the Chinese room or the person inside it cannot claim to possess human-level intelligence because it does not interact with the real world. The solution is to let the room receive input from the real world. Searle's objection is that there is no way for the input from the real world to be experienced in a way that a human experiences it [Boden, 1988].

The latter two replies pertain to what we will call "the information content" part of Searle's argument.

2.2 Two challenges from the Chinese Room

Searle's thought experiment appears to conflate two issues, and it is worth teasing those apart. One issue is the issue of information content. The formal symbol stands for an entity in the real world. Obviously, any machine which has no access to the mapping from the formal symbol to the entity cannot be said to understand what the formal symbol means.

The other issue is the type of information storage and access. It is our intuition that there is something profoundly unintelligent about interacting with the environment through a look-up table, or by following an externally generated sequence of steps. The problem with this scenario is that the intelligence observed in the behavior of the agent originates outside of the agent, which is in and of itself incapable of generating novel behavior. This is the issue of generative incapacity.

2.3 Objection to the argument from information content

The question of information content is really the question of the nature of meaning, or real-world knowledge, and its pertinence to our intuition about intelligence. Meaning is the relationship between a formal symbol and the entity (or event) it stands for. Obviously, if one is without access to sensory input from the world, or the correspondence relation between the world and the formal symbols, one cannot claim access to meaning. We will not argue this point.

However, we take issue with Searle's assumption that meaning is central to our notion of intelligence. To clarify, let us suppose that locked in the Chinese room is a Chinese speaker, who by some accident of fate has never encountered a hamburger. The interrogators hand him a story involving hamburgers, and ask him questions, which s/he answers to the best of his/her abilities. When asked about properties of hamburgers that cannot be inferred from the story, s/he claims ignorance or makes a guess. Obviously, it would not be reasonable for us to claim that the Chinese speaker does not understand Chinese simply because s/he does not know the properties of hamburgers. If anything, we would say that s/he understands Chinese, but not the world of American diners.

Similarly, the fact that the machine does not understand what a formal symbol's relationship to the world does not necessarily imply that it should be labeled "unintelligent." Rather, the design limitation of the machine, its different embodiment and experience make it differ from a human in ways irrelevant to the question at hand, just as a Chinese person who has not been exposed to hamburgers differs from a American speaker of Chinese.

Of course, we could claim that knowledge of the real world is essential to human intelligence, and that anyone who exhibits verbal behavior without accompanying knowledge does not qualify as intelligence. However, such an assertion is controversial, and can hardly be held to form a central part of our common sense intuition. For example, we usually consider congenitally blind individuals to be just as intelligent as the rest of us, even though they are deprived from a certain type of knowledge of the real world. Their inability to relate visual-based concepts to the real world is an accident, and does not bear on their intrinsic intelligence. If we accept that blind (or deaf) individuals are intelligent, the question becomes, how much real world deprivation can an entity handle while still be considered intelligent. Would we be willing to set some arbitrary threshold, for example, such that blind people are intelligent, but deaf and blind people are not, or that deaf and blind people are intelligent, but deaf and blind people with no haptic sensation are not? While imagining this gruesome scenarios is difficult, it would help us understand Searle's objection.

Our intuition regarding the intelligence of individuals who lack any non-verbal stimulation is far from obvious. For example, what if a subject of a profoundly unethical cognitive science were raised confined to a bed, blindfolded and fed through an IV, but verbally taught to do mathematics? The question whether such a person is intelligent is difficult to answer, but the intuition is not as clear-cut as Searle would like us to believe.

2.4 The argument of generative incapacity

In addition to the symbol-grounding problem, Searle's thought experiment raises another issue: to what extent are the inner workings of the computing mechanism relevant to intelligence? The intuition that the Chinese room lacks intelligence is partially due to the absence of data compression and generalization in the symbol manipulation process.

Let us say that two Chinese impostors that differ only in the type of instruction manual they have committed to memory. The first impostor's manual lists all stories in Chinese of finite length, and all questions about them, in a giant lookup table. The second impostor on the other hand has memorized a much leaner manual, which instructs the person to analyze the questions and build answers in a combinatorial fashion. While we may be reluctant to say that either person understands what Chinese words mean, it is clear that the latter understands something about how Chinese works, which the former does not. Thus, our intuitions about the Chinese Room experiment also depend on the way in which - we are told - information is represented and accessed.

3 The Aunt Bertha thought experiment

The concern is legitimate from the point of view of AI. While different people might have different intuitions regarding the contribution of real-world knowledge to intelligence, we believe that most AI researchers would find a look-up table approach to question-answering unintelligent. This intuition is best clarified by Ned Block in his Aunt Bertha argument.

Imagine that the length of the Turing test is known to us in advance, e.g. one hour. Now imagine that Block have a

machine with extremely large storage capacity, and programs it to converse by looking up the answer to any question in a giant look-up table. This is possible, Block claims, because the number of questions that can be asked in a 1-hour Turing test is finite, and of finite length. He will construct the table by consulting an actual human—Aunt Bertha on all possible conversations of some length l . Obviously, the performance of the machine on the test would not constitute a proof of its intelligence—it would merely be a testimony to Aunt Bertha's intelligence. Hence, Block argues, passing the Turing test cannot be thought of as a sufficient condition for intelligence.

To use a different metaphor: one wouldn't want to administer the Turing test to a walkie-talkie, which is remotely tuned in to Aunt Bertha. Obviously, while the answers coming from the walkie-talkie are intelligent, it is not. Essentially, a machine that recorded the answers of Aunt Bertha is merely a mechanism for transmitting Aunt Bertha's intelligence, and does not itself possess intelligence.

What is missing in both cases is **information compression and generalization on the part of the device whose intelligence we are probing**. The Aunt Bertha machine can only respond to the questions for which it was programmed, and the answers to related questions are related only because they were so in the mind of Aunt Bertha. Despite this unintelligent organization of information however, the Aunt Bertha machine is claimed to be capable of passing the Turing test.

Thus, one option is to amend Turing's definition of intelligence as follows:

If an agent has the capacity to produce a sensible sequence of verbal responses to a sequence of verbal stimuli, whatever they may be, and without requiring storage exponential in the length of the sequence, then it is intelligent [Shieber, 2006].

The problem with the revised definition is that it is no longer purely behavioral, because it requires us to examine the internal workings of the candidate entity. Therefore, Block argues, the Turing test is not a sufficient condition of intelligence.

4 Why a look-up table cannot pass the Turing test

We devote the rest of this paper to arguing against the assertion that the Turing test can be passed by a look-up table. There are, in fact, two interpretations of this argument. One interpretation is that a look-up table can be constructed such that it would be guaranteed to satisfy the Turing test. This formulation ignores the complication that no entity, even if it is intelligent, is guaranteed to pass the Turing test, due to the natural fallibility of judges. The second interpretation is that, for any look-up table, regardless of its sophistication, there is always a (very small) probability of a false positive result, if the test questions just happen to be the ones recorded in the table.

4.1 The possibility-of-construction fallacy

We begin by attacking the first interpretation of this argument, namely, that it is possible to construct a look-up table which

can definitely pass a non-trivial Turing test. By non-trivial we mean a test which is sufficiently long to allow the judge to conclude that a human possesses human intelligence.

To clarify, let us examine the notion of test length, and its influence on the argument. It is obvious that the shorter the test is, the easier it is for a machine to pass. In fact, if the test is sufficiently short, it will be passed by any machine. Suppose the test is as short as one second. No human would be able to say anything in one second, and neither would the machine. Hence, the judge would be at chance on a forced choice. Obviously, this type of failure of the Turing test is not a reason for eliminating it as a sufficient condition for the presence of intelligence. We tacitly assume that the Turing test has to be administered for a reasonable period of time.

This is the first step toward exposing the possibility-of-construction fallacy. We will show that Ned Block's argument's relies on the unwarranted assumption that real-world time and space limitations are not a factor in the Turing test. Given that we accept - and we certainly have to - that the Turing test is only meaningful beyond some minimal length, it becomes an important question whether an appropriate look-up table can be constructed to pass it.

Let us review Ned Block's proposed way of constructing the Aunt Bertha machine. He suggests to exhaustively conduct with Aunt Bertha all conversations of length one hour. Presumably, Aunt Bertha would devote her lifetime to this process. But even if Aunt Bertha lives extraordinarily long, this is impossible. Suppose that Block somehow manages to record not only Aunt Bertha's one hour conversations, but all hour long conversations that took place since humans got the ability to speak. It is clear that even in this case, the look-up table would not contain all possible one hour conversations. This is because a) the set of possible conversations depends on the natural, social and cultural environment and evolves with it and b) because future conversations can always reference those conversations that have previously occurred. For example, while a conversation like:

-Have you heard the new Federline rap song?
-Yes, I have it on my iPod.

is fairly common nowadays, it would have been impossible just five years ago. Similarly, a conversation about Plato's dialogues would have been impossible when Plato was five years old. Thus, while it is true that the set of all possible conversations of fixed length is finite at any given point in time, it is not true that it is the same set. Crucially, the set of all hourlong conversations would change from the time when Aunt Bertha's recordings would have ended to the time when the Turing test would begin.

In fairness, Block does anticipate this counter argument [Block, 1981], but dismisses it on the grounds that the Turing test is not a test of knowledge, but of intelligence, and therefore ignorance of current events does not constitute grounds for failing:

A system can be intelligent, yet have no knowledge of current events. Likewise, a machine can imitate intelligence without imitating knowledge of current events. The programmers could, if they liked, choose to simulate an intelligent Robinson Crusoe

who knows nothing of the last twenty-five years.

However, Block's reply to this challenge is inadequate. While an intelligent system need not have knowledge of current events, it should be capable of learning about them and subsequently commenting on them in the context of the test. A machine which relies on a finite look-up table will not be able to accomplish this, because it is unable to add previously non-existing entries.

The same argument holds even if the length of the test is not measured in time but in number of questions. Suppose that Block knows the test is going to contain a single question. Now he is faced with uncertainty about the length of the question, which he has to assume. Suppose he assumes that the longest question the judge can ask is the length of his/her lifetime. Unfortunately, since the judge is still alive, Block would have to estimate the length of the judge's lifetime from the length of the longest living human to date. Let us grant Block some way of generating all possible questions of that length (for example, by generating all possible strings of that length). Now he needs to provide human-like answers to them for his look-up table. Thus let us also assume that many human beings from this time on would spend their entire lives reading one question of length one lifespan and providing the answer. Even then, the machine is not guaranteed to pass the Turing test. This is because, while we know that the judge has a finite lifespan, his lifespan is unbounded. For example, there is no a priori reason to believe that the judge will not live as long as the longest living human today plus one day. If he does, Block's machine will fail the test.

Finally, let us assume that there is some maximal length of a question, which greatly surpasses a human lifespan, e.g. a question can be at most two hundred years long. If Block could create a table with all questions and answers of this length, he would surely win! Not so fast: one should not forget that each of the questions of the look-up table should be given an answer by an actual human being. Since no human being can answer a question which exceeds its own particular lifespan, it follows that the length of any question in Block's table cannot exceed the life lifespan of the longest-living human to date. Therefore, it is impossible for Block to construct his look-up table.

What these arguments go to show is that Block's assertion that a look-up table machine is guaranteed to pass a Turing test of a given length is based on two improper assumptions: 1) that the set of possible finite conversations of given length is constant in time, and 2) that the length of the questions is bounded in advance. Thus, his argument:

If the Turing test is sufficient condition for intelligence, all entities that are not intelligent should fail the test.

Premise 1: A look-up table machine is not intelligent.

Premise 2: A look-up table machine can pass the Turing test.

Therefore, the Turing test is not a sufficient condition for intelligence.

is flawed due to the fallacy of the second premise for any Turing test of unbounded length.

If the length of the test must be limited in advance, the question becomes what constitutes a reasonable length (see [Shieber, 2006]). Of course, if we severely limit the length of the Turing test, or if the length of questions is known a priori, it might be possible to construct a look-up table with guaranteed success. However, imposing such a limitation is unwarranted, as it would trivially invalidate the test. In fact, if the limit is severe enough, the Turing test would be passed by a rock! It is worth saying that the length of the test factors into validity of the first premise as well, since intelligence is a feature of a system embedded in space and time. An abstract mathematical construct manifested by a cosmic accident is difficult to relate to the common sense behind the notion of intelligence.

The limitations of a thought experiment

It is possible to defend the second premise of Block's argument by appealing to the idea that a thought experiment is not subject to the ordinary space-time limitations of the natural world. Thus, what if we imagine that Block can record not only all past, but also all future conversations of length one hour? Such a question is ill posed and causes the debate to rapidly depart the realm of usefulness. We are interested in whether or not the Turing test is a sufficient condition for intelligence in this universe, which has the limitation that time only goes forward. Indeed, it is completely unclear what the concept of intelligence would be in a universe in which we had access to the future as well as the past. If intelligence is the capacity to generate novel behavior from past observations, it might not exist as a concept in a universe in which novelty would be as absent from the future, as it is from the past.

Our point is that a thought experiment is like any other model of reality. It is an abstraction which preserves some characteristics of reality and gets rid of those characteristics which are irrelevant for the question at hand. For example, a papier-mache model of the solar system is useful if we are interested in exploring the relative sizes of planets, but is utterly useless if we are interested in exploring their gravitational fields. Similarly, a thought experiment which is meant to aid us in the question of whether the Turing test is a sufficient condition for intelligence is only useful if it preserves the kind of real-world characteristics which are relevant to intelligence, namely, space-time limitations.

4.2 The very small probability argument

The second interpretation of Block's argument is that for any look-up table, there is always a probability of a false positive result. In other words, we do not have to record all possible one hour conversations. It is enough for us to record one such conversation. There is always a chance that the judge will lead our machine into this exact conversation, and falsely conclude on the basis of the machine's verbal behavior, that the machine possesses human intelligence.

This interpretation is in some ways trivial, because the Turing test has always had a probabilistic element to it. Obviously, whenever a forced choice situation is set up among two alternatives, the response might simply reflect the judge's

prior. Turing discussed the necessity to set up that prior properly, so that positive answers are not inherently favored:

We had better suppose that each jury has to judge quite a number of times, and that sometimes they really are dealing with a man and not a machine. That will prevent them from saying "It must be a machine" every time without proper consideration (Newman, Turing, Jefferson and Braithwaite 1952).

However, it is true that this occasional fallibility of the test makes it an insufficient condition for intelligence. Furthermore, what is interesting about this particular possibility of a false positive, is that it is not in any way traceable to the fallibility of the judge. That is to say, the judge has no reason to believe that he has picked the only set of questions which the machine has on record. Hence, the judge's conclusion would be entirely reasonable on the basis of the machine's performance

It is worth thinking about whether there is a way to guarantee that the Turing test cannot be passed by a look up table (again, leaving aside the issue of judge fallibility). We suggest that the answer is yes, provided that the test is developed with reference to the moment in which the machine is considered complete. In a sense, we are reversing Block's approach. In his thought experiment, he first fixed the length of the test, and then went about constructing his machine. Here, we argue that there is a way to prepare the test after the machine is complete such that the look-up table can be guaranteed to fail.

The argument hinges on the assumption that we know what properties a look-up table constructed at time t can and cannot have. For example, the longest question such a table can contain must be shorter than t . Thus, all it takes for us to make sure that the machine is not merely using a look-up table is to ask a question which is longer than t . Obviously, this strategy could be very impractical for the reasons discussed earlier. If the time it took to construct the machine was greater than one lifetime, we might not be able to use a question longer than a lifetime in the test period (unless we have an exceptionally dedicated and longlived judge). Also, we might be unable to ascertain the time it took to construct the machine. In these cases however, it is possible to use other consequences of the observation that the set of all possible questions is dynamically evolving. For example, it is possible to ask the machine questions about events which happened after the machine was completed. Note that the machine need not observe the events—it is sufficient that the events are described to it by the judge. A human-like intelligent entity would be able to answer questions about such stories, but not a look-up table machine.

It appears that it is at least theoretically possible to make sure that any Aunt Bertha device will fail the Turing test. The question of how this could be done in reality is interesting and merits further consideration. However, it is important to bear in mind that the Turing test in practice is like any scientific measurement, in that it allows for experimental error. Therefore, it is impossible to argue that passing any particular version of the Turing test is a sufficient condition of intel-

ligence. Rather, it is the capacity to pass an ideal error-free Turing test, which is still a candidate for a sufficient condition on intelligence.

5 Conclusion

The idea that the presence of genuine thought can be conclusively determined by the ability of an entity to engage in verbal discussions can be traced back at least to Descartes *Discourse on Method*, who writes:

if there was a machine shaped like our bodies which imitated our actions as much as is morally possible, we would always have two very certain ways of recognizing that they were not, for all their resemblance, true human beings. The first of these is that they would never be able to use words or other signs to make words as we do to declare our thoughts to others. ... [O]ne cannot imagine a machine that arranges words in various ways to reply to the sense of everything said in its presence, as the most stupid human beings are capable of doing.

Nearly four centuries later, it is difficult to underestimate the foundational importance of the Turing test to the fields of Artificial Intelligence and Cognitive Science. However, the extent to which a test on verbal behavior can be taken to represent a sufficient condition for intelligence is highly controversial.

This paper assessed the validity of existing arguments against the Turing test. We presented two of the most persistent objections: the objection from information content and the objection from generative incapacity. While we agree that the Turing test has little to say about the extent to which the formal symbols manipulated by the machine relate to the real world, we showed that our intuition about the importance of this fact are not clear-cut. On the issue of generative incapacity, we argued that a machine with a finite information store could not be constructed in such a way as to be guaranteed to pass the Turing test. We also suggested that a Turing test can be constructed in such a way as to guarantee that a machine with a finite information store will fail it.

References

- [Block, 1980] N. Block. What intuitions about homunculi don't show. *Behavioral and Brain Sciences*, 3:425–426, 1980.
- [Block, 1981] Ned Block. Psychologism and behaviorism. *Philosophical Review*, 90:543, 1981.
- [Boden, 1988] M. Boden. *Computer Models of the Mind*. Cambridge: Cambridge University Press, 1988.
- [Chrisley, 1995] Ronald L. Chrisley. Weak strong ai: An elaboration of the english reply to the chinese room. In *New Directions in Cognitive Science: Proceedings of the International Symposium, Saariselk, 4-9 August 1995, Lapland, Finland, Finnish Artificial Intelligence Society: Helsinki*, 1995.
- [Copeland, 2004] B.J. (ed.) Copeland. *The Essential Turing*. Oxford University Press, 2004.

[Rey, 1986] G. Rey. What's really going on in searle's chinese room. *Philosophical Studies*, 50:169–185, 1986.

[Searle, 1980] John Searle. Minds, brains and programs. *Behavioral and Brain Sciences*, 3:417–457, 1980.

[Shieber, 2006] Stuart M. Shieber. The Turing test as interactive proof. *Noûs*, 2006.

[Turing *et al.*, 1952] Alan M. Turing, R. Braithwaite, G. Jefferson, and M. Newman. Can automatic calculating machines be said to think?, 1952. A BBC radio broadcast. Reproduced in Copeland 2004.

[Turing, 1950] Alan M. Turing. Computing machinery and intelligence. *Mind*, 59:433–460, 1950.