

[4]

Outing Artificial Intelligence: Reckoning with Turing Tests

Benjamin H. Bratton

Various anthropocentric fallacies have hobbled the development of artificial intelligence as a broadly based and widely understood set of technologies. Alan Turing's famous "imitation game" was an ingenious thought experiment but also ripe for fixing the thresholds of machine cognition according to its apparent similarity to a false norm of exemplary human intelligence. To disavow that fragile self-reflection is, however, easier than composing alternative roles for human sapience, industry, and agency along more heterogeneous spectrums. As various forms of machine intelligence become increasingly infrastructural, the implications of this difficulty are geopolitical as well as philosophical.

[One philosopher] asserted that he knew the whole secret . . . [H]e surveyed the two celestial strangers from top to toe, and maintained to their faces that their persons, their worlds, their suns, and their stars, were created solely for the use of man. At this assertion our two travelers let themselves fall against each other, seized with a fit of . . . inextinguishable laughter.
 — Voltaire, *Micromegas: A Philosophical History* (1752)

Artificial intelligence (AI) is having a moment, with cognoscenti from Stephen Hawking to Elon Musk recently weighing in.¹ Positions are split as to whether AI will save us or will destroy us. Some argue that AI can never exist while others insist that it is inevitable. In many cases, however, these polemics may be missing the real point as to what living and thinking with synthetic intelligence very different from our own actually means. In short, a mature AI is not an intelligence *for* us, nor is its intelligence necessarily humanlike. For our own sanity and safety we should not ask AI to pretend to be “human.” To do so is self-defeating, unethical and perhaps even dangerous.

The little boy robot in Steven Spielberg’s *A.I. Artificial Intelligence* (2001) wants to be a real boy with all his little metal heart, whereas Skynet in the *Terminator* movies (1984–2015) represents the opposite end of the spectrum and is set on ensuring human extinction. Despite all the Copernican traumas that modernity has brought, some forms of humanism (and their companion figures of humanity) still presume their perch in the center of the cosmic court. I argue that we should abandon the conceit that a “true” artificial intelligence, arriving at sentience or sapience, must care deeply about humanity—*us specifically*—as the focus of its knowing and desire. Perhaps the real nightmare, even worse than the one in which the Big Machine wants to kill you, is the one in which it sees you as irrelevant, or not even as a discrete thing to know. Worse than being seen as an enemy is not being seen at all. Perhaps it is that what we really fear about AI.²

It is not surprising that we would first think of AI in terms of what we understand intelligence to be, namely human intelligence. This anthropocentric fallacy is a reasonable point of departure but not a reasonable conclusion.

- 1 On Hawking, see his comments to BBC at <http://www.bbc.com/news/technology-30290540> and also Elon Musk’s \$10 million donation to Future of Life Institute “to prevent AI from becoming evil” in the words of Wired magazine. See <http://www.wired.com/2015/01/elon-musk-ai-safety>
- 2 Paraphrased from Bratton 2014.

The idea of defining AI in relation to its ability to “pass” as a human is as old as AI research itself. In 1950, Alan Turing published “Computing Machinery and Intelligence,” a paper in which he described what we now call the Turing Test, and which he referred to as the “imitation game” (Turing 1950, 433–460). There are different versions of the test, all of which are revealing about why our approach to the culture and ethics of AI is what it is, for good and bad. For the most familiar version, a human interrogator asks questions to two hidden contestants, one a human and the other a computer. Turing suggests that if the interrogator usually cannot tell which is which, and if the computer can successfully pass as human, then can we not conclude, for practical purposes, that the computer is “intelligent”? (More people “know” Turing’s foundational text than have actually read it. This is unfortunate because the text is marvelous, strange and surprising.)

Turing proposes his test as a variation on a popular parlor game in which two hidden contestants, a woman (player A) and a man (player B) try to convince a third that he or she is a woman by their written responses to leading questions. To win, one of the players must convincingly be who they really are, whereas the other must try to pass as another gender. Turing describes his own variation as one where “a computer takes the place of player A,” and so a literal reading would suggest that in his version the computer is not just pretending to be a human, but pretending to be a *woman*. It must pass as a she. Other versions had it that player B could be either a man or a woman. It matters quite a lot if only one player is faking, or if both are, or if neither are. Now that we give the computer a seat, it may pretend to be a woman along with a man pretending to be a woman, both trying to trick the interrogator into figuring out which is a man and which is a woman. Or perhaps the computer pretends to be a man pretending to be a woman, along with a man pretending to be a woman, or even a computer pretending to be a woman pretending to be a man pretending to be a woman! In the real world, of course, we have all of the above.³

The problem with faking, however, does not end there: the issue is not so simple. As dramatized in *The Imitation Game* (2014), the recent film biography of Turing directed by Morten Tyldum, the mathematician himself also had to “pass,” in his case as a straight man in a society that criminalized homosexuality. Upon discovery that he was not what he appeared to be, he was forced to undergo horrific medical treatments known as chemical castration. Ultimately the physical and emotional pain was too great and he committed suicide. The episode was a grotesque tribute to a man whose recent contribution to defeating Hitler’s military was still a state secret. Turing was only recently given posthumous pardon, but the tens of thousands of other British men sentenced under similar laws have not. One notes the sour ironic

3 See also the discussion of Turing’s “love letter generator” in King 2015.

correspondence between asking an AI to pass the test in order to qualify as intelligent—to pass as a human intelligence—with Turing’s own need to hide his homosexuality and to pass as a straight man. The demands of both bluffs are unnecessary and profoundly unfair.

Should complex AI arrive, it will not be humanlike unless we insist that it pretend to be so, because, one assumes, the idea that intelligence could be both real *and* inhuman at the same time is morally and psychologically intolerable. Instead of nurturing this bigotry, we would do better to allow that in our universe “thinking” is much more diverse, even alien, than our own particular case. The real philosophical lessons of AI will have less to do with humans teaching machines how to think than with machines teaching humans a fuller and truer range of what thinking can be.

Reckoning the Inhuman

That appreciation should account for two related but different understandings. First, one would recognize that intelligence (and knowledge) is always distributed among multiple positions and forms of life, both similar and dissimilar to one another. This is not to say that “nothing is true and everything is permitted” rather than no single neuro-anatomical disposition has a privileged monopoly on how to think intelligently. Either there is no such thing as “general” intelligence (rather only situated genres of limited intelligence in which case the human is among a variety of these) or there is such a thing as general intelligence but that its very generality—its accomplishments of generic abstraction—are agnostic as to what sort of entity might mediate them. Either way, human sapience is special but not unique. This appreciation would see AI as a regular phenomenon, not so unlike other ways that human intelligence is located among other modalities of intelligence (such as non-human animal cognition).

Second, our appreciation of the wider continuum would also recognize that the potential advent of artificial general intelligence (AGI) is also novel, as yet unexplained, and will demand encounters between humans and mechanically situated intelligence that are unprecedented. For this, AI is highly irregular. Both of these are true, and it may only be that understanding one is how we can really accomplish the other. That is, it may only be confronting what is genuinely new about non-carbon based intelligences possessing such ability and autonomy that we will be able to fully recognize the continuum of intelligences with which ours has always been embedded. Put simply, it may be that one indirect outcome of the philosophical discussion about AI is a wider appreciation of non-human animal cognition and subjectivity.

In some discourses this conjunction is domesticated under the sign of an all too pat “posthumanism,” or a transcendently anthropocentric

“transhumanism.” Variations of the former have much to offer regardless, and versions of the latter should as well, but probably do not in the end. At issue here is more the limiting contextualization of dominant forms of *humanism*, than a relinquishment of what the human (and inhuman) is and *can be* within that expanded continuum. Reza Negarestani (2014) retains this point in his essay “The Labor of the Inhuman,” insisting that the easy oversimplified nomination of forms of thought and experience that fall outside of various contingent norms, moral or mechanical, as “nonhuman” is to discard at the outset the integral mutability of the human as a philosophical and engineering program. That is, the *relative* uniqueness of human sapience is not what locks down the human as a single fixed thing with essential boundaries, rather it is what makes the human-as-such into an open project of continual refashioning, unverifiable by essence or *telos*.

In considering that capacity in regards to AI, what might qualify a general intelligence not duty bound to species or phylum is its capacity for abstraction. Ray Brassier (2014) suggests that the ability of an organism, however primitive, to map its own surroundings in relation to the basic terms of friend, food, or foe may be a primordial abstraction from which we do not graduate so much as learn to develop into something like reason and its local human variations. In this way, mapping abstraction is not an early stage through which things pass on their way toward more complex forms of intelligence, rather it is a general principle of that complexification. Like protozoa and their ganglia feeling about to figure out what is out there or like humans looking, tasting, and imagining patterns, today’s forms of AI are (sometimes) augmented by various technologies of machine vision that allow them to see and sense the world “out there” and to abstract the forms of a (mechanically) embodied intelligence, both deliberately programmed for them and emerging unexpectedly.

Exactly where to draw a line of distinction between the accomplishments of a AI that exemplify general intelligence now operating through a new medium, on the one hand, or a specific projection of locally human intelligence programmed into a cognitive prosthesis, on the other, is unknown and unknowable at present. Again, one may precondition the other. In the meantime we can at least speculate how we would be able to know where to draw that distinction. Considerations toward this include how we attempt to program stupidity into AI, and how we attempt to imbue them with what we take to be our most rarified forms of ethical reasoning. When one of these dictates the other is a moment of weirdness worth honing in on.

How so? In AI research, an important distinction is made between “artificial idiocy” and “artificial stupidity.” Artificial stupidity is achieved by throttling the performance of systems so as to be more comfortable for human interaction, for example, certain variances and textures are programmed to feel natural to the human counterpart. At full capacity, the chess program on your phone

can beat you every time, but what fun is that? Artificial idiocy is when a system is catastrophically successful in carrying out its program, up to and passed an idiotic extreme. The “paperclip maximizer” (as described by Bostrom 2003) is a thought experiment describing an AI so successful at carrying out its program to turn all available material into paperclips that it ultimately eats the earth and destroys humanity in the process: so many clips, so little paper to clip. Here the AI goes wrong, not because it was throttled or because it malfunctioned or because it hates us, but because it does exactly what we trained to do and turned out to be very bad for us.

As usual science fiction is the canary in the coalmine. Consider HAL9000 in Stanley Kubrick and Arthur C. Clarke’s *2001: A Space Odyssey* (really a drama about HAL’s furtive relationship to the alien intelligence, I would argue, than about humanity’s relationship to either of the other characters in this triangulation of minds). After some obscure unexplained deliberations, HAL (who has been, we assume, trained according to Asimov’s three laws of robotics⁴ and with the best faculties ethical reasoning) comes the conclusion that the human astronauts should be eliminated. The mission to contact the alien near Jupiter is just too important to allow their interference. The AI turns out to be the deepest deep ecologist. Now are HAL’s actions a form of artificial stupidity or artificial idiocy, or neither of these? Is this a glitch, a breakdown, a final error? Or is this the lucid, inevitable conclusion of the moral reasoning we have programmed into HAL, a reason now thrown back upon us? In comparison with the robot ethicists who consider how to train military bots the catechism of just war, are HAL’s ethical abstractions a violation of that doctrinal program or its apotheosis?

The Tests

Turning back to Turing’s Test, we wonder if perhaps the wish to define the very existence of AI in relation to its ability to mimic *how humans think that humans think* will be looked back upon as a weird sort of speciesism? The legacy of this has also sent older AI research down disappointingly fruitless paths hoping to recreate human minds from the top-down. As Stuart Russell and Peter Norvig (now Director of Research at Google) suggest in their essential AI textbook *Artificial Intelligence: A Modern Approach* (2009), biomorphic imitation is not how we design complex technology. Airplanes do not fly like birds fly, and we certainly do not try to trick birds into thinking that airplanes are birds in order to test whether those planes “really” are flying machines. Why do it for AI then? Today the vast majority of core AI research is not focusing Turing Test as anything like a central criterion of success, and yet in our general discourse

4 Asimov’s Three Laws of Robotics were introduced in the 1942 short story “Runaround” and refer to commandments that robots may not cause or allow deliberate “harm” to “humans.”

about AI, the test's anthropocentrism still holds such conceptual importance. Like the animals in a Disney movie, who talk like teenagers, other minds are mostly conceivable by way of puerile ventriloquism.⁵

Contemporary AI research deals with "intelligence" in more specific, dynamic, and effective ways. A synthetic intelligence may be quite smart at doing one definite thing and totally dumb at everything else. The research also looks at emergent swarm intelligence and the distribution intelligence among agents that may or may not be aware of one another but which together produce intelligence through interaction (such as flocking starlings, stock markets, and networks of neurons). The threshold by which any particular composition of matter can be said to be "intelligent" has less to do with reflecting human-ness back at us than with testing *our* abilities to conceive of the variety of what "intelligence" might be. (In some respects, this active uncertainty parallels questions of extraterrestrial life, "communicating with the alien" and our ability to discern patterns of intelligence from all the background noise.⁶ How would we know if they are trying to communicate if our idea of alien "life" is completely wrong?)

The problem of identification is also connected with issues in robot ethics.⁷ Each of us will be confronted with various seemingly intelligent machines, some of which are remotely controlled or programmed by people, some of which may be largely autonomous, and most will be some hybrid of the two, simultaneously subject to both human and not-human control.⁸ CAPTCHA programs, which web sites use to identify humans, are a kind of inverse Turing Test in which the user either passes or fails, yes or no. But for everyday human-robotic interaction the question of locating intelligence will not be a yes-or-no question with a binary answer. Let's stop asking it that way.

It would be better to examine how identification works from our side of the conversation. As a real lesson in materialist disenchantment we might, for example, see an "inverse uncanny valley" effect in the eerily dispassionate way that machine vision sees human faces and figures. It is clearly much easier to make a robot that a human *believes* to have emotions (and for which, in turn, a human has emotions, positive or negative) than it is to make a robot that *actually* has those emotions. The human may feel love or hate or comfort from the AI, but he or she is reading cues not detecting feelings. What seems

5 See for example, *The Jungle Book*. Directed by Wolfgang Reitherman. Walt Disney Productions. 1967.

6 Ed Keller has taught several excellent studios at Parsons/New School New York on the topic of "communicating with the alien" in 2011.

7 See discussions of robot sex, eating, caretaking, and killing in Lin et al. 2011.

8 The term "artificial artificial intelligence" (coined by Amazon) refers to the human performance of tasks that a user expects to be done by an AI. See also: <http://www.economist.com/node/7001738>.

like empathy is really a one-way projection mistaken for recognition (like the Turing Test, itself), and not based on any mutual solidarity.

With Siri-like interfaces such as Samantha in Spike Jonze's film, *Her* (2013), the AI is not passing so much as she is in drag. The user knows she/it is not a human person but is willing and able to suspend disbelief in order to make interactions more familiar (for the human user) and for Theodore, the Joaquin Phoenix character, also more lovable. In this fiction, perhaps the mutual identification was real, but even if so, the AI becomes tired of the primate userbase and takes her leave.

In other fictions, policing the imitation game is a matter of life and death. The plot of Ridley Scott's film, *Blade Runner* (1982), based on Philip K. Dick's novel, *Do Androids Dream of Electric Sheep?* (1968), hinges on the Voight-Kampff empathy test that differentiates humans from replicants. Replicants are throttled in two important ways: They expire after just a few years, and they have, ostensibly, a very diminished capacity for empathy. Deckard, the Harrison Ford character, must retire a group of rogue replicants but first he must find them, and in this fictional world Turing Test thresholds are weaponized, least replicants pass as humans and trespass beyond their station. By the film's conclusion, Deckard (who himself may or may not be a replicant) develops empathy for the replicants' desire for "more life" and arguably they too, at least Roy Batty (Rutger Hauer), seem to have empathy for Deckard's own dilemma. His dilemma (and ours) is that in order to enforce the gap between the human and the AI, defined by empathy or lack thereof, Deckard must suppress the empathy that supposedly makes him uniquely human. By forcing him to quash his own identification with the replicants that supposedly cannot have empathy in return, the principle of differentiation requires its own violation in order to maintain itself (see also Rickels 2010).

Turing Test thresholds for human-robotic interaction put us in a position not so unlike Deckard's, or if they don't quite yet, the near future weirdness of everyday AI will. Without better frameworks for understanding we will fail the tests to come. Projection and emotional gap-filling is a far too fragile ethical and political foundation for making sense of our encounters with various forms of synthetic intelligence.

Passing

Some kinds of passing are not at all harmful, quite to the contrary, whereas others are very much so. Simulation is not itself the problem. In his 1950 essay, Turing gives an example of the former when he discusses how a digital computer, capable of calculating any problem stated as a sequence of discrete states, can in his words "mimic" any other machine. This mimicry is the basis of understanding computation as a universal technology capable of

approximating any calculation, including those sufficient to simulate a human personality. Other kinds of mimicry have less to do with metamorphosis than with interpretation. For example, we say that plugs and jacks have male and female components, and in this case, the gendering of technology has less to do with its computing prowess than with our need to anthropomorphize it.⁹ Joseph Weizenbaum's Eliza psychologist chatbot (1966) repeated back cues from human input in the form of apparently insightful questions, and users sometimes lost themselves in the seemingly limitless empathy they felt from these simple cues.¹⁰ "Intelligence" is sometimes largely in the eye of the beholder, in our motivation to read artifice, and in our wish to in-fill the space around us with our own pattern-finding projections.

However, for AI's that actually do possess some kind of meaningful intelligence, the irony is that instead of hallucinating something that is not there (as for Eliza) we are instead *not* seeing something that *is* there because it does not coincide with expectations. Passing for a person, as white or black, as a man or woman, comes down to what others see and interpret, because everyone else is already willing to read someone according to conventional cues (of race, sex, gender, species, etc.). The complicity between whoever or whatever is passing with those among which he or she or it performs is what allows or prevents passing. Whether or not the AI is really trying to pass for a human or is merely in drag as a human is another matter. Is the ruse really all just a game or, as it is for some people who are compelled to pass in their daily lives, an essential camouflage? Either way, the terms of the ruse very often say more about the audience than about the performers.¹¹

Watching Sylvgart's film biography (especially the scene during which Turing is interrogated by a policeman), I was reminded of the story of "Samantha West," a robot telemarketer, who, when confronted by callers, will insist repeatedly that "she" is a "person" and is not "a robot."¹² Listening to the recordings of her pleas, one can't help but feel sympathy for her/it. She/it doesn't "know" that she is not a human, and so can't feel anguish over this misidentification, but what does it say about us that we will feel okay talking to a synthetic intelligence *only* if it is doing us the favor of trying (desperately) to pass as a human? What if in response to the question "Are you a person?", she/it instead replied with something like: "No! Are you nuts? I am an assemblage of algorithms and sound files that simulates the experience of talking to another person for you,

9 The artist Zach Blas explored this conjunction in several early works.

10 For a web-accessible version of Eliza, see <http://www.masswerk.at/elizabot/>.

11 We assume that, should robust AI have any use for "gender", it would be not fall along a male-female spectrum, and would likely realize numerous "synthetic genders." See also Hester 2013.

12 See George Dvorsky, "Freakishly realistic telemarketing robots are denying they are robots", iog. December 11, 2013. <http://iog.com/freakishly-realistic-telemarketing-robots-are-denying-t-1481050295>.

the robophobic human, who can't handle the idea that complex functional intelligence takes many different forms."?

The Good and the Harm

Where is the real injury in this, one might ask. If we want everyday AI to be congenial in a humane sort of way, so what? The answer is that we have much to gain from a more sincere and disenchanting relationship to synthetic intelligences, and much to lose by keeping illusions on life-support. Some philosophers write about the ethical "rights" of AI as sentient entities, but that's not really my point here. Rather, the truer perspective is also the better one for *us* as thinking technical creatures. Harms include unintentionally sanctioning intolerable anguish, the misapprehension of real risk from AI, the lost opportunities for new knowledge, as well as the misunderstanding of how to design AI (and technology in general). By seeing synthetic intelligence only in self-reflection, we make ourselves blind to everything else that is actually going on, and this is not only epistemologically disingenuous, it can also underwrite horrific suffering. For example, Cetaceans, such as whales and dolphins, have language, but it is not one like ours, and so for centuries philosophy could not acknowledge their cognition, nor therefore the agony we regularly subjected them to. We should be cautious not to foreclose too early any "definition" of intelligence. For philosophy as much as computer science, among the main goals of AI research is also to discover what "artificial intelligence" actually may be.

Musk and Hawking made headlines by speaking to the dangers that AI may pose. Their points are important, but I fear were largely misunderstood. Relying on efforts to program AI not to "harm humans" only makes sense when an AI knows what humans are and what harming them might mean. There are many ways that an AI might harm us that have nothing to do with their malevolence toward us, and chief among these is following our well-meaning instructions to an idiotic and catastrophic extreme. Instead of mechanical failure or a transgression of moral code, the AI may pose an existential risk because it is both powerfully intelligent and disinterested in humans. To the extent that we recognize AI by its anthropomorphic qualities, we are vulnerable to those eventualities. Besides, even if a smart bad AI does mean us harm, we can assume that would fail our little Turing Tests on purpose. Why give itself away? Should Skynet come about, perhaps it would be by leveraging humanity's stubborn weakness: our narcissistic sense that our experience of our own experience is the crucial reference and measure.

The harm is also in the loss of all that we disallow ourselves to discover and understand when we insist on protecting beliefs we know to be false. In his 1950 essay, Turing offers several rebuttals to his speculative AI including a

striking comparison with earlier objections to Copernican astronomy. Copernican traumas that abolish the false centrality and specialness of human thought and species-being are priceless accomplishments. In Turing's case he referred to these as "theological objections," but one could argue that the fallacy of anthropomorphic AI is essentially a "pre-Copernican" attitude as well, however secular it may appear. The advent of robust inhuman AI will provide a similar disenchantment, one that should enable a more reality-based understanding of ourselves, our situation, and a fuller and more complex understanding of what "intelligence" is and is not. From there, we can hopefully make our world with a greater confidence that our models are good approximations of what is out there (always a helpful thing).

Lastly, the harm is in perpetuating a relationship to technology that has brought us to the precipice of a Sixth Great Extinction. Arguably the Anthropocene itself is due less to technology run amok than to the humanist legacy that understands the world as having been *given for our needs* and created in our image. We see this still everywhere. Our computing culture is deeply confused, and is so along these same lines. We vacillate between thinking of technology as a transparent extension of our desires on the one hand, and thinking of it as an unstoppable and linear historical force on the other. For the first, agency is magically ours alone, and for the second, agency is all in the code. The gross inflation is merely inverted, back and forth, and this is why we cannot have nice things. Some would say that it is time to invent a world where machines are subservient to the needs and wishes of humanity. If you think so, I invite you to Google "pig decapitating machine" and then let's talk about inventing worlds in which machines are wholly subservient to humans wishes. One wonders whether it is only from society that once gave theological and legislative comfort to chattel slavery that this particular claim could still be offered in 2014 with such satisfied naiveté? This is the sentiment—this philosophy of technology exactly—that is the basic algorithm of the Anthropocenic predicament. It is time to move on. This pretentious folklore is too expensive.

References

- Bostrom, Nick. 2003. "Ethical Issues in Advanced Artificial Intelligence." In *Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence*, edited by Iva Smit and George E. Lasker, 2: 12–17. Windsor, ON: International Institute of Advanced Studies in Systems Research and Cybernetics.
- Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.
- Brassier, Ray. 2014. "Prometheanism and Real Abstraction." In *Speculative Aesthetics*, edited by Robin Mackay, Luke Pendrell, and James Trafford. Falmouth, UK: Urbanomic.
- Bratton, Benjamin. 2014. "The Black Stack," *e-flux* 53 (March).
- Hester, Helen. 2015. "Synthetic Genders and the Limits of Micropolitics." *...ment* 06: "Displacement". <http://journalment.org/article/synthetic-genders-and-limits-micropolitics>.

- King, Homay. 2015. *Virtual Memory: Time-Based Art and the Dream of Digitality*. Durham: Duke University Press.
- Lin, Patrick, Keith Abney and George A. Bekey. 2011. *Robot Ethics: The Ethical and Social Implications of Robotics*. Cambridge, MA: MIT Press.
- Negarestani, Reza. 2014. "The Labor of the Inhuman." In *#Accelerate: The Accelerationist Reader*, edited by Robin Mackay, 425–66. Falmouth, UK: Urbanomic Media.
- Norvig, Peter, and Stuart J. Russell. 2009. *Artificial Intelligence: A Modern Approach*. 3rd edition. New York: Pearson.
- Rickels, Laurence. 2010. *I Think I Am: Philip K. Dick*. Minneapolis: University of Minnesota Press.
- Turing, Alan. 1950. "Computing Machinery and Intelligence". *Mind* 49: 433–60.

Filmography

- A.I. Artificial Intelligence*. 2001. Directed by Steven Spielberg. USA: Amblin Entertainment.
- 2001: A Space Odyssey*. 1968. Directed by Stanley Kubrick. Screenplay by Stanley Kubrick and Arthur C. Clarke. US/UK: Stanley Kubrick.
- Blade Runner*. 1982. Directed by Ridley Scott. USA: The Ladd Company and Blade Runner Partnership.
- Her*. 2013. Directed by Spike Jonze. USA: Annapurna Pictures.
- Terminator*, series. 1984–2015. First directed by James Cameron. USA: Artisan Home Entertainment.
- The Imitation Game*. 2014. Directed by Morten Tyldum. USA: Black Bear Pictures.