


Is AI a Threat to Humanity?

By: Greg Scoble
December 30, 2014
CNN.com



Editor's note: Greg Scoble is the technology editor of PDN Magazine. Follow him on @GregScoble. The views expressed are his own. For more on the future of technology, watch the upcoming GPS "Moonshots" special on December 28 at 10 a.m. and 1 p.m. ET.

(CNN) -- Imagine you're the kind of person who worries about a future when robots become smart enough to threaten the very existence of the human race. For years, you've been dismissed as a crackpot, consigned to the same category of people who see Elvis lurking in their waffles.

In 2014, you found yourself in good company.

This year, arguably the world's greatest living scientific mind, Stephen Hawking, and its leading techno-industrialist, Elon Musk, voiced their fears about the potentially lethal rise of artificial intelligence. They were joined by philosophers, physicists and computer scientists, all of whom spoke out about the serious risks posed by the development of greater-than-human machine intelligence.

In a [widely cited op-ed](#) co-written with MIT physicist Max Tegmark, Nobel laureate Frank Wilczek and computer scientist Stuart Russell, Hawking sounded the AI alarm. "One can imagine (AI) outsmarting financial markets, out-inventing human researchers, out-manipulating human leaders, and developing weapons we cannot even understand. Whereas the short-term impact of AI depends on who controls it, the long-term impact depends on whether it can be controlled at all."

Musk [was reportedly more emphatic](#), expanding on his [tweeted warnings](#) by calling AI humanity's biggest "existential risk" and likening it to "summoning the demon." The debate over AI was given a big boost this year by the publication of philosopher Nick Bostrom's "[Superintelligence: Paths, Dangers, Strategies](#)," which makes a close study of just why and how AI may be so catastrophically dangerous (2013's "Our Final Invention" by documentarian James Barrat makes a similar case).

Bostrom is the director of the Future of Humanity Institute at Oxford, one of several new institutions devoted to studying existential threats to the human race, of which AI figures centrally. In May, the Massachusetts Institute of Technology christened its own [Future of Life Institute](#). In the academic community at least, AI anxiety is booming.

They're right to be worried.

The first and most immediate issue is the potential for AI to put large numbers of humans out of work. [A study by Carl Frey and Michael Osborne](#) of Oxford's Program on the Impacts of Future Technology put the matter starkly. In their analysis of over 700 jobs, almost half could be done by a computer in the future. This wave of computerization could destroy not simply low-wage, low-skill jobs (though those are in acute danger) but some white-collar and service sector jobs previously thought to be immune as well. Technology is marching on both our manual and mental labor.

As serious a threat as widespread job loss is, we've seen this movie before. During past technological upheavals, humans have cleverly created jobs and industries from the ashes of obsolete ones. We may be able to keep our collective heads above water even if AI encroaches on more creative and intellectual industries (heck, we may even start working less). What we should be more concerned about is humanity losing its perch as the Earth's foremost intelligence.

For those anxious about AI, current efforts to develop self-correcting algorithms ("machine learning"), coupled with the relentless growth in computer power and the increasing ubiquity of sensors collecting all manner of intelligence and information around the world, will push AI to human and ultimately superhuman intelligence. It's an event that's been dubbed "[the intelligence explosion](#)," a term invoked in 1965 by computer scientist Irving John Good in a paper outlining the development path for artificial intelligence.

What makes an intelligence explosion so worrisome is that intelligence is not a tool or a technology. We may think of AI as something that we use, like a hammer or corkscrew, but that's fundamentally the wrong way to think about it. Sufficiently advanced intelligence, like ours, is a creative force. The more powerful it is, the more it can reshape the world around it.

Artificial intelligence does not need to be malevolent to be catastrophically dangerous to humanity. When computer scientists talk about the possible threat to humanity from superintelligent AI, they don't mean the Terminator or Matrix.

Instead, it's typically a more prosaic end: humanity wiped out because an AI tasked with a simple goal (say, creating paper clips,[an example that is often used](#)) requisitions all the energy and raw materials on Earth to relentlessly churn out paper clips, outsmarting and outmaneuvering all human attempts to stop it. In Hollywood's telling, there are always humans left to fight back, but such an outcome is implausible if humanity is faced with a truly superior intelligence. It would be like mice attempting to outwit a human (we're the mice). In that event, AI researchers like Keefe Roedersheimer [see a less inspiring finale](#): "All the people are dead."

Needless to say, not everyone shares this bleak forecast. In the AI optimist camp, futurist and Google director of engineering [Ray Kurzweil](#) also sees intelligent machines precipitating human extinction, of a sort, only in Kurzweil's telling, humanity is not exterminated but subsumed into a superintelligent machine. Kurzweil's human-machine symbiosis is not a techno-catastrophe but the ultimate liberation from humanity's biological frailties.

[Others are skeptical](#) that AI will ever reach human levels of intelligence and cognition, let alone surpass it. Some, like New York University's Gary Marcus, are on the fence. "I don't know of any proof that we should be worried," [Marcus told me this year](#), "but nor of any proof that we should not be worried."

Irving John Good [famously described](#) the development of an ultraintelligent machine as the "last invention that man need ever make," for after that, humanity would cede innovation and technological development to its smarter progeny. Even if it's not a straight line from Siri to extinction, we humans should probably be watching our machines just a bit more closely.