Detecting Intelligence:
*The Turing Test and Other Design Detection Methodologies**

George D. Montañez[1]
[1] *Machine Learning Department, Carnegie Mellon University, Pittsburgh PA, USA*
gmontane@cs.cmu.edu

Keywords: Turing Test, Design Detection, Intelligent Agents

Abstract: "Can machines think?" When faced with this "meaningless" question, Alan Turing suggested we ask a different, more precise question: can a machine reliably fool a human interviewer into believing the machine is human? To answer this question, Turing outlined what came to be known as the Turing Test for artificial intelligence, namely, an imitation game where machines and humans interacted from remote locations and human judges had to distinguish between the human and machine participants. According to the test, machines that consistently fool human judges are to be viewed as *intelligent*. While popular culture champions the Turing Test as a scientific procedure for detecting artificial intelligence, doing so raises significant issues. First, a simple argument establishes the equivalence of the Turing Test to intelligent design methodology in several fundamental respects. Constructed with similar goals, shared assumptions and identical observational models, both projects attempt to detect intelligent agents through the examination of generated artifacts of uncertain origin. Second, if the Turing Test rests on scientifically defensible assumptions then design inferences become possible and cannot, in general, be wholly unscientific. Third, if passing the Turing Test reliably indicates intelligence, this implies the likely existence of a designing intelligence in nature.

1 THE IMITATION GAME

In his seminal paper on artificial intelligence (Turing, 1950), Alan M. Turing, the brilliant, persecuted, and posthumously celebrated hero of Bletchley Park (Copeland, 2012) asked a seemingly simple question: "Can machines think?" His answer was that while this question was too vague to admit an exact answer, it could be replaced by a precise and more meaningful question: can a machine trick a human judge into believing it is human? To answer the second question, he described a variant of a popular party game, the imitation game. In his version of the game, two subjects (one human and the other a computer) are placed in separate rooms. Questions are asked of the subjects, and responses are typed out on sheets of paper (or displayed on screens). The machine's "goal" is to fool the judge into believing it is the human, and if it can reliably do so, we posit that the machine is intelligent. This game and variants of it are now known as the Turing Test (French, 2000).

While challenges to the Turing Test abound (Saygin et al., 2003), this paper presents a novel critique of the Turing Test in the spirit of a *reductio ad absurdum*. I will argue that the Turing Test procedure cannot be viewed as scientifically valid by most consistent scientists and that even its necessary assumptions have problematic consequences. This is shown in three steps. First, the core logic and assumptions of the Turing Test are drawn out. Second, it is shown that if the Turing Test is a valid and reliable indicator of intelligence, then this implies a designing intelligence behind physical nature with high probability. Lastly, even if we discard the Turing Test itself as unreliable, the assumptions necessary for any test like it to function allow for intelligent design to be formulated as scientifically sound methodology. These consequences are explored and we are either forced to reject the test and its underlying assumptions, or reevaluate the scientific standing of a theory the mainstream academic consensus has declared pseudo-scientific.

Before we begin, let us return to Turing's famous paper. The question "Can machines think?" raises a logically prior question, namely, could we even in principle distinguish between machines that think and those that do not? Without an affirmative answer to the second question, the first question remains unan-

*This is a slightly modified version from the original presented at ICAART 2016. It includes typo fixes and an updated reference to the work of Warwick and Shah, 2015.

swerable. Turing proposed the imitation game as a method of distinguishing between thinking and non-thinking machines, between intelligent and unintelligent causes. For the Turing Test to have any value, there must be some empirical basis for distinguishing between intelligent and unintelligent machines. To this we now turn.

2 SIGNS OF INTELLIGENCE

The Turing Test consists of three core components:

1. A human judge (or judges),
2. An unknown message sender (machine or human), and
3. A set of artifacts produced by the sender (e.g., words on a screen).

The job of the judge is to decide, based on observation of the artifacts, whether the sender is a human or machine. Other telltale signs and clues are removed, such as vocal articulation, physical appearance and differences in how likely the sender is to be human or non-human. This leaves the artifacts as the only observable evidence available to discern the nature of the sender. In the language of graph theory, the chain from sender to judge forms a Markov chain,

$$S \to W \to \hat{S},$$

where S is the true sender, W are the words that appear on the screen and $\hat{S}$, a function of W, is the estimate from the judge as to whether the sender is a computer or a human. The test forms a Markov chain by design, since Turing's choices of placing the sender and receiver in different rooms and not allowing vocalized communication were purposefully made to remove all side information. Because of this Markovian structure, the data processing inequality (Cover and Thomas, 2006) ensures that $\hat{S}$ contains no more information about S than given by W.[2]

The set of artifacts, W, comprising the sum total of information regarding S available to the judge, must either contain information that allows for reliable discrimination between intelligent and unintelligent senders or not. If it does not, then the Turing Test cannot work even in principle, since there can exist no reliable means of distinguishing between dumb causes (machines without intelligence) and smart ones (agents with intelligence). If no discrimination is possible between artifacts sent by unintelligent causes and intelligent causes, then one cannot distinguish between unintelligent computers and intelligent computers, which are the cause of the artifacts in question.

[2]Formally, $I(S;\hat{S}) \leq I(S;W)$, where $I(\cdot;\cdot)$ denotes the mutual information between two random quantities.

On the other hand, if artifacts *can* convey information sufficient to allow for discernment of intelligent causes, then the Turing Test might possibly work, since it opens the possibility of determining whether the cause of an artifact was intelligent or unintelligent. The test was proposed to be exactly such a discriminator: when a judge repeatedly believes the machine-made artifacts come from a human, this is offered as an indicator of computer intelligence (i.e., an intelligent cause). Thus, the purpose of the Turing Test as a test for intelligence rests on the assumptions that there is a difference between the powers of unintelligent systems and intelligent agents and that such differences can be reliably reflected in physical objects. Furthermore, reliable discrimination must be possible based on physical artifacts alone.

The logic underlying the Turing Test is laid out as follows:

1. Intelligent causes, which we call agents, can be distinguished from unintelligent causes based on their effects.
   (a) Intelligent computers can be distinguished from unintelligent computers based on the artifacts they produce.
2. Minds cannot be directly observed, while physical objects and artifacts can.
3. Physical artifacts by themselves are sufficient for discrimination between intelligent agents and unintelligent causes.
   (a) Based on the observation of physical artifacts, a procedure can distinguish between intelligent and unintelligent computers.
4. The Turing Test is proposed as a reliable procedure for performing such discrimination.

Denial of the any of the points (except 2) renders the Turing Test unworkable and invalid. Furthermore, if some procedure allowed us to perceive minds directly then the Turing Test would no longer be needed. Point 3 may be stated more strongly as "Physical artifacts by themselves, *without any knowledge of the true sender and her attributes or capabilities*, are sufficient for discrimination between intelligent and unintelligent causes," given that the test depends essentially on the judge *not* having side information about the sender. This becomes important when discussing the connection to intelligent design.

3 INTELLIGENT DESIGN CRITICISMS OF STRONG AI

Intelligent design advocates have publicly voiced skepticism to the prospects of Strong Artificial Intelligence and computational persons (Kurzweil et al., 2002; Egnor, 2014; Richards, 2011; Larson, 2015; Torley, 2014; Smith, 2015). Humans, we are told, are unique in their intellectual abilities and mind cannot be reduced to computation; thus, mere computational devices can never become persons nor ever become truly intelligent beings. The Turing Test, insofar as a machine could pass it, would not work as advertised, since a machine could not be intelligent in principle. Thus, skepticism towards the prospect of Strong AI has been openly voiced within the intelligent design community.

However, we will see that this perceived incompatibility actually obscures a striking similarity between the Turing Test and the intelligent design research program.

4 UNEXPECTED IMPLICATIONS

Given that the Turing Test is rooted in the assumed capability of physical structures to carry observable signs of intelligence, we come to an unexpected connection. Intelligent design theory posits that features of the natural world, such as biological organisms, bear the hallmarks of being caused by an intelligent agent and that by observation of physical structures one can reliably discriminate between designed and undesigned objects. If the Turing Test is a scientifically and philosophically sound test then we cannot rule out that this test might be applied to artifacts other than words. Indeed, it takes special pleading to prevent the test from being applied to other kinds of objects, such as poems, musical scores, works of architecture and molecular machines. Given the little contested fact that humanity has a strong tendency to attribute design to structures in nature, this would further indicate that an intelligence behind nature has already been judged as present: when judges repeatedly attribute artifacts to an intelligent mind, the Turing Test says the cause behind the artifacts is indeed some form of intelligence.

While one could personify natural selection as an intelligent agent, we would also need to account for cosmological structures, which cannot as easily be ascribed to the powers of Darwinian natural selection (though some have argued that this is also the case, e.g., (Smolin, 2004)). However, invoking natural selection as an intelligent agent acting across all space and time (wise, omnipresent and eternal) brings us remarkably close to the Charybdis of classical theism. Claiming natural selection as a god-like intelligent mind seems a poor way of avoiding the idea that natural systems are the result of an intelligent mind. Doing so essentially concedes the point at issue, even if arguing against its interpretation.

A more sober rejoinder is that just because humans repeatedly mistake artifacts as being the result of an intelligent mind does not mean they are. This is true; but then what becomes of the Turing Test itself? The test is based on the premise that if humans repeatedly mistake artifacts as being the result of an intelligent agent then they likely are. However, if humans consistently judge the cause of observed artifacts to be intelligent, and yet this does not force us to consider the cause intelligent, then a successfully passed Turing Test becomes meaningless: fooled human judges would serve as no indicator of prior intelligent activity. Thus, an inherent tension exists between the logic of the Turing Test and unexpected consequences that follow from its validity.

Taking a step back, we might argue that while the applied logic of a generalized Turing Test is not suitable for discriminating between intelligent and unintelligent causes, other methods may be. While this avoids the problem of having already confirmed the likely existence of design in nature, it leaves open the possibility of design inference in general. If physical objects can reliably transmit signals of a designing intelligence, even absent of knowledge concerning their origins or the additional capabilities and attributes of their true causes, then research into methods of empirical design detection cannot be ruled out a priori. Furthermore, such methods might one day be applied to artifacts of nature, such as biological systems and cosmological structures. If reliable methods of design detection are developed, then we cannot rule out the possibility that natural objects themselves might contain such signs of intelligence. If they did, intelligent design theory would not only give a sound scientific methodology but also a true theory of origins. If such possibilities cannot be entertained then neither the methodology nor the logic underlying the Turing Test can be treated as sound.

We now see that the Turing Test and intelligent design methods attack different variants of the same underlying problem: detecting intelligent causes through the examination of generated artifacts. While there are differences between the two methods, the similarities are striking. A solution to one problem set provides a solution to the other, since reliable methods of design detection can be applied to words written on a computer screen as easily as to engineered objects like machines. In both cases, uncertainty surrounds the true sender and his or her capabilities. In both cases, physical objects are examined to determine whether the generative cause was an intelligent agent or merely an unintelligent physical system. Fundamentally, the Turing Test is applied intelligent design: a special case of design detection applied to computational artifacts. Thus, the scientific validity of one endeavor is intrinsically tied to the scientific validity of the other.

5 POSSIBLE OBJECTIONS

Let us now propose some arguments that might salvage the Turing Test and its underlying assumptions.

5.1 The Language Objection

It may be argued that the textual and verbal aspects of Turing Test are inseparable from it and that generalized Turing Tests that rely on other types of artifacts such as machines, artwork, or other non-verbal compositions are not possible, or at very least, are not of equal validity to the classical form of the Turing Test. Furthermore, it can be argued that because nature is *not* a written text, then we can safely prohibit the application of Turing's logic to aspects of nature, sparing one from potentially embarrassing consequences, such as seeing design in the cosmos.

We have already briefly mentioned that this restriction has all the appearances of special pleading. If verbal artifacts in isolation can carry sufficient information to allow for agent discrimination, why can't other artifacts? Why not machine code or machine designs? We can recognize the work of genius in a superbly constructed machine or an emotionally moving portrait as readily as we can in written text. It is difficult to argue otherwise. For example, if we encountered a planet where non-verbal beings were busily constructing skyscrapers, vehicles, and sculptures, would we hesitate to ascribe intelligence to the beings simply because they could not speak? Do we normally accuse the mute of not possessing minds? Words, while possibly sufficient, do not appear to be necessary to convey the presence of an intelligent mind.

Nature is replete with awe-inspiring marvels of engineering, containing self-adapting and self-replicating code systems, all following the mathematically elegant (and surprisingly comprehensible (Wigner, 1995)) forces of physical law. However, even ignoring the apparent genius of natural machines and systems, our objection is defeated by the presence of an actual textual system in biology itself: DNA (Hood and Galas, 2003). DNA forms the basis for a chemical language, which is transcribed, read, and directs the construction of objects within the cell. The four bases, A-G-C-T, act as chemical letters, forming "words" that convey instructions and regulate processes. Thus, if strings of text are sufficient for conveying intelligence, then the language objection cannot exempt biological artifacts from the application of intelligence detection methods.

5.2 Known Origin of Biological Systems

One might also argue that since we *know* how biological systems were produced, namely through unguided evolution over millions of years, then we can independently rule out any "false positives" caused by the potential application of Turing-like tests to the natural world.

Such an argument fails; if knowing the mechanical causes of an artifact rules out the possibility that it was produced by an intelligent mind, then computers could never be judged as intelligent. We know, down to the transistor and logic-gate level, the mechanical workings of computer systems. Any artifact produced by a modern computer can be traced back to exact sequences of binary pulses and voltage oscillations in logic gates. Thus, if the objection were to hold, computers cannot in principle be ruled intelligent, since any hint of design in their artifacts would be immediately overruled by our knowledge of their mechanical origins, turning all such inferences into false positives. This objection applies to intelligent design and the Turing Test with equal force.

5.3 Objection of the Unknown Agent

It might be further argued that we only recognize design and intelligence by analogy to human capabilities and that we can recognize purported artificial intelligence only as far as it parallels human intelligence. Because artificial intelligence developed by man will be constructed with the explicit goal of mimicking man, we have some hope of being able to detect it. However, were the intelligent mind extraterrestrial, immaterial (e.g., a spirit), or otherwise wholly unknown, then we may not be able to detect design in its artifacts or infer its intelligence. Thus, one might maintain that the Turing Test can reliably detect certain types of intelligence, namely agents of man-made origin (human or computer) that are sufficiently human or human-like, but that intelligent design methods cannot reliably detect the actions of non-human intelligent agents in nature, precisely because they are *not* human-like. I will call this the objection of the unknown agent.

This is superficially plausible as an objection, since it is true that an unknown mind might not display its intelligence in ways that are humanly perceivable or detectable by our tests. However, it might; being unknown, we cannot rule out this possibility. Furthermore, if we had independent evidence that the mind in question *did* in fact generate artifacts similar to those produced by humans, then the force of the objection would vanish. Even barring such independent evidence, the degree to which the unknown actor produced artifacts similar to those produced by humans is the degree to which our tests might (correctly) detect its intelligence. Thus, in regards to intelligent design in nature, we have an empirical question: how similar are the artifacts of nature to human artifacts? The fact that springs (Shin and Tam, 2007), gears (Burrows and Sutton, 2013), compasses (Qin et al., 2015), boolean logic networks (Robinson, 2006), digital codes (Hood and Galas, 2003) and other human inventions have been found to preexist in biology at least suggests the possibility that a mind behind natural phenomena would be sufficiently similar to human intelligence to allow for its detection through observation of its engineered artifacts.

Lastly, the objection, at its core, concerns false negatives, not false positives. The objection is basically that an unknown agent may escape detection, not that an unknowing system might be mistaken as intelligent. Given the danger of falsely attributing intelligence to unintelligent systems, this is exactly the form of bias we want. Furthermore, if we did detect an unknown intelligence behind a set of objects, the detection itself would be evidence in favor of it being sufficiently similar to human intelligence, since such tests only output positive classifications when artifacts are sufficiently similar to those anticipated by humans.

Therefore, the possibility of unknown agents does not create sufficient separation between the Turing Test and other methods of inferring design to warrant any distinction.

5.4 The Interrogation Objection

Lastly, there remains a seemingly powerful argument that can be raised against the full equivalence of the Turing Test to other design detection methods: the call-and-response nature of the Turing Test seems absent from other methodologies, and thus might be used to validate one while invalidating the other. The Turing Test presupposes that one can *interrogate* the subject in question, gathering specific responses to specific questions, whereas one cannot demand specific answers from nature.

To ensure this argument does not point to a distinction without a difference, we must specify exactly what it is about the call-and-response that makes it uniquely suitable to convey the existence of intelligence, and ensure that this same quality cannot be attained by other means. One such feature could be that in responding to posed questions, a system is forced to respond to and overcome unanticipated or novel challenges. However, the same can be said of biological systems, which have shown adaptive architectures for overcoming a variety of fluctuating and novel environmental challenges. When viewed from this perspective, we have systems which are given input configurations (e.g., questions from a judge or environmental structure in nature) and create response configurations that convey intelligence (e.g., an apropos response or clever environmental adaptation). Furthermore, nature can be probed for specific answers in the form of experiments and tested hypotheses, thus mimicking, albeit imperfectly, the call-and-response structure of the Turing Test. Thus, this distinction might not present any real difference.

Examining this critique in greater depth, we see that even given the call-and-response structure of the Turing Test, one cannot fully control the responses given, thus lessening the force of this objection. While an interrogator is free to ask the system in a Turing Test any question, he cannot be guaranteed that the response will be sufficient or even relevant. For example, consider the responses of Eugene Goostman, a computer system claimed to have passed the Turing Test by fooling a subset of human judges into thinking it was human (Warwick and Shah, 2015). The following excerpts were taken from transcripts of conversations held with an online version of Goostman (Auerbach, 2014):

Judge: Why do birds suddenly appear?

Eugene: Just because 2 plus 2 is 5! By the way, what's your occupation? I mean - could you tell me about your work?

...

Judge: It is ok, I get sick of sick people. How is your stomach feeling today? Is it upset maybe?

Eugene: I think you can't smile at all. I bet you work in a funeral agency.

Clearly, the output of such systems may have little or no correlation to the input questions. When the responses are unrelated to the questions asked, we are essentially given an arbitrary collection of sentences, having no say in their contents, similar to the situation in nature. However, even such a rigid collection might persuade judges, if it displayed empirical signs of intelligent thinking. A one-way communication channel, such as a mayday signal or SETI-received broadcast from an alien civilization, may still contain sufficient information to convey an intelligent agent at work. It is difficult to argue that *only* two-way communication can convey intelligence when such counterexamples are easily found.

However, if the interrogative nature of the Turing Test is a necessary aspect for intelligence detection and there exist no other means of attaining the same property, then we would be forced to demote full equivalence to a mere uncanny similarity.

6 STRUCTURED ARGUMENTS

Having stated the main arguments and some possible objections, we will now review the arguments in a more structured manner. Let P denote premises, C conclusions and R corollaries.

1. The first argument can be stated as follows:

P1. The Turing Test procedure is based solely on information contained in observable artifacts.

P2. That such artifacts could reliably indicate the causal action of an intelligent agent entails that:

A. The causal powers of intelligent agents differ from those of unintelligent systems in empirically detectable ways; and

B. Signals of these differing causal powers can be reliably contained in physical objects, absent of knowledge of their origins or characteristics of their generative systems.

P3. Nature contains physical objects of unknown origin.

C1. From (P1)-(P3), it is possible that natural physical objects may reliably indicate the presence of an intelligent cause in empirically detectable ways.

R1. Reliable signals of an intelligent cause may exist in natural objects such as biological systems or cosmological structures.

2. The second argument can be stated as:

P4. The Turing Test is a reliable test for detecting the activity of intelligent agents.

P5. The Turing Test is scientifically legitimate.

P6. The primary objective of intelligent design theory is to develop reliable tests for detecting the activity of intelligent agents.

C2. (P4) and (P5) together entail that developing reliable tests for detecting the activity of intelligent agents can form a scientifically legitimate pursuit.

C3. From (P6) and (C2), the primary objective of intelligent design theory might represent a scientifically legitimate pursuit.

R2. From (P4) and (P5), there exists at least one reliable, scientifically legitimate test for detecting the activity of intelligent agents.

R3. Design inferences in general cannot be unscientific.

3. The third argument can be stated as:

P7. Successfully passing the Turing Test indicates the probable existence of an intelligent cause.

P8. Passing the Turing Test requires that physical artifacts be repeatedly mistaken for the products of an intelligent cause.

P9. Natural cosmological and biological objects have repeatedly been mistaken for the products of an intelligent cause.

C4. From (P8) and (P9), the cause of natural cosmological and biological objects has successfully passed the Turing Test.

C5. From (P7) and (C4), the cause of natural cosmological and biological objects is probably intelligent.

7 CONCLUSION

Where does this leave us? For the Turing Test to work, one must be able to distinguish intelligent causes from unintelligent causes based solely on observable artifacts. But this leads to the conclusion that intelligent design cannot be simultaneously disregarded, since its methodological structure rests on the same foundation. Furthermore, if the Turing Test is a reliable procedure for detecting intelligence, then the cause of biological origins is likely an intelligent mind, having passed a generalized Turing Test since the time of Cicero (Cicero, 45BC). Appealing to strong materialism, known mechanistic origins and unknown agents does nothing to remedy the situation. This forces us to either reject the reliability of the Turing Test and its underlying logic, or face what many would consider a creeping creationism.

To allow for the possibility of intelligent design would be to deny a scientific consensus, and we are often reminded that true scientists and intelligent laymen cannot deny any fact established by the consensus of experts. To do so is derided as a strong form of denialism, for suggesting the majority of scientists might be wrong on a well-studied issue. Surely such could not be the case (Ioannidis, 2005; Baker, 2015).

REFERENCES

Auerbach, D. (2014). A computer program finally passed the Turing Test? slate.com/articles/technology/bitwise/2014/06/turing_test_reading_university_did_eugene_goostman_finally_make_the_grade.html/. Accessed: 2015-11-10.

Baker, M. (2015). First results from psychology's largest reproducibility test. http://bit.ly/1bk6qDF. Accessed: 2015-05-13, nature.com/news.

Burrows, M. and Sutton, G. (2013). Interacting gears synchronize propulsive leg movements in a jumping insect. *Science*, 341(6151):1254–1256.

Cicero, M. T. (45BC). *Of the Nature of the Gods*. Book II, chapters XXXVII, XLIV, and XLVII.

Copeland, B. J. (2012). *Turing: pioneer of the information age*. Oxford University Press.

Cover, T. M. and Thomas, J. A. (2006). *Elements of information theory*. Wiley & Sons. Section 2.8.

Egnor, M. (2014). Steven Novella Doesn't Trust His Computer One Bit. evolutionnews.org/2014/12/steven_novella092181.html. Accessed: 2015-05-11.

French, R. M. (2000). The Turing Test: the first 50 years. *Trends in cognitive sciences*, 4(3):115–122.

Hood, L. and Galas, D. (2003). The digital code of DNA. *Nature*, 421(6921):444–448.

Ioannidis, J. P. (2005). Why most published research findings are false. *PLoS medicine*, 2(8):e124.

Kurzweil, R., Gilder, G. F., and Richards, J. W. (2002). *Are We Spiritual Machines?: Ray Kurzweil vs. the Critics of Strong AI*. Discovery Institute.

Larson, E. J. (2015). Group Delusions Aside, Sentient Robots Aren't on the Way. evolutionnews.org/2015/03/sorry_nature_gr094741.html. Accessed: 2015-05-11.

Qin, S., Yin, H., Yang, C., Dou, Y., Liu, Z., Zhang, P., Yu, H., Huang, Y., Feng, J., Hao, J., Hao, J., Deng, L., Yan, X., Dong, X., Zhao, Z., Jiang, T., Wang, H.-W., Luo, S.-J., and Xie, C. (2015). A magnetic protein biocompass. *Nat Mater*, advance online publication. Article.

Richards, J. (2011). I, for one, welcome our new robot overlords. aei.org/publication/i-for-one-welcome-our-new-robot-overlords/. Accessed: 2015-05-19.

Robinson, R. (2006). Mutations change the boolean logic of gene regulation. *PLoS biology*, 4(4).

Saygin, A. P., Cicekli, I., and Akman, V. (2003). Turing test: 50 years later. In *The Turing Test*, pages 23–78. Springer.

Shin, J. H. and Tam, B. (2007). Unleashing of a biological spring. In Magjarevic, R. and Nagel, J., editors, *World Congress on Medical Physics and Biomedical Engineering 2006*, volume 14 of *IFMBE Proceedings*, pages 2824–2827. Springer Berlin Heidelberg.

Smith, W. J. (2015). Machines Will Always Be Things, Never "Persons". evolutionnews.org/2015/04/machines_will_a095201.html. Accessed: 2015-05-11.

Smolin, L. (2004). Cosmological natural selection as the explanation for the complexity of the universe. *Physica A: Statistical Mechanics and its Applications*, 340(4):705–713.

Torley, V. (2014). Why a simulated brain is not conscious. http://www.uncommondescent.com/intelligent-design/why-a-simulated-brain-is-not-conscious/. Accessed: 2015-05-19.

Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, pages 433–460.

Warwick, K. and Shah, H. (2015). Can machines think? A report on Turing test experiments at the Royal Society. *Journal of Experimental & Theoretical Artificial Intelligence*, pages 1–19.

Wigner, E. P. (1995). The unreasonable effectiveness of mathematics in the natural sciences. In *Philosophical Reflections and Syntheses*, pages 534–549. Springer.