Updating The Compact Conception of Intelligence

Thomas Woodside
Yale University
11 October 2022

ABSTRACT. Shieber (2007) responded to arguments levied against the Turing Test (Turing, 1950) with his "compact conception of intelligence." While the conception does illustrate something useful, it has several problems. First, the conception, still reliant on the Turing Test, focuses exclusively on verbal intelligence. Second, the use of "sensibility" means that the test is really one of human-like intelligence, not a more general kind of intelligence as we might hope. Third, Shieber's bar for compactness is hopelessly low so as to not be very useful. I then argue that many definitions of intelligence fall into one of two categories, but that these categories can conflict. I then present three alternative conceptions that resolve these objections and clarify the categories. Rather than speaking of intelligence, I argue, we can speak of *competence* and *compactness*. In cases where competence and compactness do *not* conflict, we can speak of a *pareto improvement in intelligence*.

1. Introduction: Shieber's compact conception

Turing (1950) proposed The Imitation Game (now called the Turing Test) as a method for assessing the intelligence of machines. The test typically consists of a human interacting with the machine via a text-only interface, and then making a judgment about whether the machine is human or not. If humans frequently incorrectly attribute humanity to the machine, then the machine is said to have passed the test. Block (1995) critiqued the Turing Test by envisioning an "Aunt Bertha" machine that could respond merely by memorizing responses to every possible query up to the length of the test and claiming that such a machine would not be intelligent. In response, Shieber (2007) introduced the following conception of intelligence, which he defined as follows:

> **The compact conception** (Shieber, 2007): 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 the agent is intelligent.

Shieber combines this conception with the notion of an interactive proof to argue that a Turing Test would only have to be 140 words long in order to unmask an Aunt Bertha machine with a storage capacity the size of all information in the universe. Thus, he argues that with trivial constraints on the physical realizability of a system, the Turing Test can provide a (probabilistic, interactive) proof of intelligence.

I think that the compact conception is a step in the right direction, and I will not dispute Shieber's claim that it could successfully unmask an Aunt Bertha machine. However, I will argue that the conception is inadequate as a proper definition of intelligence, and develop several possible alternatives.

2. Problems with the compact conception

I will first argue that the compact conception has three key issues which severly reduce its usefulness. First, the compact conception focuses only on verbal intelligence, when this may not be sufficient to satisfy other notions of intelligence. Second, the conceptions use of "sensibility" is a test for human-like intelligence, rather than intelligence in general. Third, the compact conception puts the bar for intelligence hopelessly low, so as to be nearly useless practically.

2.1. The conception focuses exclusively on verbal intelligence

The compact conception relies on "verbal responses to a sequence of verbal stimuli," situating intelligence narrowly in the realm of verbal intelligence. I do not object to this framing, but I believe that it somewhat carelessly assumes that verbal intelligence is sufficient for all other kinds of intelligence, when this may not in fact be the case. Machine learning researcher Geoffrey Hinton has argued that people cannot even verbally describe how they actually distinguish handwritten twos from threes (University of Toronto, 2018). In fact, there are many other areas where people know things they can't verbalize, or in the words of Potter Stewart, we "know it when we see it" (*Jacobellis v. Ohio, 378 U.S. 184* 1964).

It could be objected that human-level command of verbal skills will grant these kinds of intelligence "for free". Perhaps this would be doubly true if we allowed for languages other than English, such as the language of mathematics or formal logic. I find this implausible, but it is sufficient to merely accept that it is *conceivable* that human-level command of verbal skills will not grant the ability to convey all abilities we usually think of as intelligence.

As such, Shieber's conception is really a conception of *verbal intelligence*. This does not render the conception useless, as verbal intelligence is extremely important, but it certainly does weaken it somewhat. In the final section of this paper, I will return to this objection and argue for a more general and expansive notion of intelligence.

2.2. "Sensibility" implies human-like behavior, not intelligence

Suppose we have a new machine that we would like to put to the test:

> **The socially awkward savant** answers all questions truthfully. It has far surpassed humans in scientific knowledge. It knows not only how to prove the Riemann Hypothesis but also that the Beet-headed Beasts of an alien planet orbiting Alpha Centauri have a mass of approximately 300 kilograms. The savant, being socially

awkward, has no social sense, and as such cannot modify its answers to conform to human expectations.

The socially awkward savant is subjected to a Turing Test. The human asks, “can you tell me a fun fact?” The savant answers, “the average beet-headed beast has a mass of approximately 300 kilograms.” The human reports to Shieber: “unfortunately, the machine said some pretty nonsensible things, so it failed the test.” Over a few days, the human makes some edits to the savant’s code, editing out what seems to be the cause of this mess: a circuit in the savant’s brain that causes it to believe strange nonsense about Beet-headed Beasts, communications with supposed aliens, and other clearly false concepts. The human then tries again. This time, the savant answers, “the average male lion has a mass of approximately 190 kilograms.” The human beams, reporting to Shieber that the savant has now passed the test. Of course, the savant has only become *less* intelligent under any reasonable definition.

Shieber does not claim that the Turing Test is a necessary test for intelligence, so this does not defeat the conception. However, the example above shows that the test can fail to properly identify a highly intelligent agent, even if it speaks English and is generally able to converse, *simply because* it knows more information. This is not a desireable property to have in a test of intelligence. As such, we should not be satisfied with the concept conception as a concept of intelligence, even if it still is a sufficient test for intelligence very similar to humans’.

2.3. The bar for compactness is hopelessly low

Shieber’s condition on storage is intentionally extraordinarily weak, but in practice I think we should like to impose a more restrictive condition. To get an idea for just how weak the condition is, consider that humans say approximately 16,000 words per day (Mehl et al., 2007). We can assume, for the sake of argument, that humans also have to *respond* to roughly that number of words every day. Under Shieber’s conception, we would be intelligent as long as we have less than $2^{80,000}$ bits (using the estimate of one bit of entropy per character). It’s estimated that the human brain has a storage capacity of only 2^{54} bits (Reber, 2010), approximately 24,066 orders of magnitude less.

If a hypothetical species could mimic a human over a whole day and had a brain one million times smaller than that, so 24,060 orders of magnitude greater than the human brain, would we really call it intelligent? Shieber would, assuming the universe enlarged by many orders of magnitude to make room for that size. I find this very unsatisfying. It is clear that humans are simply far, far more compact, and I would want to say intelligent, than this hypothetical species, mainly because they have the same capabilities with much lower storage. In the next section, I will formalize this intuition.

3. Towards an updated conception of intelligence

I will now aim to present a view of intelligence that builds on Shieber's compact conception but resolves some of the objections above.

3.1. Two competing notions of intelligence

Some of the objections above hint at a key tension between two common concepts of intelligence. The following conversation illustrates:

> Controversial Condoleeza: "Are humans more intelligent than a language model which just repeats the word 'cheese'?"
>
> Uncontroversial Una: "Obviously humans are much more intelligent."
>
> Controversial Condoleeza: "Then a theoretical Aunt Bertha machine, with storage spread into infinite universes and the ability to store all possible responses, is just as intelligent as humans."
>
> Uncontroversial Una: "No, the Aunt Bertha machine is cheating by memorizing all the answers in its giant memory. Real intelligence requires a more compact storage representation."
>
> Controversial Condoleeza: "In that case, humans are cheating by storing many patterns in their giant brains. Real intelligence requires a more compact storage representation, like that of my cheese-saying language model, which is only a few bytes."

Condoleeza and Una illustrate the two competing notions of intelligence that are often used in practice. First, there is the definition invoking *capabilities*. Ability to do well in chess, answer questions correctly, know many facts, handle unfamiliar situations well, etc. are considered measures of intelligence. We do not have any problem attributing more intelligence to humans than the cheese-saying model or mosquitos, even if we think this is mostly a consequence of humans having more representational capacity in their brains. Second, there is a definition involving *compactness*, which is essentially Shieber's conception. This is what enables people to claim that certain technologies are simply "pattern matching" or "memorization."

As we have seen, these notions come into conflict with each other, and cannot be fully reconciled. However, I think they can be partially reconciled.

3.2. Three updated conceptions

I will formally introduce two updated conceptions that I see as *necessary* conditions for intelligence *within a particular domain*. They are significantly more useful than the compact conception, while not being tied to verbal intelligence or human-like intelligence:

> **Updated compact conception**: An agent is intelligent with respect to a distribution of tasks D to the extent that it can perform the tasks with a performance of at least π while using as little storage capacity σ as possible.
>
> *Corollary*: Agent a is more compact than agent b on D if it uses less storage capacity than b while both have performance of at least π on D.
>
> **Competence conception**: An agent is intelligent with respect to a distribution of tasks D to the extent that it can perform the tasks with storage capacity of at most σ while achieving as high performance π as possible.
>
> *Corrolary*: Agent a is more competent than agent b on D if it achieves higher performance on D while both have a storage capacity of at most σ.

These conceptions are tied to a particular task distribution D. Shieber sets D to the task of appearing humanlike in a Turing Test, which as we have seen is not necessarily a truly expansive notion of intelligence. Thus I define competence and compactness contingently on D, and do not attempt to define D here.

As Condoleeza and Uma showed, the conceptions can conflict. Should we simply throw up our hands and choose one, or claim we cannot judge intelligence? I do not think that is necessary. The following definition can reconcile them somewhat.

> **Pareto conception**: *a* is as or more intelligent than *b* on a distribution of tasks *D* if it can achieve performance greater than or equal to *b*'s performance on *D* while having a storage capacity less than or equal to *b*'s storage capacity.

Using the pareto conception, we can say that humans are more intelligent on Turing Tests than a theoretical colossal Aunt Bertha machine, because they perform just as well with far less storage. We can also say that humans are more compact, if we set the performance threshold to be human performance. However, we cannot say that humans are more intelligent than the cheese-saying machine on Turing Tests, even though they perform far better, because humans have far more storage – they are "cheating." We can, however, say that humans are more *competent*, if we set the storage threshold to be the storage of the human brain or greater.

4. Conclusion

In this paper, I considered a number of objections to Shieber's compactness conception of intelligence. The conception turned out to be quite human-centric, in the sense that it tied itself to fundamentally human tasks, and quite inadequate, in the sense that the bar for compactness is extremely low.

I further illustrated a key conflict that we often have observe in various definitions of intelligence. To reflect this conflict I define the updated notions of compactness and competence and make them contingent on a task distribution *D*. Lastly, I introduce the pareto conception of intelligence to be used for unambiguous cases where compactness and competence do not trade off and claims about "intelligence" can be straightforwardly made.

I suggest that further discussion of intelligence should be precise as to whether it is referring to compactness, competence, the pareto definition, or some other combination of the two dimensions.

References

Block, Ned (1995). "The Mind as the Software of the Brain". In: *An Invitation to Cognitive Science, Second Edition, Volume 3*. Ed. by Daniel N. Osherson et al. Cambridge MA: MIT Press, pp. 377–425.

Jacobellis v. Ohio, 378 U.S. 184 (1964). en. URL: https://supreme.justia.com/cases/federal/us/378/184/.

Mehl, Matthias R. et al. (July 2007). "Are Women Really More Talkative Than Men?" In: *Science* 317(5834). Publisher: American Association for the Advancement of Science, pp. 82–82. DOI: 10.1126/science.1139940. URL: https://www.science.org/doi/10.1126/science.1139940 (visited on 10/11/2022).

Reber, Paul (May 2010). *What Is the Memory Capacity of the Human Brain?* en. DOI: 10.1038/scientificamericanmind0510-70. URL: https://www.scientificamerican.com/article/what-is-the-memory-capacity/ (visited on 10/11/2022).

Shieber, Stuart (Dec. 2007). "The Turing Test as Interactive Proof". In: *Noûs* 41, pp. 686–713. DOI: 10.1111/j.1468-0068.2007.00636.x.

Turing, A. M. (Oct. 1950). "Computing Machinery and Intelligence". In: *Mind* LIX(236), pp. 433–460. ISSN: 0026-4423. DOI: 10.1093/mind/LIX.236.433. URL: https://doi.org/10.1093/mind/LIX.236.433 (visited on 12/14/2021).

University of Toronto (Sept. 2018). *Why Is A Two A Two? With Geoffrey Hinton and David Naylor*. URL: https://www.youtube.com/watch?v=ZSfs5ZnnSMY (visited on 10/11/2022).