


AI and the future of your mind

February 18 2020



Credit: CC0 Public Domain

Susan Schneider, associate professor of philosophy and cognitive science and director of the AI, Mind and Society (AIMS) Group at UConn, has gained a national and international reputation for her writing on the philosophical implications of artificial intelligence (AI). She writes about the nature of the self and mind, AI, cognitive science, and

astrobiology in publications including the *New York Times*, *Scientific American*, and *The Financial Times* and her work has been widely discussed in the media, such as *Science*, *Big Think*, *Nautilus*, *Discover* and *Smithsonian*. She was named NASA-Baruch Blumberg Chair for the Library of Congress and NASA and also holds the Distinguished Scholar Chair at the Library of Congress. In her new book, *Artificial You: AI and the Future of Your Mind* (Princeton University Press, 2019), she examines the implications of advances in artificial intelligence technology for the future of the human mind.

Q: What is the focus of your newest book?

A: This book is about the future of the mind. It explores the nature of the self and consciousness in a not so distant future, using today's work in artificial intelligence and brain enhancement technologies.

Consciousness is the felt quality to experience—what it feels like to be you. When you smell the aroma of your morning coffee, hear the sound of a Bach concerto, or feel pain, you are having conscious experience. Indeed, every moment of your waking life, and even when you dream, it feels like something from the inside to be you. This book asks: assuming we build highly sophisticated artificial intelligences at some point in the future, would they be conscious beings? Further, how would we detect consciousness in machines? These questions are addressed in the first half the book. The second half of the book is on the nature of the self. I illustrate that AI isn't just going to change the world around us. It's going to go inside the head, changing the human mind itself, but I'm concerned about the potential uses of invasive AI components inside of our heads. I urge that we need to understand deep philosophical questions about the self, consciousness, and the mind before we start playing with fire and start replacing parts of our brains with artificial components. When it comes to the self and mind, we are faced with vexing philosophical questions that have no easy solution.

Q: You report about such experimentation with neural implants for things like Alzheimer's disease but return to the question of, if there's an artificial intelligence when does it become aware of itself?

A: There are all kinds of impressive medical technologies underway, and I'm very supportive of the use of invasive brain chips to help individuals with radical memory loss or locked in syndrome, in which individuals entirely lose their ability to move. I think innovations to help these people are important and exciting. What I get worried about, though, is the idea that humans should engage in widespread and invasive AI-based enhancement of their brains. For instance, Elon Musk has recently declared that we will eventually need to keep up with super-intelligent AI—a hypothetical form of AI that vastly outsmarts us—and we need to do that by enhancing our brains. He also thinks doing so will help us keep up with technological unemployment that many economists claim will happen because AI will outmode us in the workforce. Musk and others talk about "merging with AI and I" through gradually augmenting intelligence with AI technology until, in the end of the day, we are essentially AIs ourselves. Musk has recently founded a company to do this, and Facebook and Kernal are also working on this. But I argue in the book and in op-eds for the New York Times and the Financial Times that the idea we could truly merge with artificial intelligence in the ways that a lot of tech gurus and transhumanists advocate is actually not philosophically well-founded. We have to think things through more carefully

Q: You use examples of AI from science fiction, including one with the Star Trek: Next Generation character Lt. Commander Data, who is under attack on a planet and he uploads his brain's memories to a

computer on the Enterprise. You ask: Will he still be the same Data that he was before being destroyed? Will he really survive?

A: I think people assume that AIs will have the capacity to be immortal because they can just keep uploading and downloading copies of themselves whenever they are in a jam. By this they mean the android be practically immortal, living until the end of the universe. This makes them almost God-like. I am skeptical. In the book I use the Data example to illustrate that if Data found out that he was on a planet that was about to be destroyed, he couldn't upload and genuinely survive. I think the idea that you could transfer your thoughts to a different format and still be you, surviving impending death, is conceptually flawed. It is flawed in both the human case and the case of androids. Believe it or not, there are advocates of uploading the human brain to survive death at places like the Oxford Future of Humanity Institute. I am skeptical.

Q: One of the points that you make in the book is that we have come far technologically but haven't heard anything yet from an alien culture. You suggest we should prepare for alien contact by including the involvement of sociologists and anthropologists and philosophers.

As the NASA chair at NASA and the Library of Congress, I love to think about the Fermi paradox, which is the question: Given the vast size of the universe, where is all the intelligent life? Where is everybody? Nowadays, the question can be framed in terms of all of the intriguing exoplanet research that identifies habitable planets throughout the universe, but are these exoplanets actually inhabited (not just inhabitable), and if they are inhabited, does life survive into

technological majority? Or are we alone? Why haven't we heard anything? To the extent that we even do find life out there, my guess is that we will first find microbial life. There's dozens of gloriously fun answers to the Fermi paradox.

Q: In the work that you're doing with Congress, what kinds of questions are you being asked and what we should be thinking about going forward with all this technology?

There's been a lot of concern over the last few years about deep fake videos. Nobody likes it; your career could be ruined by a deep fake video that has you saying something really rotten that you never said. Algorithmic discrimination is a big issue, the fact that algorithms that are based on deep learning technologies will be data-driven, so if the data itself has implicit bias, hidden biases in it, it can actually lead to a bad result which discriminates against certain groups. There are many members of Congress who've been concerned about that. That's why we really need AI regulations. AI regulation could do tremendous work. And so I do hope we move forward on all of these issues.

Provided by University of Connecticut

Citation: AI and the future of your mind (2020, February 18) retrieved 26 April 2024 from <https://phys.org/news/2020-02-ai-future-mind.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.