Artificial intelligence: Should we worry?

The epic science fiction film *2001: A Space Odyssey* features a sentient computer, *Hal 9000*, whose job is to control a ship carrying astronauts to Jupiter, but a malfunction turns the computer against the crew. That film came out in 1968, around the same time when comedian Woody Allen was performing a stand-up routine, *Mechanical Objects.* His routine was about a talking, apparently sentient, elevator bullying him, but, unlike Hal's malfunction, the reason for the elevator's behavior was made clear: Allen had hit a television.

The late 1960s was also the time period when NASA was sending astronauts to the Moon with significantly less computing power than the smart phones that people use today for broadcasting cat videos. And so, whether a comedy scenario or a dramatic sci-fi thriller was on their mind, nobody thought the prospect of an artificial intelligence (AI) making trouble for humans would be a real-world issue so soon. Maybe, hundreds of years in the future, as imagined on *Star Trek*, but in the early 21st century? That seemed farfetched in the *Apollo* era.

But is it still ridiculous with how far technology has come and the rate at which it is advancing? That really depends on whom you ask.

Some leading experts on the consequences of science and innovation—Stephen Hawking (astrophysicist), Bill Gates (founder of Microsoft, philanthropist), Elon Musk (engineer, innovator, space entrepreneur), and Sam Harris (author, neuroscientist, philosopher), to name a few—are seriously worried about the prospect of artificial general intelligence (AGI). AGI is achieved when a machine mind ponders and cares about numerous topics at a level similar to our own brains. Others—author and Arizona State physicist Lawrence Krauss for instance—think the worriers should chill out, that we should be careful about how we integrate increasingly intelligent machines into military command and control, but that it will be a century or more before a computer mind could be created at, or beyond, the level of a human brain.

But even if AGI is far off, more specialized forms of machine intelligence are already operating and will affect us very soon in profound ways. Certain cars now feature driver-assist and collision avoidance systems, including some that break automatically for pedestrians if the driver fails to do so. How long will it be until a car can take other evasive measures, such as veering onto the sidewalk, if there's not enough time to stop before hitting the pedestrian? Probably sooner than you think. People are working on this, because the demand for increased traffic safety and the economic factors are strong drivers, and we'd all agree that they should indeed work on it.

Now, imagine that a great-grandmother is in the road and a baby stroller on the sidewalk. What should the accident avoidance algorithm do? Who should live? If there are 10 people in the road and just one on the sidewalk, what then? Does the needs of the many outweigh the needs of the few, or the one? The AI systems, now under development for cars are specialized, not generally intelligent like the human brain, nor are they sentient. But in the years to come, if they can react to the above scenarios, they will be making *moral* decisions. Just like Hal.

Engineers on the spectrum

Keep in mind that chill-outers like Krauss are not saying that we'll never have to worry about AGI. Krauss is just saying that it's not likely something we'll have to worry about anytime soon. To support a synthetic, sentient brain built on an electronic computing principle, Krauss has calculated that we'll need a huge amount of power, 10 Terawatts, a little more than half of humanity's total power use at present. Thus, he thinks that we'd need a qualitative change in computing technology, something that emulated the cellular and tissue architecture of our biological brains, but such an artificial brain would not have the talents that are particular to electronic computers: the memory and instant recall capacity, the computational ability. You could merge the two types of computing technologies to have an artificial mind that is both generally intelligent like a human (possessing AGI) and computes like a computer, but the incentive to develop the AGI component is questionable. Why do that if you can keep improving specialized AI—the kind needed to drive cars, fly planes, perform robot surgery, or whatever you want it to do.

Can the specialized AI approach lead to AGI, and thereby a scenario where machines start behaving in ways that are not in our interest? Stephen Hawking and Bill Gates think we should be careful. Elon Musk thinks AI has the potential to become more dangerous than nuclear weapons. That may sound hyperbolic but considering the source—Musk is one of the greatest inventors and forward thinkers of our time— if he's worried, it's worthwhile to listen.

Sam Harris is one of the people saying that we ought to be very careful. He's cautioned against AGI, but even if he's wrong about the dangers of AGI, some of his arguments could be applicable to out of control AI. In particular, Harris suggests that the worst nightmare scenario is that a super-intelligence develops that is NOT sentient. It's just effectively omniscient and omnipotent and thus starts doing a lot of things that it decides should be done, based on its own calculations. You could not reason with such a "mind", because it's not sentient, conscious, or self-aware.

That's the type of AI that could develop based on incremental progress in electronic computing, the kind of progress the accident avoidance systems in cars. And who is designing those programs? Curious about the issue, Harris attended a major AI conference and noticed that for the most part it's computer engineering whizzes who tend to be on the spectrum –meaning the autism spectrum (he did not mean it as hyperbole)– and who spent most of their time sitting around working in front of a screen drinking Red Bull.

Ponder that a little bit to balance out the chilling out message.

David Warmflash is an astrobiologist, physician and science writer. Follow @CosmicEvolution to read what he is saying on Twitter.