# QUARTZ



GETTY IMAGES

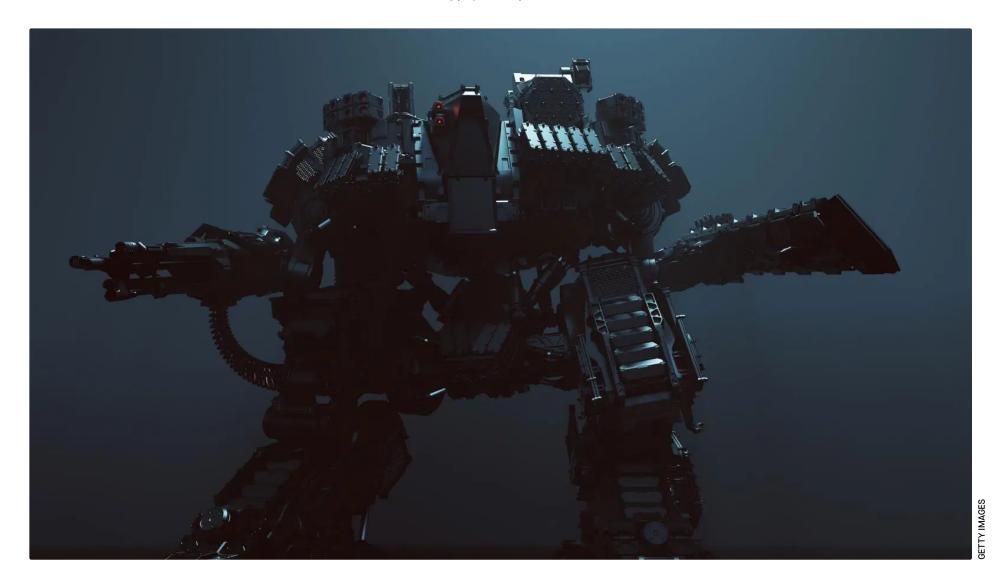**MAN VS. MACHINE**

# The robots are coming, the robots are coming

✦ Member exclusive by Justin Rohrlich for The future of war

The idea that technology will eventually surpass human intelligence and that humans will then lose control of it is called "the singularity." The term originates from a 1993 essay by science fiction writer Vernor Vinge titled, "The Coming Technological Singularity: How to Survive in the Post-Human Era."

It pretty effectively sums up the existential fear now taking hold as both militaries and non-state actors across the globe develop robotic weaponry that can think for itself.

"Within 30 years, we will have the technological means to create superhuman intelligence. Shortly after, the human era will be ended," is how the essay's abstract begins.

That was 26 years ago.

The specter of lethal autonomous weapons systems has generated worldwide debate over their merits and dangers. And that debate is growing more intense as technologies continue to progress virtually unencumbered by law or regulation.

Afraid of what's to come, activists are now pushing for a ban not just on the use of so-called "killer robots," but even on the research and development of them. Others think the hysteria is exaggerated, fueled by the work of science fiction writers like Vinge.

Utah state representative Adam Robertson, who also serves as chief technology officer for the artificial intelligence software developer Fortem Technologies, thinks everyone just needs to calm down.

"There's a lot of hype around autonomy and artificial intelligence right now, a lot of excitement around what happens when machines become sentient and take over humanity," Robertson told Quartz. "But the reality is that even our

most sophisticated systems are very low level. We've had billions of [dollars in] investments in the automobile industry just to try to get cars to drive themselves, and it's taking a lot longer than they thought."

We can't let ourselves get hung up on "scary" terminology, argues Lindsey Sheppard, an expert in the military use of emerging technologies at the nonprofit Center for Strategic and International Studies.

"There are lots of unresolved questions regarding appropriate or inappropriate roles for machines and the appropriate level of delegation and assessment and decision-making for things we traditionally think of as very human endeavors," Sheppard said. "I think there's a lot of uneasiness because those questions remain unanswered, and have remained unanswered for decades."

It might not matter which side of the debate you fall on, though. The development and use of fully autonomous weapons is already happening, and they'll be near impossible to stop or even contain.

## They're here

While the debate over these weapons centers mostly on what could be, Stuart Russell, the director of the Center for Intelligent Systems at the University of California-Berkeley, is more than a little unnerved by what already exists.

The US military is already using autonomous weapons systems in a defensive capacity, such as Lockheed Martin's Aegis Combat System, which without human instruction is supposed to detect and intercept incoming ballistic missiles. The US Army is developing an AI-capable cannon, which would select and engage targets on its own, as well

as AI-assisted tanks that, as Quartz first reported, will be able to "acquire, identify, and engage targets" at least three times faster than any human.

Russell notes with alarm that the Turkish military is set to deploy in Syria autonomous quadcopters with swarming technology (the ability for groups of unmanned vehicles to make their own decisions using information shared between them), facial recognition, and anti-personnel attack capabilities.

A highly regarded expert on the implications of AI, Russell believes the further pursuit of technology like this will lead to disaster. And because they do not require human supervision, lethal autonomous weapons are potentially scalable weapons of mass destruction. Virtually unlimited numbers of "killer robots" could be launched by a small number of people, something that isn't possible with conventional arms.

"I estimate, for example, that roughly one million lethal weapons can be carried in a single container truck or cargo aircraft, perhaps with only two or three human operators, rather than two or three million," Russell continued. "Such weapons would be able to hunt for and eliminate humans in towns and cities, even inside buildings. They would be cheap, effective, unattributable, and easily proliferated once the major powers initiate mass production and the weapons become available on the international arms market. As a result, we expect that autonomous weapons will reduce human security at the individual, local, national, and international levels."

## The human element

The Pentagon, for its part, is attempting to avoid these doomsday scenarios.

Last month, it announced its intention to bring a full-time ethicist on board at its new Joint Artificial Intelligence Center (JAIC), a multi-service command set up in 2018 to harness emerging AI capabilities for the US military. The command's director said at the time that an ethical component lies at the core of the Pentagon's overall AI strategy.

"There is no doubt that machines will advance in their ability to learn and make more sophisticated decisions," US Navy lieutenant commander Arlo Abrahamson, a spokesperson for the JAIC, told Quartz. "But the Department of Defense does not field technology for mission-critical applications before we have substantial evidence that it will work as intended."

To that end, Abrahamson said the Defense Department has spent the past several years experimenting with AI in non-lethal areas, like intelligence and humanitarian assistance. "The lower consequence applications have allowed us to learn valuable lessons about best practices in AI program management and collect operational data without disrupting existing operations and without putting lives at risk," he said.

But the most effective way to mitigate the dangers of deploying artificial intelligence in a conflict, activists say, is to always inject a human into the decision-making process. The use of unmanned drones by the US military and the CIA in places like Afghanistan, Pakistan, Yemen, Somalia, and elsewhere has allowed the US to more precisely target terrorist networks (with varying degrees of success). But a soldier is still behind the controls, even if not on the battlefield.

REUTERS/ANNEGRET HILSE

Activists from the Campaign to Stop Killer Robots protest in Berlin.

Paul Scharre, director of the technology and national security program at the bipartisan Center for a New American Security and the author of *Army of None: Autonomous Weapons and the Future of War*, said that in this scenario drone pilots can still see their targets, and the aftermath of their decisions to use force. That's a good thing, he says. Machines, no matter how intelligent, don't have the kind of moral judgment to determine things like proportionality when considering a response.

Abrahamson agrees that human judgment is "central" to the safe and lawful use of any weapons system. "A human will always be accountable, and our activities will be guided by our policies, such as Department of Defense Directive 3000.09, Commanders' Intent, and the Law of War," he said.

Issued in 2012, Directive 3000.09 requires that all "commanders and operators...exercise appropriate levels of human judgment over the use of force." However, while it lays out a series of guidelines that apply to AI, and specifically calls

for the use of human decision-making in combat, Scharre said it does not actually forbid fully autonomous lethal weapons.

"I certainly wouldn't describe it as a prohibition and the Defense Department would not either," Scharre said. "It outlines a review and approval processes that it would have to go through at various levels of authority, and then a set of procedures and questions that people would have to ask about anything that was new that would use autonomy in a novel way."

Ultimately, it leaves the door open for the development and use of fully autonomous weapons. And once the human element is removed, Scharre said, everything changes.

"What would be the effect if we fought a war and no one felt responsible for the killing that occurred?" Scharre asked. "I think that's an important question. There are some compelling reasons to say we would want people involved in lethal decisions, regardless of what the technology is capable of, for both legal and ethical reasons."

## It's probably too late

While activists would applaud the Pentagon's hiring of an ethicist, preventing or even containing the use of autonomous weapons systems at this point is going to be difficult.

As technology companies like Microsoft, Google, and Amazon further develop and refine AI, and it becomes more accessible to anyone anywhere, the world's militaries will have little choice but to move quickly if they want to stay competitive on the battlefield.

Russia and China are doing just that, according to Sheppard. And they've been developing theirs "without consideration to the operational risks and technology shortfalls," she said. Israel, South Korea, and the UK, meanwhile, are all also developing their own weapons systems that have "significant autonomy" to select and engage targets, according to the Campaign to Stop Killer Robots, a six-year-old coalition of international nonprofits. "If left unchecked the world could enter a destabilizing robotic arms race," the group warns.

The US is already falling behind.

"DoD is really behind the curve compared to the private sector on this task of digital modernization, and that terrifies me," Sheppard said. "You see this example in talking to service members that come back from Syria, and how quickly our adversaries' technology is all commercial, off-the-shelf, and software-driven. These tools and technologies are rapidly developed, rapidly iterated on, and driven by software. Our acquisitions system, the whole defense enterprise, is slow and cumbersome and we're not sending the best technology and the best capabilities forward. The further we fall behind in that curve, the worse it's going to get."

The responsibility lies with the US to apply AI technology in ways that align with American norms and values while still securing a tactical advantage, she said.

"Just because it's a robot doesn't mean it suddenly has unlimited authority to go clean a battlefield," Sheppard said.

## The problems with a ban

Past attempts to globally regulate dangerous weapons technologies have been mixed. Some have outright failed, like Pope Urban II's attempt to ban the crossbow in 1097. On the other hand, most countries have adhered to the ban on chemical and biological weapons, as well as the 1995 preemptive ban on blinding laser weapons.

Efforts to regulate nuclear weapons, however, have been less successful.

"Nuclear weapons, by any reasonable metric, are far worse [than chemical weapons], but they're also much more valuable to countries," Scharre said. "Even if they're not used, they have tremendous value in deterrence. To the point where you see some countries like North Korea willing to go through incredible hardship in order to acquire nuclear weapons."

That isn't the case with chemical weapons, which are not as decisive. Scharre believes nations see fully autonomous weapons like they do nukes, as "very valuable."

The Campaign to Stop Killer Robots has called on countries to at least adhere to a set of guardrails, like retaining "meaningful" human control of all weapons systems. But the US and Russia have so far stymied attempts at regulation, and any potential ban would also be hamstrung by the fact that most countries don't even agree on the definition of "killer robots." Is their software hard-coded, or do they rely on learning-based solutions? Can they be remotely operated? Can they function once remote operators are no longer in contact with them?

"The conditions under which a system is operated is very much defined by how you programmed the software, and verifying would require some level of access or understanding of that," Sheppard said. "And I don't see the reality of countries opening up their systems in that way."

And so without a regulatory way forward, the robot wars remain a possible and terrifying future reality.

"Our chess programs are already smart enough to make their own decisions. They play in ways that are incomprehensible and far superior to humans," said Stuart Russell. "They decide where to move and which enemy pieces to kill, so to speak."