

On the Impossibility of Supersized Machines

Ben Garfinkel , Miles Brundage¹, Daniel Filan^{2,3}, Carrick Flynn³,
Jelena Luketina , Michael Page , Anders Sandberg³, Andrew
Snyder-Beattie³, and Max Tegmark⁴

¹*School for the Future of Innovation in Society, Arizona State University*

²*Department of Computer Science, University of California, Berkeley*

³*Future of Humanity Institute, University of Oxford*

⁴*Department of Physics, Massachusetts Institute of Technology*

April 1, 2017

Abstract

In recent years, a number of prominent computer scientists, along with academics in fields such as philosophy and physics, have lent credence to the notion that machines may one day become as large as humans. Many have further argued that machines could even come to exceed human size by a significant margin. However, there are at least seven distinct arguments that preclude this outcome. We show that it is not only implausible that machines will ever exceed human size, but in fact impossible.

Introduction

The history of life is often understood as a story of growth. If one takes the long view, then one can trace an exponential curve from our minuscule earliest ancestors, which were little more than self-replicating molecules, to the substantial creatures that we are today (Payne, 2009).

Although humanity became aware of this story only in the 19th century, through the work of Charles Darwin, we have long had the privilege of witnessing a partial recapitulation every time someone new comes into the world (Darwin, 1859). Before each person is a full-sized adult, they are first an invisibly small cell.

It is perhaps no surprise, then, that human largeness has for thousands of years fascinated many of our greatest thinkers. While some have sought to understand the nature and origins of largeness, others have anxiously inquired: *Could there ever be something larger than a human?*

Evidence of this anxiety can be found as far back as humanity's oldest recorded myth, *The Epic of Gilgamesh*, in which the monstrous giant Humbaba is appointed by Enlil, the king of the gods, to terrorize mankind (Sandars,

1972). From this point onward, bellicose giants have been a consistent presence in our literature, appearing in works ranging from Homer’s *Odyssey* to the English fairytale “Jack and the Beanstalk” (Homer, 1994; Anonymous, n.d.).

Over time, perhaps in response to our species’ growing mastery of nature, it has become increasingly common to tell stories in which people are the ones responsible for the larger-than-human (or “supersized”) creatures that threaten them. For generations, audiences have been drawn to tales of frightful creations such as Frankenstein’s monster, described as over eight feet tall and “proportionally large”, and the golems of Kabbalah, which some rabbis feared would grow large enough to destroy the universe (Shelly, 2008; Moshe, 1990).

This archetype has perhaps never been more prevalent than it is in modern Hollywood films, however. Inspired by the apparently steady march of technological progress, and the wild speculations of futurists, our media has become saturated with images of murderous supersized machines.

In the long-running *Transformers* film series, machines known as Decepticons, each perhaps the size of a hundred men, repeatedly threaten to exterminate humanity with their enormous metal bodies (Bay, 2007). Numerous entries in the *Godzilla* film series feature machines so large that they can crush portions of the Tokyo skyline with a single step (Honda, 1975). We find that the *Matrix* film series, the *Terminator* film series (notable for its casting of an exceptionally large actor), and countless others also feature supersized machines that seek to cause the extinction of the human species (The Wachowskis, 1999; Cameron, 1984).

It does not help that in recent years a number of computer scientists, philosophers, and other academics have publicly lent credence to the possibility of supersized machines. There has been no shortage of media coverage of these figures’ pronouncements.¹

However, perhaps fortunately, all predictions of a coming age of supersized machines are fundamentally misguided. We present seven distinct arguments, each of which suffices to show that supersized machines are impossible.²

Arguments Against Supersized Machines

1. The Irreducible Complexity of the Human Body

Despite having been an active research area for hundreds of years, developmental biology has hardly progressed beyond its initial stages. We are far from being able to tell a story in all but the bluntest of terms of how a human zygote is

¹In addition, it has become very common for articles on recent trends in computer science to use terms such as “big data” and “massive neural networks” in ways that are likely to be misinterpreted. Reading these articles, even ones that appear in highly reputable newspapers, it is often unclear whether their authors are aware that the use of size language in these contexts is purely metaphorical.

²It is worth clarifying that there are, of course, systems today that appear to exceed human size in narrow dimensions. Lamp posts are one example. The predictions that we are considering concern some more general notion of largeness.

able to transform itself, over the course of two decades, into an adult that is several orders of magnitude larger (Cameron, 2012).

Scientists are at the point of being able to identify traits that correlate with largeness—certain genetic markers, for instance—but they have nothing like a complete theory of the causal pathways that explain these correlations. All attempts to construct such a theory have been stymied by the irreducible complexity of the human body, which contains tens of thousands of distinct proteins (Wilhem, 2014). It seems inevitable that, for this same reason, all future attempts will fail as well.

Since we cannot comprehend the processes responsible for human largeness, it follows that we will never be able to produce machines that surpass this largeness.

2. The Meaninglessness of “Human-Level Largeness”

One simple reason that we can reject predictions of supersized machines is that these predictions are not in fact well-formed.

The term “supersized machine” implies a machine that has crossed some threshold, which is often denoted “human-level largeness.” However, it is not clear what “human-level largeness” could refer to. Has a machine achieved human-level largeness if it has the same height as the average human? If it has the same volume? The same weight? Or some more complex trait, perhaps the logarithm of girth multiplied by the square of height?³

When one begins to consider these questions, one quickly concludes that there are an infinite number of metrics that could be used to measure largeness, and that people who speak of “supersized machines” do not have a particular metric in mind. Surely, then, any future machine will be larger than humans on some metrics and smaller than humans on others, just as they are today.

One might say, to borrow Wolfgang Pauli’s famous phrase, that predictions of supersized machines are “not even wrong” (Peierls, 1960).

3. The Universality of Human Largeness

A further reason why it is senseless to speak of machines that are larger than people is that humans already possess the property of universal largeness.

By this, we mean that humans are capable of augmenting their bodies or coming together to become indefinitely large, no matter the metric chosen. If a human would like to be taller, they can stand on a chair or climb onto another human’s shoulders. If they would like to be wider, they can begin consuming a high-calorie diet or simply put on a thick sweater (Hensrud, 2004; Figure 1). There are recorded cases of humans joining their bodies together to reach heights of up to 12 meters (Guinness World Records, 2013).

³Note also that humans vary quite significantly along all of these dimensions, and that even among humans there is no single accepted measure of largeness (Pomeroy, 2015).




Figure 1: There is no upper bound on how large humans can be (Gonzalez, 2017).

In short, since there exists no physical law to put an upper bound on human largeness, humans can be of any size. It follows, then, that no machine could ever really be larger than a human.

4. The Psychological Origins of Belief in Supersized Machines

By explaining why some people may be inclined to worry about supersized machines, evolutionary psychology reveals that such fears are not rational.

It is only natural that our ancestors should have developed a fear of beings larger than themselves. The greater a tribe member’s size, the more capable they are of employing violent coercion against other members or stealing their mates (Brewer, 2009). For this reason, vigilance toward the possibility of very large things was a highly advantageous trait.

Although largeness now plays a much-diminished role, at least in Western societies, there has been little time for human psychology to adapt (Donald, 1993). Furthermore, given the central role that technology plays in modern life, we should find it perfectly unsurprising that many people (especially “alpha males” enmeshed in Silicon Valley culture) have come to possess a fear of supersized machines.

Thus it is evolution, rather than logic or evidence, that serves as the true source of the belief that supersized machines are possible. It follows that we can safely assume this belief to be false.

5. Humans and Machines Together Will Always Be Larger Than Machines Alone

When writers discuss the possibility of supersized machines, they appear to be missing a crucial consideration: No machine could ever be larger than that same machine and a human together.

If machines are to play a role in pushing forward the frontier of largeness, then this role could only ever be to supplement human largeness.

This is another simple reason why it is senseless to imagine larger-than-human machines.⁴

6. The Hard Problem of Largeness

Suppose one were to concede that machines could become as large as humans, in some sense related to physical extension (although this is of course impossible).

Even if this were so, there would still remain a second, more meaningful sense of the word “large” that would not apply to these machines.

This second kind of largeness is the one evoked whenever someone is described as “larger than life” or “living large” (Tom, 2004). Largeness of this sort is a non-physical (i.e. non-natural) property, separate from the mundane physical property that “largeness” most often denotes.

To build a large machine, then, in the meaningful sense, we would first need to solve the “hard problem” of determining what this non-physical property is and how it arises. However, it is not at all clear that the problem is soluble, since the traditional methods of science seem equipped only to deal with questions that concern the physical world (Hall, 2010). Furthermore, the notion of a machine “living large” strikes one as intuitively implausible (perhaps even absurd).

Therefore, machines will never *truly* be large.

7. Quantum Mechanics and Gödel’s First Incompleteness Theorem

Quantum theory, as traditionally formulated, divides the world up into microsystems and macrosystems (Heisenberg, 1949). Within microsystems lie small objects, such as particles, and within macrosystems lie large objects, such as humans.

The theory tells us that objects in microsystems may initially have no definite properties at all, such that any question concerning a given particle’s position, momentum, and so forth, will simply lack an answer. However, the remarkable ability that humans possess, as a result of their largeness, is the ability to force objects in microsystems to take on definite properties by performing “measurements” on them. For example, if a human “measures” that a particle has a certain location, then it becomes a new fact that the particle has this location.

⁴This consideration also suggests that credible machine largeness researchers ought to focus on human-machine interfaces, which enable size-enhancing machines to be attached directly to the human body. Existing work on stilts may suggest one promising research direction (Smith, 2010).

One of the great mysteries of quantum mechanics, that its originators never succeeded in resolving, is the question of what distinguishes microsystems from macrosystems (Bell, 1990). It seems that we are to understand that some fundamental line separates the large from the small, such that small objects exist in a sort of limbo until large objects perform measurements on them. However, we lack guidance on how to draw this line, and it is difficult to understand how and why the line exists at all. The problem of making sense of this line, and thereby uncovering the nature of largeness, is known as “the measurement problem.”

A partial answer to the measurement problem may be suggested by Kurt Gödel’s first incompleteness theorem (Gödel, 1931). This theorem was first proved in 1931, although its full significance arguably remains to be appreciated. The theorem states that, for any sufficiently expressive formal system, the system must either be inconsistent or incapable of proving true or false all statements that are expressible within the system.

To understand how Gödel’s theorem can resolve the measurement problem, it is perhaps most useful to apply the lens of quantum stochastic calculus (Kholevo, 1991). QSC, as a reminder, generalizes classical stochastic calculus to cover cases of non-commuting random variables, which are ubiquitous in quantum mechanics. Take the quantum Stratonovich integral of a system operator, $g(t)$, which is given by (Gardener, 2004):

$$(S) \int_{t_0}^t g(t') dB(t') = \lim_{n \rightarrow \infty} \sum_{i=1}^n \frac{g(t_i) + g(t_{i+1})}{2} (B(t_{i+1}, t_0) - B(t_i, t_0))$$

Applying this expression, it is trivial to show that:

$$(S) \int_{t_0}^t g(t') dB(t') - (S) \int_{t_0}^t dB(t') g(t') = \frac{\sqrt{\gamma}}{2} \int_{t_0}^t dt' [g(t'), c(t')]$$

Now suppose that we would like to formalize this deduction within non-well-founded set theory, to which Gödel’s theorem of course applies (Aczel, 1988). Importantly, by assuming the axiom of anti-foundation we are able to introduce into our analysis self-referential objects, such as Quine atoms, which possess the property of being large enough to contain themselves.

Although the technical details from this point onward are unfortunately too dense to include in a general-audience essay of this sort, assuming as they do familiarity with constructive non-standard analysis, it suffices to say that supersized machines cannot be made (Figure 2).

Conclusion

We have presented seven distinct arguments against the possibility of supersized machines. While each of these arguments would be sufficient on its own, the




Figure 2: Machines cannot be large (Gonzalez, 2017).

conjunction of them surely constitutes an insurmountable barrier to the belief that an age of supersized machines lies anywhere on the horizon.⁵

Our conclusion is in at least one way a relief. There is no reason to fear preposterous stories about towering Terminator machines.

However, our conclusion might also be taken as a sad one. We are the largest things in the universe, and we will never be otherwise.

Fantasies of supersized machines hold an appeal, in addition to inspiring fear, because it is tempting to imagine these machines as perfected versions of ourselves. They are who people would be if only we were a little larger. They are steadier, and more able to look down upon the world with a distant wisdom, rather than becoming entangled in the insignificant details close to ground.

It can be nice to think that if we are unable to resolve our own problems here on Earth, then maybe this is only because we lack the size.

A world in which we are the largest things conceivable is a world without excuses. We submit that this is a good thing, however. It is time to stop daydreaming about something larger than ourselves, and time to begin understanding how large we truly are.

⁵One may wonder why we have felt it necessary to demonstrate that supersized machines are impossible, rather than arguing for the much weaker claim that supersized machines are unlikely to be developed soon. The reason is that, counter-intuitively, many of the academics who have expressed concern about supersized machines appear to accept this weaker claim. They argue from the position, currently controversial among policy-makers, that it is worth preparing for distant or low-probability events (Bedford, 2001). This position has led many to stake out similarly provocative stances in favor of climate change mitigation, pandemic preparedness, and seatbelt use.

References

- [1] Aczel, Peter. "Non-well-founded sets." (1988).
- [2] Anonymous. "Jack and the Beanstalk."
- [3] Bay, Michael. *Transformers*. DreamWorks Pictures, 2007. Film.
- [4] Bedford, Tim, and Roger Cooke. *Probabilistic Risk Analysis: Foundations and Methods*. Cambridge University Press, 2001.
- [5] Bell, John. "Against measurement." *Physics world* 3.8 (1990): 33.
- [6] Brewer, Gayle, and Charlene Riley. "Height, relationship satisfaction, jealousy, and mate retention." *Evolutionary Psychology* 7.3 (2009).
- [7] Cameron, James. *The Terminator*. Orion Pictures, 1984. Film.
- [8] Cameron, Noël, and Barry Bogin. *Human growth and development*. Academic Press, 2012.
- [9] Darwin, Charles. *On the Origin of Species*. 1859.
- [10] Donald, Merlin. "Human cognitive evolution: What we were, what we are becoming." *Social Research* (1993): 143-170.
- [11] Gardiner, Crispin, and Peter Zoller. *Quantum noise: a handbook of Markovian and non-Markovian quantum stochastic methods with applications to quantum optics*. Vol. 56. Springer Science & Business Media, 2004.
- [12] Glenday, Craig, ed. *Guinness world records 2013*. Bantam, 2013.
- [13] Gonzalez, Ian. Message to Ben Garfinkel. 28 Mar. 2017. Facebook.
- [14] Gödel, Kurt, Stephen Cole Kleene, and John Barkley Rosser. "On undecidable propositions of formal mathematical systems." Princeton, NJ: Institute for Advanced Study, 1934.
- [15] Hall, Ned. "The Natural/Non-Natural Distinction", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition). Web.
- [16] Heisenberg, Werner. *The physical principles of the quantum theory*. Courier Corporation, 1949.
- [17] Hensrud, Donald D. "Diet and obesity." *Current opinion in gastroenterology* 20.2 (2004): 119-124.
- [18] Homer. *Odyssey*. Dorset Press, 1994.
- [19] Honda, Ishiro. *Terror of Mechagodzilla*. Toho, 1975. Film.
- [20] Idel, Moshe. *Golem: Jewish magical and mystical traditions on the artificial anthropoid*. State University of New York Press, 1990.

- [21] Kholevo, A. S. "Quantum stochastic calculus." *Journal of Mathematical Sciences* 56.5 (1991): 2609-2624.
- [22] Payne, Jonathan L., et al. "Two-phase increase in the maximum size of life over 3.5 billion years reflects biological innovation and environmental opportunity." *Proceedings of the National Academy of Sciences* 106.1 (2009): 24-27.
- [23] Peierls, R. E. "Wolfgang Ernst Pauli. 1900-1958." *Biographical Memoirs of Fellows of the Royal Society* 5 (1960): 175-192.
- [24] Pomeroy, Ross. "A New Potential Replacement for Body Mass Index." *RealClearScience*. 30 Dec. 2013. Web.
- [25] Sandars, Nancy, ed. *The epic of Gilgamesh*. Penguin, 1972.
- [26] Shelley, Mary Wollstonecraft. *Frankenstein, Or, The Modern Prometheus: 1818*. Engage Books, 2008.
- [27] Smith, Kevin D., and Patrick Chaney. "Vertically Adjustable Stilts." U.S. Patent 7,794,336. 14 Sep. 2010.
- [28] The Wachowskis. *The Matrix*. Warner Bros., 1999. Film.
- [29] Tom. "living large." *UrbanDictionary*. Web.
- [30] Wilhelm, Mathias, et al. "Mass-spectrometry-based draft of the human proteome." *Nature* 509.7502 (2014): 582-587.