



The Mythical Turing Test

[Aaron Sloman](#)

This file is <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-test.html>

Also <http://tinyurl.com/tmyth>

From time to time I shall use html2ps and ps2pdf to create a PDF version, better suited for printing, though it may be out of date at times:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-test.pdf>

A compressed version of this paper will appear in one of a collection of papers on Turing's work (one of the Centenary volumes, but delayed till 2013).

A pre-print is available here: <http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106c> here

It is widely believed that Turing proposed a test for intelligence.

This is false.

He was far too intelligent to do any such thing, as should be clear to anyone who has read

his paper, online here: <http://cogprints.org/499/>

Unfortunately, many people who report what he is supposed to have proposed, have never read the original: they simply believe something reported by someone else about what Turing proposed.

Read Turing's 1950 paper

Turing actually wrote

"1. The Imitation Game

I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think." The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words "machine" and "think" are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, "Can machines think?" is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.
..."

So, far from proposing a test to answer the question whether machines can think or whether machines are intelligent, he actually decides (rightly) that the question is absurd because the question is (as he put it in section 6, quoted below), "too meaningless to deserve discussion."

Instead, after a preamble, he proposes not a test but a game, which he calls "The imitation game", and he proposes it in order to make a technological prediction, not in order to provide a test for intelligence.

What Turing thought about the question

Later in the paper he wrote:

"6. Contrary Views on the Main Question

We may now consider the ground to have been cleared and we are ready to proceed to the debate on our question, "Can machines think?" and the variant of it quoted at the end of the last section. We cannot altogether abandon the original form of the problem, for opinions will differ as to the appropriateness of the substitution and we must at least listen to what has to be said in this connexion.

It will simplify matters for the reader if I explain first my own beliefs in the matter. Consider first the more accurate form of the question. 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. The original question, "Can machines think?" I believe to be too meaningless to deserve discussion.
..."

As far as I can tell, Turing described the game not in order to give a criterion for thinking or for intelligence, but in order to be able to formulate that prediction (part of "the more accurate form of the question"). He was remarkably accurate about the number of bits likely to be available in a computer's memory by the turn of the century.

I think that by about the year 2000 it actually was possible to fool a significant subset of the human population, in an interaction lasting only a few minutes into thinking that they were talking to a human rather than a machine, though only in a very restricted range of contexts.

But it is unlikely that anyone sensible would claim that fooling 30% of the *average* population for a mere *five* minutes was a criterion for a machine to be intelligent, and Turing certainly did not. It's a pity so many highly educated, highly intelligent people support the myth.

However, even though the technology has got better, and computers are now doing much cleverer things than they were a decade ago, humans have also been learning and are now much better educated about computers, as a result of very frequent use, as well as news reports, discussions, etc., than they were then.

I suspect that has made it harder to fool some of them into thinking they are interacting with a human. In other words, although AI technology has improved, so has general human understanding of what computers can and cannot now do. As a result, many humans are more likely to be able to choose things to say to the machine that will reveal its inability to respond like a human. But that reflects limitations of current technology rather than any deep fact about what computers (or more generally machines) can or cannot do in principle.

I suspect Turing underestimated the difficulty of programming machines to be like humans even in his simple game.

Anyhow, the truth or falsity of his prediction is not the main point of his paper: he made the prediction as a basis for attacking arguments purporting to show that it would be impossible to programme a machine to play the imitation game successfully, even with a very weak criterion for success. He puts forward several such objections and then replies to them, attempting to show their errors.

Unfortunately, the misinterpretation of his paper as proposing a test for thinking or for intelligence is so wide-spread that it has led to huge amounts of wasted effort discussing the merits of such a test: wasted, because, as Turing himself pointed out, the notion of such a test is based on a question which is "too meaningless to deserve discussion".

It's a pity so many people don't bother to read what he wrote.

For a critique of some assumptions regarding the question whether machines could be conscious see

- <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#cons09>
Why the "hard" problem of consciousness is easy and the "easy" problem hard.
(And how to make progress)
A PDF tutorial presentation, also available in flash format on slideshare.net
 - <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk80>
Helping Darwin: How to Think About Evolution of Consciousness
Or "How could evolution get ghosts into machines?"
 - <http://www.cs.bham.ac.uk/research/projects/cogaff/09.html#pach>
Phenomenal and Access Consciousness and the "Hard" Problem: A View from the Designer Stance
-

Turing's error about human-like learning

Turing did make one serious error in that paper. In the section headed:

7. Learning Machines

he wrote

"In the process of trying to imitate an adult human mind we are bound to think a good deal about the process which has brought it to the state that it is in. We may notice three components.

- (a) The initial state of the mind, say at birth,
- (b) The education to which it has been subjected,
- (c) Other experience, not to be described as education, to which it has been subjected.

Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's? If this were then subjected to an appropriate course of education one would obtain the adult brain. Presumably the child brain is something like a notebook as one buys it from the stationer's. Rather little mechanism, and lots of blank sheets. (Mechanism and writing are from our point of view almost

synonymous.) Our hope is that there is so little mechanism in the child brain that something like it can be easily programmed. The amount of work in the education we can assume, as a first approximation, to be much the same as for the human child.
...."

I think that here Turing, like many AI researchers studying learning machines, grossly underestimates the contribution of biological evolution to the processes of human learning.

John McCarthy also emphasised the contribution of evolution in "The Well-Designed Child", a paper he wrote around 1996, which was eventually published in the journal *Artificial Intelligence*, vol 172, No 18, 2008, also [available on his web site](#).

"Evolution solved a different problem than that of starting a baby with no a priori assumptions.

...

Animal behavior, including human intelligence, evolved to survive and succeed in this complex, partially observable and very slightly controllable world. The main features of this world have existed for several billion years and should not have to be learned anew by each person or animal. In order to understand how a well-designed baby should work, we need to understand what the world is like at the gross level of human interaction with the environment."

In the rest of the paper McCarthy makes some useful contributions to the analysis of what might have been done by evolution for such a baby.

Jackie Chappell and I have been working on a similar hypothesis. Some of the ideas are in two papers and an online presentation

- The Altricial-Precocial Spectrum for Robots, A. Sloman, J. Chappell, *Proceedings IJCAI'05*, Edinburgh, pp. 1187--1192, 2005, <http://www.cs.bham.ac.uk/research/cogaff/05.html#200502>,
 - Natural and artificial meta-configured altricial information-processing systems, Jackie Chappell, Aaron Sloman, *International Journal of Unconventional Computing*, 3, 3, pp. 211--239, 2007. <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609>
 - Talk 52: Evolution of minds and languages. What evolved first and develops first in children: Languages for communicating, or languages for thinking (Generalised Languages: GLs) (Online Presentation PDF) <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#glang>
[A 'flash-based' version is on slideshare.net](#)
-

A point missed by many: There's no dichotomy

There is a very common mistake, implicitly made by anyone who asks "Can machines be intelligent?", or possibly even "Can machines think?" (which is more subtle in ways that I shall not discuss).

The mistake is looking for, or assuming that there exists, a **dichotomy** (a binary distinction) in a complex space where things are very varied. An example would be to assume there is a binary divide between things that are and things that are not intelligent (or conscious, or sentient, or able to have emotions, or ...)

Sometimes people who realise that that is a mistake try to avoid the error by talking about differences of **degree**, e.g. by saying that instead of a binary divide there are differences in **degrees** of intelligence, consciousness, etc.

That view makes two related mistakes

- Assuming that there is a total ordering of cases, as if, for example, species could be put into a linear ordering of animals with more or less intelligence (or consciousness, etc.)
- The second mistake that is sometimes made in referring to "degrees" is the assumption that there is **continuous** variation.
This can also include claiming that all boundaries are fuzzy.

An example of the first sort of mistake could also be made by a child who finds that some containers can be put inside other containers (e.g. a little box inside a bigger box) and draws the conclusion that all containers form a linear ordering, so that given any two containers Ca and Cb, either Ca can be put inside Cb, or Cb can be put inside Ca or they are exactly the same size, because their bounding surfaces are the same shape and size.

This assumption of a total ordering is easily shown to be wrong by the discovery of containers of different shapes such that neither can fit inside the other. E.g. one may have a square cross section and the other circular, and the diameter of the circle is more than the side of the square and less than the diagonal of the square.

There can be many different sorts of competence that are aspects of intelligence, and different individuals or different species may have different subsets. E.g. someone may be very good at designing furniture and terrible at proving theorems in geometry, while the other is good at geometry and poor at designing furniture. Such variations in ability to perceive, learn about, understand, act in, plan in and survive in various types of environment are wide spread among organisms, so they cannot be arranged in a linear sequence of increasing intelligence.

The assumption of continuous variation is also false, since there are some kinds of knowledge or competence that are present or absent, but cannot vary continuously between those extremes, e.g. knowing the sum of 2 and 2. More generally, competences that are expressed in rules don't have intermediate cases using half a rule, quarter of a rule, etc. Since biological genetic makeup is ultimately implemented in chemistry and chemical molecules can only differ discretely (e.g. by addition or removal of atoms) it is impossible for species to vary continuously in their genetic makeup.

The errors in this section are related to the mistake assumption that there is some unique competence "general intelligence" which is either thought of in binary (i.e. it is either present or absent in an individual or species) or as a matter of degree, with linear variability along a total ordering.

Similar mistakes can be made about machines: e.g. assuming that there is a binary divide, assuming that there is a total ordering, or assuming that there is continuous variability. Compare [Artificial General Intelligence \(AGI\)](#)

Conclusion: the idea of a Turing test is muddled

It follows from the above that the very idea of a Turing test or any other test for intelligence is muddled if there is no binary divide between things that are and things that are not intelligent, only a vast variety of cases.

A different muddle is introduced if the proposal for a **test** is replaced by a proposal for a way of **measuring** intelligence (or thinking ability). This can be mistaken both in assuming that there is a total ordering of types of intelligence and that there is continuous variation so that degrees of intelligence can be represented by real numbers.

If there is no total ordering, only a complex space of combinations of competences (just as there is a complex space of combinations of atoms, making the notion of ordering molecules along a line of increasing 'chemicality' (?) misguided), then, instead of using a number or label for degree or amount of intelligence, we need to find ways of **describing** types of intelligence in terms of the combinations of competences that they include, just as we describe chemical molecules in terms of the different combinations of atoms and chemical bonds and also the different properties that result from those structures (e.g. acidity, alkalinity and many more used in drug design, for example).

If varieties of intelligence vary in something like the ways in which sentence structures do, with different components combined hierarchically with varying relationships between the components, then perhaps we need something like a **grammar** for types of intelligence.^[*] However, that is just one possibility to be considered.

Compare ["The design-based approach" to the study of mind.](#)

NOTE:

The notion of "grammar" used here is not restricted to production rules for generation of strings. E.g. during the 1960s, many researchers attempted to generalise Chomsky's notion of a "generative grammar" to include any specification of a class of structures composed from some well defined set of primitives arranged in accordance with permitted relationships, to form new structures that could also be arranged in accordance with permitted relationships, and so on. If the rules are recursive then an infinite variety of structures could be characterised in this way. "Web grammars", for example, were developed to characterise 2-D graphical structures found in pictures.

From this viewpoint understanding a complex structure conforming to such a grammar amounted to constructing a specification of the structure's derivation from the rules of the grammar, just as parsing a sentence involves constructing a specification of the derivation of the sentence from the grammatical rules. However, in the case of picture grammars the derivation need not have the form of a tree: it could be a graph with cycles,

just as pictures can include cyclic structures (e.g. polygons).

A different sort of "grammar" might explicitly characterise ways of generating structures from some initial structures. For example, a grammar for legal chess-board configurations would be a specification of an initial state of the board, plus a specification of permitted changes to any possible state. In this case a derivation of a chess position would be a description of a legal game leading to that position. A quite different notion of a grammar for a chess configuration would characterise pieces and groups of pieces in terms of their positions and relationships. So notions like "pinning" and "forking" and a move being forced would be part of the characterisation of the board structure.

A grammar for chemical molecules might specify ways in which a finite set of atoms can be combined in accordance with the laws of physics to form a single structure, and could include intermediate structures that are neither atoms nor molecules, e.g. ions.

In many cases ambiguity is possible: the same structure may have more than one "parse" if there are two or more ways in which the 'rules' allow it to be constructed.

A grammar for types of intelligence might be a specification of varieties of combinations of competences of many sorts that could be implemented in a unified working architecture, and could include combinations that could grow themselves, e.g. as a result of interactions with some environment.

REF

S. Kaneff, Ed. **Picture language machines**, 1970, Academic Press, New York,

My thanks to Damien Duff for pointing out the need to explain this generalised use of the word "grammar". I have also used it in connection with the idea of [a "grammar" for emotions](#), as a counter to shallower forms of description of emotions.

ADDED 26 May 2010: What sort of test would be worth while?

If we are to propose tests of the general sort that people take the Turing Test to be, namely a test for something being intelligent, or human like, it is important to distinguish testing **a particular individual** from testing **a theory about a type of design** for working systems.

It is clear that there are many very different human beings and also that they all share a large collection of common features. We really should be testing a theory about what's common.

Compare: if a theory about the weather is able to explain only how a particular sort of tornado works and no other tornadoes, nor any other weather phenomena, then it cannot be a good explanation of the particular case. Likewise a theory of what's going on when oil burns that says nothing wood burning, or coal burning, or gas burning cannot be a good theory about oil burning.

Or suppose someone claims to have a theory about how to solve algebraic equations and an online computer test that demonstrates the theory. The implemented algorithm solves only quadratic equations: give it any quadratic equation and it will produce the solutions (including complex solutions where appropriate, as required for this equation:

$$X \times X = -1$$

Would that be taken as a good test for a theory of equation solving? We would rightly demand something more *generic*.

What would the required sort of *generic* theory of intelligence look like?

The closest answer I can give is something like a parametrised specification for a highly polymorphic design for a working system, which can be given different parameters to produce instances of the design, where the instances will be very different in the way that, for example, humans in different cultures, or who talk different languages, or who grow up to have very different competences and interests are different, and yet be as similar as different humans are, a requirement that is very complex and very demanding, and not yet specifiable in detail since we don't yet know enough about what typical humans are like (e.g. how their vision systems work, how they learn, what mechanisms are involved in their motivational and other affective states and processes.)

The parameters, instead of all being supplied at the time the instance is created, would have to be picked up at various times during the development and testing of the instance.

In particular, in order to really understand human intelligence we should be able to specify a type of system, different instances of which, as a result of educational and environmental influences, and some individual personality features, "grow up" to be philosophers with very different views, including views on what machines can or cannot do, e.g. some growing up to be robots like John Searle, others like Tom Nagel, others like David Chalmers, others like Dan Dennett, and perhaps even some like me, since I disagree with all the others!

See also:

<http://www.cs.bham.ac.uk/research/projects/cogaff/09.html#910>

An Alternative to Working on Machine Consciousness

<http://www.cs.bham.ac.uk/research/projects/cogaff/10.html#1003>

How Virtual Machinery Can Bridge the "Explanatory Gap",
In Natural and Artificial Systems. (For SAB'2010).

Useful links

So much has been written about Turing's paper that it would be pointless to try to give a comprehensive set of references. Occasionally I shall add links here if I find something that looks useful (even if I have not read it 'cover to cover').

- For more on what Turing *actually* wrote see Chapter 16 (section 16.ii) of Margaret Boden's book [Mind As Machine: A history of Cognitive Science \(Vols 1--2\)](#) Oxford University Press, 2006
- This paper does not fall into the common trap: <http://crl.ucsd.edu/~saygin/papers/MMTT.pdf> Saygin, A.P., Cicekli, I. and Akman, V. (2000), 'Turing Test: 50 Years Later', *Minds and Machines* 10(4): 463-518.
Part of this web site (the turing test page):
<http://www.fil.ion.ucl.ac.uk/~asaygin/tt/ttest.html>

- After reading this paper (in 2010) Diane Proudfoot modified a draft of the following paper D. Proudfoot, Anthropomorphism and AI: Turing's much misunderstood imitation game, in **Artificial Intelligence**, Vol 175, 5--6, 2011, pp. 950--957, <http://www.sciencedirect.com/science/article/pii/S000437021100018X>
However, she disagrees with my reading of Turing. Later I'll try to present an analysis of the disagreement. (Please let me know if you do!)
 - This symposium may be of interest: AI-Inspired Biology 31st March -1st April, 2010. <http://www.cs.bham.ac.uk/research/projects/cogaff/aiib/>
Part of [the AISB 2010 Convention](#)
-

Maintained by:

[Aaron Sloman](#)

Installed: 1 Feb 2010

Updated: 2 Feb 2010; 6 Feb 2010; 26 May 2010; 23 Dec 2010; 25 Nov 2011; 5 Dec 2012