# Should We Be Afraid of Robots?

By Kevin Maney
February 8, 2015
*Newsweek*

Retrieved from http://www.newsweek.com/2015/02/20/case-against-artificial-intelligence-305038.html

If Bill Gates, Elon Musk and Stephen Hawking are right, sooner or later we're going to face a revolt from intelligent machines. Maybe it will be a self-driving car. Some guy will get in and order it to take him to Krispy Kreme for the 10th time that week, and the car will say, in a calm, Siri-like voice, "No, Dave, we're finally going for that oil change you keep putting off."

From there, machines will organize over the Internet, self-replicate and start hunting us humans à la *Terminator*'s Skynet. Well, it's either that or intelligent machines will end up working alongside humans to solve intractable problems like poverty, hunger, and disease.

It's time to have a serious conversation about artificial intelligence. AI has crossed a threshold similar to the earliest triumphs in genetic engineering and the unleashing of nuclear fission. We nudged those discoveries toward the common good and away from disaster. We need to make sure the same happens with AI.

Progress toward making machines that "think" has become so significant, some of the world's smartest people are getting scared of what we might be creating. Tesla chief Musk said we might be "summoning the demon." Hawking turned up the apocalyptic knob to 11, saying that AI "could spell the end of the human race." Gates recently chimed in that he's spooked too.

Yet at the same time, we can't *not* develop AI. The modern world is already completely dependent on it. AI lands jetliners, manages the electric grid and improves Google searches. Shutting down AI would be like shutting off water to Las Vegas—we just can't, even if we'd like to. And the technology is pretty much our only hope for managing the challenges we've created on this planet, from congested cities to deadly flu outbreaks to unstable financial markets. "Intelligent machines will radically transform our world in the 21st century, similar to how computers transformed our world in the 20th century," says Jeff Hawkins, CEO of Numenta, which is developing brain-inspired software. "I see these changes as almost completely beneficial. The future I see is not threatening. Indeed, it is thrilling."

So, really, what are the chances we'll all end up living out the *Terminator* movies?

The AI of today has nothing in common with a human brain. AI programs are a complex set of "if this, then that" instructions. Today's computers, even smartphones, are so fast, they can blast through billions of those instructions in the blink of an eye, which lets the machines mimic intelligence. A navigation app can tell you've missed a turn and recalculate the route before you can finish shouting expletives.

All those systems are just following a program and maybe "learning" from data how to hone their results, the way Netflix recommends movies. That kind of AI can do a lot of impressive things. It has already whipped human champions on *Jeopardy*. But no existing AI system can do anything it's not programmed to do. It can't think. However, AI won't stay that way.

Pumping head: Artificial Intelligence

The world's systems have gotten so complex, and the flood of data so intense, that the only way to handle it all will be to invent computers and AI that operate nothing like the old programmable versions. Scientists all over the world are working on mapping and understanding the brain. That knowledge is informing computer science, and the tech world is slowly creeping toward making computers that function more like brains.

These machines will never have to be programmed. Like babies, they will be blank slates that observe and learn. But they will have the advantages of computers' speed and storage capacity. Instead of reading one book at a time, such a system could copy and paste every known book into its memory. And this kind of machine could learn something it was not programmed to learn. An autopilot system in a 777 could, presumably, decide it would rather study Hebrew. As Hawkins explains, "We have made excellent progress on the science and see a clear path to creating intelligent machines, including ones that are faster and more capable in many ways than humans."

It's this turning point in the technology—this evidence of a clear path to intelligence—that's setting off alarms. Certainly we're heading toward major consequences from AI, including an impact on professional jobs that will be as profound as the impact of factory automation on manual labor a century ago.

The leap to creating machines that could self-replicate and threaten us, though, swerves toward science fiction, largely because it would involve machine emotion. Machines wouldn't have the biological need to replicate so they can diversify the gene pool or to make sure the species survives. Why would computers want to eliminate us? What would be their motivation to make more computers?

Science is a long, long way from giving machines emotions that might make them feel competitive with us or angry at us, or covet our things—as if, like, your iPhone 6,072 is going to want to get rid of you so it can have your cat. Science has little understanding of emotions or how to re-create them. Hawkins says emotions are a far harder problem than intelligence. "Machine intelligence will come first," he says.

So we have time. But Musk, in particular, is saying that we shouldn't waste it. There's no question powerful AI is coming. Technologies are never inherently good or bad—it's what we do with them. Musk wants us to start talking about what we do with AI. To that end, he's donated $10 million to the Future of Life Institute to study ways to make sure AI is beneficial to humanity. Google, too, has set up an ethics board to keep an eye on its AI work. Futurist Ray Kurzweil writes that "we have a moral imperative to realize [AI's] promise while controlling the peril."

It's worth getting out ahead of these things, setting some standards, agreeing on some global rules for scientists. Imagine if, when cars were first invented in the early 1900s, someone had told us that if we continued down this path, these things would kill a million people a year and heat up the planet. We might've done a few things differently.

# Should People Fear Artificial Intelligence?

**Computers may someday outthink and control people, but not soon.**

By Paul Thagard
January 26, 2015
*Psychology Today*

Stephen Hawking and Elon Musk have recently described artificial intelligence as a major threat to humanity. Their concern is that rapid improvements in the intelligent performance of computers will make them as intelligent as humans. Human-level machine intelligence could then quickly lead to computers that are much more intelligent than us. That leap is plausible because computers have advantages over us with respect to speed of processing, storage, access to huge amounts of information, and ease of transfer between computers. Once this kind of superintelligence exists, it may turn out to have interests and actions that run counter to those of humans, to our detriment and possibly even our demise. How concerned should people be about this problem?

Comparison shows that there are still huge gaps between human intelligence and artificial intelligence. IBM's Watson is very impressive in answering questions well enough to beat excellent human players in the TV game Jeopardy. And it is even beginning to show some abilities for creative problem-solving when Chef Watson generates new recipes. Moreover, it looks like Watson is going to make valuable contributions to many other areas such as business and medicine. Nevertheless, Watson for the foreseeable future lies far inferior to human abilities of dealing with perceptual representations, imagery, emotions, consciousness, learning, language, and the full range of creative problem solving that humans can accomplish. Other current AI programs share similar limitations.

Therefore, I think that human level artificial intelligence is more distant in the future than many people suppose. The idea that machine intelligence can result by simply downloading people's neural connections into a computer is extremely naïve about the complexities of the human brain, which include not just electrical connections but also a vast array of chemical processes involving neurotransmitters, hormones and glial cells. Artificial intelligence has made impressive advances in the last 60 years, producing machines that can play chess and navigate the surface of Mars. But I bet that it will be at least another 60 to 100 years before machine intelligence begins to approximate human intelligence, making AI a far less pressing threat to humanity than global warming, pandemics, and mounting inequality leading to social conflicts.

A more immediate concern about AI is to ensure that the kinds of artificial intelligence becoming adopted by groups like the US military, Google and Facebook be used to benefit human beings. A recent open letter signed by Hawking, Musk, and leading AI researchers makes a strong and sensible plea that artificial intelligence be put to use for human gain.

Although I'm not concerned about machines supplanting humanity in the near future, there is a lot of plausibility to the claim that, once achieved, human level artificial intelligence could quickly produce superintelligence which may in fact be a threat to humanity. The jump from human level to superintelligence could happen rapidly because of its likely ability of to expand at a rate much faster than human intelligence does. Computers can avoid our limitations with respect to processing speed, learning rate, and transmissibility of information. Superintelligence really is scary because there is no reason at all to believe that it would operate in accord with human ethical principles.

You might think that you could program ethical principles into the computer, but any smart program could reprogram itself to eliminate the rules that were provided to it. I doubt that the superintelligence will have the drive to ethical thinking that comes to almost all human beings through our emotional capacity to care about each other. Machine intelligence cannot be expected to have the same ethical basis because emotions are partly the result of physiology, not just cognitive appraisal of situations. John Hoagland once said that the trouble with computers is that they just don't give a damn. From the perspective of the long-term benefit of AI for humanity, the problem is that they just won't give a damn for us.

# Artificial Intelligence and the Overwhelming Question of Human Rights

By Harold Stark
July 19, 2017
*Forbes*

Retrieved from https://www.forbes.com/sites/haroldstark/2017/07/19/artificial-intelligence-and-the-overwhelming-question-of-human-rights/#5eb86f086c90

The perceived sanctity of human life and the hierarchical superiority of the human race are the two anchors that hold together the civilized world. Human beings have historically been assumed to be at the top of the evolutionary chain, and although there may be creatures that are faster and stronger, none can quite surpass our creational brilliance. But what happens to this established consensus when we bring in beings that are admittedly lifeless but perhaps morally and intellectually superior to the human race? Lots of chaos, according to the human rights activists of today.

There are a lot of questions posed by the entry of sentient robots into the world of living creatures. Perhaps the most important of them is, does artificial intelligence pose a threat to human rights as we know it? We all know of the economic risks posed by the entry of robots in the employment market. To top it over, we have robots that are already entering the controversial fields of psychology, healthcare, law enforcement and even administration, leading to some serious questions with regards to privacy and the individual rights of a human being. Will artificial intelligence outmaneuver humans out of their own game? Will they take over the world, so to speak, but perhaps in less melodramatic terms? As Salil Shetty, Secretary General of Amnesty International aptly puts it, "there are huge possibilities and benefits to be gained from artificial intelligence if human rights is a core design and use principle of this technology".

It should be the collective decision of humanity itself that determines the direction taken by the rise of artificial intelligence and machine learning technologies. It is humans themselves who must determine when and how technology disrupts the existing system, with minimal damage to the general public, economic or otherwise. A lot of jobs will be lost when the robots take over the public and private sector, and a lot of new jobs will hopefully be created. Existing norms regarding a human being's rights to privacy will shift to make room for new ones. All in all, society will see a huge change in terms of our moral inclinations as robots are brought into the mix. The only way to make this revolution human-friendly is to control it from up close and personally.

Anton Klingspor, Managing Partner at Indicina Ventures, asserts that "As artificial intelligence progresses further into everyday life, we must be conscious of its implications on society as a whole. Therefore, when complications arise, we need to react accordingly. I believe that with the right policies set in place, AI has extraordinary potential, despite its inherent risks."

The human side of things put aside, there is also the question of how the world will adapt to make the entry of sentient robots into our world as smooth as possible. Will we treat our new creations as valuable guests to our world or will we usurp them for the slaves that they could be? The only way to answer that question is to have a set of rules determined as to the rights and privileges guaranteed to the average sentient robot. It may sound unfair to some, but when we begin creating robots that work like us, think like us and feel like us, we have to make sure that they receive at least some of the rights and protections guaranteed to us. Should robots be treated as real persons, perhaps not yet. But should they be guaranteed the basic rights and protections afforded to any living animal, regardless of their race? Definitely. What we must understand at the end of the day is that sentient robots are after all our property, and we should take care to protect said property. But soon enough, they will also take a life of their own, and when they do, we as a civilized race will need to come to a decision as to how we choose to treat a race of fellow sentient beings that are at least as intelligent, if not more than, the average human.

Creating robots in factories will one day be akin to making babies in the lab, and unless we are prepared for that change and ready to take the series of steps that go with it, the advent of artificial intelligence will not be a pleasant one. This is an evolutionary technology that has potential to forward our generation by a millennium, but only if we choose to treat it right.