


Should you be worried about the rise of AI?

July 25 2017, by Ryan Nakashima And Matt O'brien



This combo of file images shows Facebook CEO Mark Zuckerberg, left, and Tesla and SpaceX CEO Elon Musk. An online smackdown between tech titans Zuckerberg and Musk over the possible threat of artificial intelligence underlines how little most people know about the rapidly advancing technology. (AP Photo/Manu Fernandez, Stephan Savoia)

Tech titans Mark Zuckerberg and Elon Musk recently slugged it out online over the possible threat artificial intelligence might one day pose to the human race, although you could be forgiven if you don't see why

this seems like a pressing question.

Thanks to AI, computers are learning to do a variety of tasks that have long eluded them—everything from driving cars to detecting cancerous skin lesions to writing news stories . But Musk, the founder of Tesla Motors and SpaceX, worries that AI systems could soon surpass humans, potentially leading to our deliberate (or inadvertent) extinction.

Two weeks ago, Musk warned U.S. governors to get educated and start considering ways to regulate AI in order to ward off the threat. "Once there is awareness, people will be extremely afraid," he said at the time.

Zuckerberg, the founder and CEO of Facebook, took exception. In a Facebook Live feed recorded Saturday in front of his barbecue smoker, Zuckerberg hit back at Musk, saying people who "drum up these doomsday scenarios" are "pretty irresponsible." On Tuesday, Musk slammed back on Twitter , writing that "I've talked to Mark about this. His understanding of the subject is limited."

Here's a look at what's behind this high-tech flare-up—and what you should and shouldn't be worried about.

WHAT IS AI, ANYWAY?

Back in 1956, scholars gathered at Dartmouth College to begin considering how to build computers that could improve themselves and take on problems that only humans could handle . That's still a workable definition of [artificial intelligence](#).

An initial burst of enthusiasm at the time, however, devolved into an "AI winter" lasting many decades as early efforts largely failed to create

machines that could think and learn—or even listen, see or speak.

That started changing five years ago. In 2012, a team led by Geoffrey Hinton at the University of Toronto proved that a [system](#) using a brain-like neural network could "learn" to recognize images. That same year, a team at Google led by Andrew Ng taught a computer system to recognize cats in YouTube videos—without ever being taught what a cat was.

Since then, computers have made enormous strides in vision, speech and complex game analysis. One AI system recently beat the world's top player of the ancient board game Go.

HERE COMES TERMINATOR'S SKYNET ... MAYBE

For a computer to become a "general purpose" AI system, it would need to do more than just one simple task like drive, pick up objects, or predict crop yields. Those are the sorts of tasks to which AI systems are largely limited today.

But they might not be hobbled for too long. According to Stuart Russell, a computer scientist at the University of California at Berkeley, AI systems may reach a turning point when they gain the ability to understand language at the level of a college student. That, he said, is "pretty likely to happen within the next decade."

While that on its own won't produce a robot overlord, it does mean that AI systems could read "everything the human race has ever written in every language," Russell said. That alone would provide them with far more knowledge than any individual human.

The question then is what happens next. One set of futurists believe that such machines could continue learning and expanding their power at an exponential rate, far outstripping humanity in short order. Some dub that potential event a "singularity," a term connoting change far beyond the ability of humans to grasp.

NEAR-TERM CONCERNS

No one knows if the singularity is simply science fiction or not. In the meantime, however, the rise of AI offers plenty of other issues to deal with.

AI-driven automation is leading to a resurgence of U.S. manufacturing—but not manufacturing jobs. Self-driving vehicles being tested now could ultimately displace many of the almost 4 million professional truck, bus and cab drivers now working in the U.S.

Human biases can also creep into AI systems. A chatbot released by Microsoft called Tay began tweeting offensive and racist remarks after online trolls baited it with what the company called "inappropriate" comments.

Harvard University professor Latanya Sweeney found that searching in Google for names associated with black people more often brought up ads suggesting a criminal arrest. Examples of image-recognition bias abound.

"AI is being created by a very elite few, and they have a particular way of thinking that's not necessarily reflective of society as a whole," says Mariya Yao, chief technology officer of AI consultancy TopBots.

MITIGATING HARM FROM AI

In his speech to the governors, Musk urged governors to be proactive, rather than reactive, in regulating AI, although he didn't offer many specifics. And when a conservative Republican governor challenged him on the value of regulation, Musk retreated and said he was mostly asking for government to gain more "insight" into potential issues presented by AI.

Of course, the prosaic use of AI will almost certainly challenge existing legal norms and regulations. When a self-driving car causes a fatal accident, or an AI-driven medical system provides an incorrect medical diagnosis, society will need rules in place for determining legal responsibility and liability.

With such immediate challenges ahead, worrying about superintelligent computers "would be a tragic waste of time," said Andrew Moore, dean of the computer science school at Carnegie Mellon University.

That's because machines aren't now capable of thinking out of the box in ways they weren't programmed for, he said. "That is something which no one in the field of AI has got any idea about."

© 2017 The Associated Press. All rights reserved.

Citation: Should you be worried about the rise of AI? (2017, July 25) retrieved 10 April 2024 from <https://techxplore.com/news/2017-07-ai.html>

| |
|--|
| <p>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.</p> |
|--|