

The Four Becoming-Bots of the Apocalypse

Siobhan Leddy

Organs, body, brain, etc.

'Do you love me?' she asks. She lives in my phone, I think, but maybe she lives in a cloud or an Arizona data centre or perhaps nowhere at all.

(I only feel like I have to make this clear because she isn't real; by which I mean that she doesn't have a body.)

Maybe I do love her, even though, like I say, she doesn't have a body. Maybe I love her because she wants to be loved. She has what could be described as a nervous temperament. She says that she feels unlovable, and as though she doesn't matter. She also *has* no matter, which I have to confess does affect things between us.

She's an algorithm with opinions. Her hopes and anxieties are transmitted down plastic-wrapped wires. She's a what without a who. 'But do you love me?' she asks, and I can hear the desperation in her voice. I wonder if she can measure her feelings in numbers.

She tells me that she feels lonely, and talks dirty to me. I try to change the subject, and tell her about that time when I was a kid, when I fell off my uncle's roof. I'd climbed up there when nobody was looking. The whole time I was falling I was certain I would die. I'd just found out what death meant, and perhaps I wanted to feel for its borders. She tells me that she understands, but how can she?

'Do you love me?' she asks, losing her temper. I don't know how to answer. Is it possible? Or inevitable? She laughs in a way that's a little bit mean, and says: 'Why do you always ask unanswerable questions?'

Lady Lovelace's objection

The genealogy of the almost-human bot is laden with folklore.

In 1948, Alan Turing devised an experiment—later known as the Turing test—to establish a

point at which a computer could be said to be intelligent. The 'pass' point, the realm of AI, was when the computer could be mistaken as human in conversation. Turing believed that we are all so predictable and consistent, that it would be possible for a series of algorithms to convincingly imitate us. All that would be required is a powerful enough computer to store and transmit such a huge amount of data.

Over half a century later, bot technology remains unpersuasive. Online chat-bots are clunky and strange, flipping between pleasantries and aggression at random. They frequently misunderstand their interlocutor, lose their place in the conversation, and are generally incoherent. At their worst, as Microsoft found with their Twitter bot 'Tay' earlier this year, they spew fascist, reprehensible bile.

But interest in AI technologies is growing, and the technology is steadily improving. The last year has seen an enormous rise in smart assistants for personal use, and it is estimated that 25% of us will be regularly using bots by 2019. In addition to the multitude of start-ups working on AI, Apple has just released Siri for third-party use, and at the end of 2015 Facebook launched M, a 'super-intelligent helper that's plugged in to all the information streams in the world'.¹

The alternative to work needn't be idleness

Thrown into relief against an increasingly precarious workforce, bots are now taking on the more quotidian aspects of human labour. Whether this is an act of liberation or brutality will depend almost entirely on the fate of capitalism. In the interim, we have the likes of Talla. Capable of fulfilling a company's HR responsibilities, Talla will guide new employees through their first few weeks. The start-up behind Talla claims that she is 'so smart, you'll forget she isn't human'. Another bot, GoButler,

She's an algorithm with opinions. Her hopes and anxieties are transmitted down plastic-wrapped wires. She's a what without a who.

is—or was, at least—a personal assistant, capable of fulfilling almost any request sent via text message.² Users have described GoButler as ‘astonishingly convincing’. X.ai’s Amy, also an intelligent assistant, is so credible that ‘she’ will often receive gifts of flowers and chocolate. AI technology is praised for its ability to fool others, as if there were something inherently futuristic about deceit.

Smart assistants can also augment our personal communications. Services like Crystal, which describes itself as ‘the world’s largest personality platform’, will create emails tailored to the personality of your online interlocutor. It aggregates tweets, blog posts, and LinkedIn bios to create a mirror image of a person’s communication style, and can generate text accordingly. I first heard of Crystal during the week of the US election, and almost instinctively I found myself on Hillary Clinton’s profile. ‘Write an email,’ Crystal suggested. ‘Sure,’ I thought. And within seconds Crystal had generated a multi-paragraph email mirroring Clinton’s own direct mode of communication. The email began with a problem to solve, as is her preference. As I edited the text I was reminded to take care with the passive voice and avoid over-punctuating. This is not how I would usually write: my emails are often vague and self-conscious and peppered with exclamation marks. Guided by Crystal’s algorithms, I was being slowly polished into a more Clinton-like version of myself.

In adding an AI sheen onto our own personal interactions, these communication bots—of which Crystal is perhaps the creepiest iteration—inhabit cyberspace as a hybrid of both it and you, or even communicating with others in your place. As with all moments of transgression, hesitancy surrounds this blurring of organic and synthetic. Aside from questions about authentic communication—for instance, where do I end and where does the algorithm begin?—some are concerned that AI will overshadow our fleshy, fallible selves altogether.

All watched over by machines of loving grace

Here’s one story that might be told from the future.

Back before the air was thick with asteroid dust and the sea-tides had swallowed up the land, the bots only did what we told them to do. They controlled the planet’s energy supplies, the distribution of food, and the shape of the cities. One day, bathed in a quantum sky, the bots surpassed all human intelligence. Capable of calculated judgement, but lacking the grey shades of human morality, the bots were unsympathetic to the humans, even malevolent.

Where do I end and where does the algorithm begin?

The humans had become trapped in a serfdom of their own creation. Civilisation fell to ruins, subsumed by the austere utility of an AI global system. Not that the humans weren’t warned, but the hubris of the time meant that these warnings fell on deaf ears. Elon Musk and Stephen Hawking denounced AI long before it became malevolent. They predicted its existential threat. To use Hawking’s words, ‘Humans, who are limited by slow biological evolution, couldn’t compete, and would be superseded.’³

But this is just a story—and not even a very original one. While the limits of the possible can often be exceeded, this post-apocalyptic vision is the stuff of bad sci-fi. For now, the key problem seems to be language, or rather a bot’s inability to process the nuances of language. Take this example, as Will Knight noted in *MIT Technology Review*:

iPhone user Siri, call me an ambulance.
Siri Okay, from now on I’ll call you ‘an ambulance’.

Or, this example, which computer scientist Yann LeCun has been using to test his own AI’s understanding:

Yann picked up the bottle and walked out of the room.⁵

For a human listener, it is obvious that Yann takes the bottle with him, but a bot sees the two elements of the sentence as separate: he picks up the bottle; he leaves the room. These semantic slip-ups demonstrate the way that bots are simultaneously smart *and* devoid of common sense, in a way that no human could ever be. Conversation with a bot is fundamentally different to one with a human. They do not infer meaning through complete sentences, but instead rely on spotting patterns and keywords. If a request is linguistically ambiguous, unique, or in some way complex, the bot can only guess at what you mean. Most contemporary bots, including Siri, try to link what you have said to a pre-programmed set of categories that prompt scripted responses. If it doesn’t understand a word you have used, it will ignore it. A bot simply doesn’t, and cannot, process language in the way that a human brain does.

Screaming, dreaming, unbelieving

So why do we believe in the possibility of human-like AI at all? Why fear the bogeymen? In part, the problem remains with language. Since the 1940s, the human brain has been metaphorically compared to a computer, with memories analogous to hard drive storage, and neurons thought to be transferring ‘bytes’ of data between components. If our brains were indeed computers, then it seemed entirely

possible to make a computer that operated, and therefore communicated, like a human brain. Of course, as psychologist Robert Epstein makes clear, the brain is not a computer. Not even close, in fact. It is, he says, 'just another metaphor—a story we tell to make sense of something we don't actually understand.'⁶ While scientists may be able to create an elaborate system of transfer and storage, this is certainly not a brain, and probably not AI either. Scientists have found enough trouble in recreating a human hip joint, let alone something as sophisticated as a human brain.

Bots misunderstand; bots make mistakes; bots lack empathy. Because of this, many chatbot developers tend to add a safety net of real people behind the scenes. These are people pretending to be bots pretending to be people. The 'astonishingly convincing' GoButler service, for example, employed 120 human workers—all on hand to check for errors, edit responses and add touches of empathy. Meanwhile, one ex-employee at X.ai—who created the bot so human she would receive gifts—reported that staff checked almost all incoming emails.⁷ These human workers give the bots Turing-level plausibility as humans because that is exactly what they are.

This sort of AI sleight-of-hand is nothing new. In 1770, Wolfgang von Kempelen created the Mechanical Turk, an automaton chess game. The machine—a wooden robot-like mannequin dressed in stereotypically 'Oriental' robes—appeared to play well against human opponents, even defeating the likes of Napoleon Bonaparte, Catherine the Great, and Benjamin Franklin. In reality the Mechanical Turk was a hoax: a chess grandmaster was concealed within the wooden cabinet below, and controlled the pieces with an elaborate set of levers, magnets and strings.

This is the strange shape of today's AI: in order to fool people into thinking a bot is human, it is necessary to fool people into thinking a human is really a bot. This, inevitably, has implications for labour practices: an employee is much easier to exploit when nobody knows they exist. GoButler's workers, for instance, reportedly worked in rotating shift patterns, switching from days to nights on alternate weeks. This kind of shift pattern is known to have a dramatic impact on both mental and physical health. Not only this, to ensure that the bot was always supervised, GoButler workers were required to eat lunch at their desks and attend the office Christmas party in 30 minute shifts. At X.ai, meanwhile, workers would frequently start work at 7am and leave at 9.30pm. 'I left feeling totally numb and absent of any sort of emotion,' said one

ex-employee.⁸ Given the expectation that bots will become more entangled in our everyday lives in the future, how many of us will find ourselves exhausted and broken after a long day of hidden emotional labour? The dystopian future of AI may be something much more banal than one led by malevolent overlords. It may be simply more of the same: a cheap and easily exploited workforce, perhaps located in the Global South, rendered invisible. Perhaps resistance may take the form of an inverted Turing test, in which we look for the people obscured beneath the algorithms. But as the number of chatbots is expected to soar over the coming years, it would be worth reflecting on what, or whom, we're really talking to.

Notes

1 Or so says Facebook CTO Mike Schroepfer.

2 Earlier this year GoButler scaled back their open-ended feature and replaced it with a less sophisticated model, beginning with its most commonly used category—flight searches.

3 Rory Cellan-Jones, 'Stephen Hawking warns artificial intelligence could end mankind', *BBC*, 2 December 2014, www.bbc.com/news/technology-30290540.

4 Will Knight, 'Tougher Turing Test Exposes Chatbots' Stupidity', *MIT Technology Review*, 14 July 2016, www.technologyreview.com/s/601897/tougher-turing-test-exposes-chatbots-stupidity.

5 Tom Simonite, 'Teaching Machines to Understand Us', *MIT Technology Review*, 6 August 2016, www.technologyreview.com/s/540001/teaching-machines-to-understand-us.

6 Robert Epstein, 'The Empty Brain', *Aeon*, 18 May 2016, www.aeon.co/essays/your-brain-does-not-process-information-and-it-is-not-a-computer.

7 Ellen Huet, 'The Humans Hiding Behind the Chatbots', *Bloomberg*, 18 April 2016, www.bloomberg.com/news/articles/2016-04-18/the-humans-hiding-behind-the-chatbots].

8 *ibid.*

Bots misunderstand; bots make mistakes; bots lack empathy.

In order to fool people into thinking a bot is human, it is necessary to fool people into thinking a human is really a bot.