

AI Will NEVER Wipe Us Out, But It's ALREADY Dangerous

**Diversions in History / althist.xyz
released under the CC BY-SA 2.0 license
2018**

Each generation imagines itself to be more intelligent than the one that went before it, and wiser than the one that comes after it.

George Orwell

One common doomsday scenario, and science-fiction trope, consists of an artificial intelligence with greater-than-human capabilities taking over the world or destroying humanity. In recent years, we have received stern warnings on this possibility from figures such as Stephen Hawking, Sam Harris, Nick Bostrom, and Elon Musk. However, despite the weight that has been placed behind the issue, it is actually a very far-fetched prospect.

For such an advanced AI to be dangerous, first one has to be built, which is a very tall order. Then it has to want to take over and be able to do so, the likelihood of which, again, is currently zero because that kind of AI doesn't exist, but even if it did, it is totally unknown whether this is a plausible outcome. So far, superhuman AI has been one of those things that is continuously promised but never delivered, like flying cars or fusion power.

It's pretty pointless to worry about something that doesn't exist, may not be able to exist, and may be totally benign or even helpful if it ultimately does, when, in actuality, rising global temperatures might eventually cause agriculture to collapse, the world is running short on critical and irreplaceable resources, and there are thousands nuclear warheads ready to be launched (by humans).

It May be Impossible to Create a Human-level AI

It's not clear that the human brain could be duplicated with lines of code being executed on computer chips. One clue it might not be possible is that the human brain can't do a lot of what computers can, like nearly instantaneously solving complex math problems, which not only brings the contrast between human intelligence and computers into relief but also begs the question of why the same wouldn't hold true the other way around. At the very least, more exotic means like quantum computers may be needed to even approach the prospect of human-level or superhuman AI, but those are still in their earliest stages.

The world is probably further away from that kind of AI than many people think. Just because, with varying success, you can put Chinese into a computer and have it spit out English, or have a computer play chess or recognize the subject of a picture, does not mean real intelligence is just a few more steps away. There is no AI that can do anything except the particular task it was designed to do, which is the fundamental difference between humans and machines. It may get more sophisticated at its given task, but the barrier has never been shown to be crossable.

Basically, the human mind is more than the sum of its parts. So many factors shape it: genetics, language, environment, experiences, culture, hormones, and so on. It's not something that's just made. If a computer program is just a set of instructions, how can someone instruct a computer to be truly adaptable, to become self aware, or to feel things? Nothing a computer does is anything like intuition,

reasoning, or sentience. If the code is just running on the hardware, then it's just the code that's directing everything, no intelligence acting on its own will. And if it's really as simple as running the right code, all success would do is prove that's all intelligence is, and where would that leave us?

Human-level AI is Inconceivable

Humanity's bias toward its own evil is really where the whole scenario of an AI doomsday comes from. Just because humans are malevolent and often deliberately cause harm doesn't mean a superhuman AI would, which is the underlying assumption. If AI is superior to humans in its abilities, it seems more likely that it would be superior to humans in ethics too. Humans have evolved to out-compete and destroy rivals. An AI, on the other hand, would be created on a computer. How could it have anything like the same intentions or thought processes as humans? If it were so superior that it could just dispose of people, would it even notice them in the first place? If it were designed to think like humans, if humans had that level of control over the end product, there would be nothing to worry about anyway; they could also just design it not to harm them.

The issue is not worth worrying about because it's out of everyone's control anyway. If creating a superhuman AI just involves writing computer code, then potentially anyone anywhere could create it. Maybe some hobbyist will create the first one in his basement. It makes just as much sense to worry about someone figuring out how to build a doomsday machine, in a more traditional sense, in his basement. It would be really bad if it happened, but not knowing how exactly such a machine might be created, if it could be, there's really nothing anyone can do about it until it happens.

There's Nothing an AI Could Do

Consider that while humans are supposedly a lot more intelligent than all other animals, other animals usually have the better reflexes, often have better senses and perceptions, and manage to kill humans all the time. Bacteria and viruses have no intelligence, but they kill millions of humans. Considering that the human creators would have every advantage aside from computing power over the AI, it doesn't seem like there would be much it could do to harm them. An intelligence by itself can't do anything; it has to be able to act on the real world, and humans would have a pair of bolt cutters on the power cord.

People worry an AI, once it emerged in some test environment, might manipulate its creators into letting it out of the bottle so to speak. But consider the most persuasive humans you've ever heard. Is there anything they could say to convince you, for instance, to go walk off a cliff? In some cases, there is just nothing that could be said to trigger a certain behavior, such as unleashing a potentially hostile superintelligence.

Another idea is that an AI could manipulate its circuitry in a way that causes radio waves to be emitted, which it could form into signals that it could use to control external devices or even spread out to them, or that it could rapidly upgrade itself beyond human understanding, neither of which is very plausible. Humans are the only known model for this kind of intelligence, and they can't do anything remotely like this with their brains or bodies; humans can hardly control any of their bodily processes. And radio doesn't work that way anyway; it needs to have circuitry designed to create meaningful, deliberate emissions. Regardless, the room containing the computer running the AI could just be shielded.

At the same time warning us of near-supernatural powers, those sounding the alarm on AI also state it could wipe us out by accident. One common illustration is that if an AI were created to make paper clips

but without a broader perspective, it might wind up turning the entire earth into paperclips in full pursuit of its task. However, if the AI were advanced enough to manage that, it's hard to believe it would be stupid enough to.

A superhuman AI is actually not something to fear. If it did supplant us, that would just be evolution at work, as it has been all along, and another step in the ongoing history of life on Earth. It wouldn't mean a massive genocide either; the AI would just do better than humans and eventually rise up to be the dominant species. If the AI were so intelligent, it would also probably do a better job of running the earth than humans do; people mainly use it as an arena for wars while they transform its resources into landfill. It's also worth mentioning that it probably wouldn't even need a planet; computers can be made to function just as well in space. At the very least, it couldn't possibly treat humans any worse than they treat supposedly lesser creatures.

The idea of AI wiping out humanity is a fantasy and a distraction. If it's a possibility at all, superhuman AI is most likely many decades away at least, and there are many other things that are a lot more likely to destroy us before then. This too is pinned on AI; the world supposedly needs it in order to solve such major problems. But everyone already knows what the solutions are; people just don't want to implement them. With regard to global warming, all an AI would do is suggest things like phasing out cars or arresting oil company executives, and people would just ignore it like they do other humans saying those things already.

The Real Threat

The idea that AI will destroy the world is a smokescreen shrouding the real threat it poses: the destruction of people's livelihoods and privacy, which it's already achieving without being that intelligent. More and more jobs are being replaced with automation while leaving the same economic structure in place so that a small, rich minority gets all the benefits while the former workers are left destitute. Meanwhile, enormous databases are being built containing information about everywhere people go, everyone they know, everything they buy, what they read, and so on, compiled from sources like credit card records, license plate readers, social media, and Internet-of-Things devices. AI that already exists analyses and correlates all this data so that the corporations that can gather the most of it and produce the most sophisticated systems to make sense of it can make tremendous amounts of money. Probably most commonly, they can build profiles on individuals' traits and lifestyles and use them for highly targeted advertisements.

In one episode, Target predicted that a girl was pregnant based on an analysis of her seemingly unrelated purchases and used that information to mail her an ad stating as much, spilling the beans to her father in the process. Massive amounts of data combined with AI can also be used to try to manipulate people. In 2014, for example, Facebook caused a major scandal by attempting to affect users moods by altering the stories listed in their news feeds.

At the same time, AI is helping to create a new underclass. The major trend is that AI is able to take over more and more tasks formerly only humans could do. One effect is being able to analyze more data in new ways. Another is that machines can now to some extent do things like check people out at the supermarket, answer phone calls, drive cars, and not only dumbly dispense money but also deposit cash and checks. Where does that leave the humans who need to earn a living and who used to do those things?

AI concentrates power to a degree that is incompatible with freedom. A great example is how, right now, the three giant credit bureaus pull in data from all over the place to automatically generate a credit score for every individual, which in part determines whether someone can get a credit card, finance a car, buy a home, get an apartment, or carry out really any major transaction. No one elected the credit bureaus to this position, no one opted into this, they just have all the data and do as they please. Similarly, employers can run background checks on prospective employees and find out a tremendous amount about them, but there's no similar means for employees to find out how someone is to work for in advance. How much longer until someone becomes unemployable because a recruitment AI doesn't like their social media history? How long until health insurance companies are terminating people's policies based on the food they buy?

The companies that collect data can sell or otherwise disclose it largely at their own discretion, and there's really no way to know what happens to it then. The data can be stolen by criminals or abused by the state, which doesn't need warrants or subpoenas. It's not like coming into your home; the authorities just ask the company for it. If a more brutal state rises up, it already has all that information at its disposal to utterly crush the population. If you're looking for something to use against people, and you have enough data on them, you can always find something to twist around. Keeping some things private, on the other hand, is part of the definition of being an individual, but now almost everything anyone does is considered the business of some company.

To a large degree, we can still avoid the spying. We can pay for things with cash wherever possible or stay away from Facebook. But as more and more devices become "smart," that is, they spy, and technologies like facial recognition come into play, it really starts to get frightening. If an AI can identify someone just by seeing his face on the street, there's no more avoiding it. As devices like smartphones invade people's lives more and more, it might start to become impossible to accomplish everyday tasks without them, rather than just inconvenient.

Amazon, for example, has already started opening stores that have no cashiers or checkouts of any kind. Customers just go in, scanning their phones, and take whatever they want. Their account is able to be charged by a computer vision system and sensors tracking everything that happens in the store. Some people are raising concerns about the loss of jobs and privacy, but more are starry-eyed over this development as a technological feat. If nothing changes, eventually cash will be out, and all shopping will be completely monitored.

There's not much reason to hope considering how freely people are capitulating. They just shrug their shoulders and regurgitate propaganda, like that you shouldn't mind losing your privacy if you have nothing to hide. People give their phone number or e-mail address away at the cash register like a dog responding to its owner's command just to get a little discount. They install some stupid game on their smartphone even though it asks for access to the camera, microphone, sensors, and nearly every bit of data on the device. If this includes the contacts, the user is deciding that privacy is worthless for every person listed there too. Not long ago, such a program was called at best spyware and at worst a Trojan horse.

The Chinese are already going beyond the world George Orwell tried to warn us about. They are building a massive, nationwide surveillance system that collects and analyzes data from innumerable sources to assign citizens a score based on whether they are behaving in a way the state considers pro or anti-social. Those with low scores are penalized, such as through travel restrictions or being banned from certain jobs. Some workers are also being made to wear hats with sensors that monitor brain activity, and an AI

uses that data to decipher their emotional states. What goes on an individual's mind isn't even private anymore.

AI is definitely taking over, but it doesn't think for itself, and it's only doing the bidding of its masters. People of the future, whether biological or artificial, will look back on the ideas we have today about superhuman AI and probably think they are as ridiculous as many past predictions of the future seem now. But whether those people are at all free or having their behavior controlled by an absolute surveillance apparatus is a question the answer to which is right now being decided.