


Friend or foe: five questions outlining the future of artificial intelligence



Is artificial intelligence going to usher in a golden age of humanity by solving our most intractable problems and creating vast amounts of wealth? Or is it the ultimate menace, a threat to our jobs and, ominously, our very survival as a species? It depends on whom you ask... and what they happen to believe about technology and humanity.

Few, if any, technologies elicit such a wide range of predictions as artificial intelligence. [Elon Musk](#), [Stephen Hawking](#) and [Bill Gates](#) fear AI, a sentiment that [Hawking](#) sums up by saying, "(...) creating AI would be the biggest event in human history... Unfortunately, it might also be the last..."

Meanwhile, an equally illustrious group of technologists including [Mark Zuckerberg](#), [Andrew Ng](#) and [Oren Etzioni](#) finds this viewpoint so far-fetched as to be hardly worth a rebuttal. Zuckerberg goes so far as to call those who peddle doomsday scenarios "[pretty irresponsible](#)", while Ng says that such concerns about AI are like worrying about "[overpopulation on Mars](#)."

The tasks AI currently performs are mundane in the extreme. It filters e-mail spam, predicts the weather and routes drivers around traffic. However, the type of AI that has everyone in a tizzy is quite different from what we have today. It is a technology called AGI, or artificial general intelligence, which is an AI as smart and flexible as a human, capable of solving problems it hasn't seen before. It is this form of AI that elicits so much fear and hope.

Why is there such disagreement about this technology? How do the two sides see the world so differently? It boils down to five key issues:

How far are we from an unsupervised learner?

Right now, the kind of narrow AI we have must be manually trained by humans using carefully prepared and curated data sets. To teach an AI to identify photos of cats, for instance, you need a large database of photos, each of which has been labeled by a human as a cat or non-cat. And even when you have this, and the AI can spot cats as accurately as any human, the AI can only recognise cats. To teach it to identify dogs requires starting over from scratch.

How quickly would an AGI advance?

Many of those who are concerned that an AGI is coming soon believe we will build an unsupervised learner that will improve itself recursively at a speed that we can neither control nor even understand. Day by day, then hour by hour, then minute by minute, the AGI would get smarter. It would improve itself at the speed of light while humans can advance only at the speed of life, meaning that before long, it would become a superintelligence, as far ahead of us as we are ahead of ants.

Those on the other side of the debate maintain that this is little more than a tired science fiction trope and point out that while the AGI may improve on its own in certain narrow domains, the idea that it will spiral into omniscience is to misunderstand what intelligence is to begin with.

How complex is intelligence?

That brings us to the third point of disagreement. Those who believe an AGI is a looming threat tend to think that intelligence may be relatively straightforward and simple. This is not without precedent. Isaac Newton determined that three simple laws are all it takes to understand motion, and, likewise, a few straightforward laws have been found to govern electricity and magnetism. Is intelligence the same? Are we just one “eureka moment” away from understanding the basic laws of intelligence? Many of those who are apprehensive about AGI believe that this is the case and that with a single breakthrough, we will let the AGI genie out of the bottle.

Those on the other side maintain that the human brain’s hundreds of specialised functions suggest that our intelligence consists of countless clever evolutionary hacks, each of which will need to be sorted out and understood before we can teach it to machines.

How special are humans?

An AGI is an artificial intelligence that is smarter and more versatile than a human. Those who think AGI is coming soon regard this as a pretty low bar. They don’t see human abilities such as creativity as all that mysterious or difficult. Those on the other side look at the capabilities of humans, from intuition to emotion to consciousness itself, and are persuaded that we are far from understanding how we do those things, let alone building machines to do them.

Have we started building a true AI yet?

As stated earlier, the kind of AI we have today is called narrow AI because it can do just one specific thing. Is narrow AI the first step to building a generalised AGI? Many of those who are troubled about an AGI believe we are well on our way to piecing together one, using the same techniques that gave us narrow AI.

The other camp sees narrow AI as an entirely different technology from what an AGI will require. While they both happen to have the words “artificial” and “intelligence” in their names, that’s where the similarities end.

These are the five questions on which the future of AI hangs.

Interestingly, the debate around automation and jobs falls exactly along the same lines. The concerned-about-AGI camp believes that with AI, robots and automation will soon take all the jobs, making almost all humans unemployable. The not-concerned camp maintains that automation and AI won’t compete directly with humans and will instead boost human productivity, creating better and higher-paying jobs.

We are in uncharted waters, to be sure. One thing is certain: These sorts of questions are a harbinger of things to come. When we start grappling with digital life and machine consciousness, we will be on even shakier ground.



Notes:

- *This blog post is based on the author’s book “[The Fourth Age: Smart Robots, Conscious Computers & The Future of Humanity](#)”, Atria Books, 2018. Read an excerpt here on [The Case for \(and against\) an AGI](#).*
- *The post gives the views of its author(s), not the position of LSE Business Review or the London School of Economics.*
- *Featured image credit: [Image](#) by [Mike MacKenzie](#) via [www.vpnstrus.com](#), under a [CC-BY-2.0](#) licence*

-
- *When you leave a comment, you're agreeing to our [Comment Policy](#).*
-



Byron Reese is the CEO and publisher of Gigaom, and hosts the [Voices in AI](#) podcast. He has been building and running Internet and software companies for twenty years and has obtained or has pending patents in disciplines as varied as crowdsourcing, content creation and psychographics.