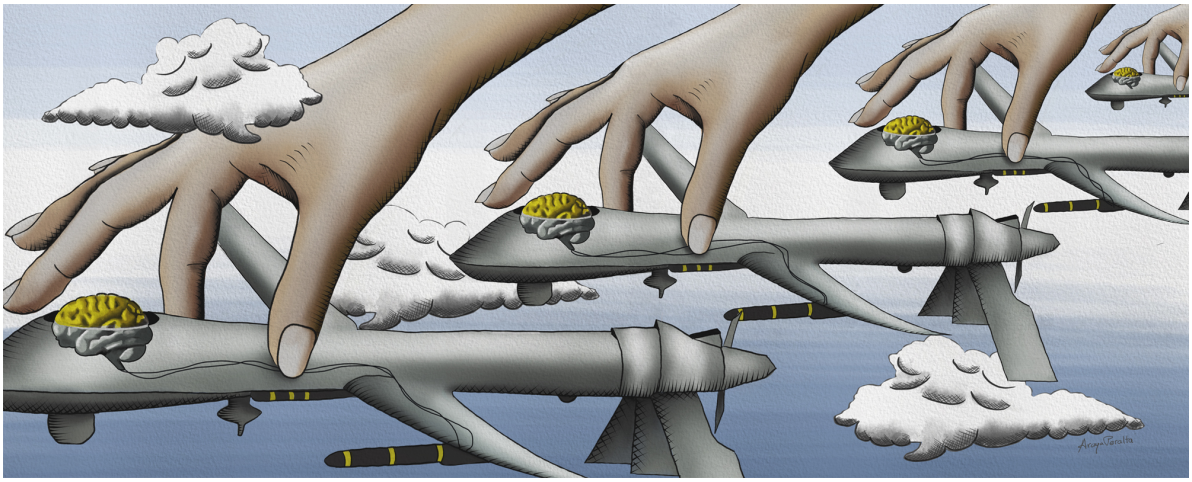
MILITARY USE OF AI

Killer robots

Artificial Intelligence and Autonomous Weapons

Toby Walsh



Araya Peralta

The world faces a very critical choice about the future of warfare. This is not due to the growing political movement against fully autonomous weapons. 28 nations have called on the UN to ban such weapons pre-emptively. Most recently, the European Parliament voted in support of such a ban, whilst the German Foreign Minister Heiko Maas called for international cooperation on regulating autonomous weapons. And in the same week that Maas called for action, Japan gave its backing to international efforts to regulate the development of lethal autonomous weapons at the United Nations.

At the end of 2018, the UN Secretary General, Antonio Guterres addressing the General Assembly offered a stark warning.

> «The weaponization of artificial intelligence is a growing concern. The prospect of weapons that can select and attack a target on their own raises multiple alarms – and could trigger new arms races.Diminished oversight of weapons has implications for our efforts to contain threats, to prevent escalation and to adhere to international humanitarian and human rights law. Let’s call it as it is. The prospect of machines with the discretion and power to take human life is morally repugnant».

No, it’s not this growing political concern that illuminates the critical choice facing the planet. Nor is it the growing movement within civil society against such weapons. The Campaign to Stop Killer Robots, for instance, now numbers over 100 non-governmental organizations such as Human Rights Watch who are vigorously calling for regulation.

It’s also not the pressure of such NGOs to take action. Nor is it the growing concern of the public. A recent IPSOS poll shows that opposition to fully autonomous weapons has increased 10% in the last two years as understanding of the issues grows.

Six out of every ten people in 26 countries polled strongly opposed the use of autonomous weapons. In Spain, for example, 65% of those polled were strongly opposed, whilst less than 20% supported their use. Opposition to and support for autonomous weapons was similar in France, Germany and other European countries.

No, the reason that we face a critical choice today about the future of warfare is that the technology to build autonomous weapons is ready to cross out of the research lab (disclaimer alert: where I work) and to be implemented by arms manufacturers around the world.

In March 2019, for instance, we saw the Royal Australian Air Force announce a partnership with Boeing to develop an unmanned air combat vehicle, a loyal “wingman” to take air combat to the next step of lethality. In the same week, the US Army announced ATLAS, the Advanced Targeting and Lethality Automated System which will be a robot tank. The US Navy also announced that its first fully autonomous ship, the Sea Hunter had made a record breaking voyage from Hawaii to the Californian coast without human intervention.

Unfortunately the world will be a much worse place if, in a decade’s time, militaries around the world are using such lethal autonomous weapons systems (LAWS), and there are no laws regulating LAWS.

> The world will be a much worse place if, in a decade’s time, militaries around the world are using such lethal autonomous weapons systems and there are no laws regulating them

The media like to use the term “killer robot” rather than a wordy expression such as lethal or fully autonomous weapon. But the problem with the media’s term “killer robot” is that this conjures up a picture of the Terminator. And it is not the Terminator that worries me or thousands of my colleagues working in AI. It is much simpler technologies that we see being announced right now.

Take a Predator drone. This is a semi-autonomous weapon. It can fly itself much of the time. However, there is still a soldier, typically in a container in Nevada, who is in over all control. And importantly, it is still a soldier who makes the final life-or-death decision to fire

one of its Hellfire missiles.

But it is a small technical step to replace that soldier with a computer. Indeed, it is technically possible today. And once we build such simple autonomous weapons, there will be an arms race to develop more and more sophisticated versions. Indeed, we can see the beginnings of this arms race. In every theatre of way, in the air, on land, on and under the sea, there are prototype autonomous weapons under development.

This will be a terrible development in warfare. But it is not inevitable. In fact, we get to choose whether we go down this particular road. For over five years now, I and thousands of my colleagues, other researchers in Artificial Intelligence and Robotics have been warning of these dangerous developments. We've been joined by founders of AI and Robotics companies, Nobel Peace Laureates, church leaders, politicians and many members of the public.

Strategically, autonomous weapons are a military dream. They let a military scale their operations unhindered by manpower constraints. One programmer can command hundreds of autonomous weapons. This will industrialise warfare. Autonomous weapons will greatly increase strategic options. They will take humans out of harm's way opening up the opportunity to take on the riskiest of missions. You could call it War 4.0.

There are many reasons, however, why the military's dream of lethal autonomous weapons will turn into a nightmare. First and foremost, there is a strong moral argument against killer robots. We give up an essential part of our humanity if we hand over the decision of whether someone should live to a machine. Machines have no emotions, compassion or empathy. Are machines then fit to decide who lives and who dies?

> To build a nuclear bomb requires technical sophistication. You need the resources of a nation state, and access to fissile material. You need some skilled physicists and engineers. Nuclear weapons have not, as a result, proliferated greatly. Autonomous weapons require none of this

Beyond the moral arguments, there are many technical and legal reasons to be concerned about killer robots. In my view, one of the strongest is that they will revolutionise warfare. Autonomous weapons will be weapons of immense destruction. Previously, if you wanted to do harm, you had to have an army of soldiers to wage war. You had to persuade this army to follow your orders. You had to train them, feed them, and pay them. Now just one programmer could control hundreds of weapons.

Lethal autonomous weapons are more troubling, in some respects, than nuclear weapons. To build a nuclear bomb requires technical sophistication. You need the resources of a

nation state, and access to fissile material. You need some skilled physicists and engineers. Nuclear weapons have not, as a result, proliferated greatly. Autonomous weapons require none of this.

Autonomous weapons will be perfect weapons of terror. Can you imagine how terrifying it will be to be chased by a swarm of autonomous drones? They will fall into the hands of terrorists and rogue states who will have no qualms about turning them on civilians. They will be an ideal weapon with which to suppress a civilian population. Unlike humans, they will not hesitate to commit atrocities, even genocide.

You may be surprised but not everyone is on board with the idea that the world would be a better place with a ban on killer robots in place. “Robots will be better at war than humans,” they say. “Let robot fight robot and keep humans out of it.” Yet these arguments don’t stand up to scrutiny in my view, and that of many of my colleagues working in AI and robotics. Here are the five main objections I hear to banning killer robots — and why they’re misguided.

Objection 1. Robots will be more effective than humans

They’ll be more efficient for sure. They won’t need to sleep. They won’t need time to rest and recover. They won’t need long training programs. They won’t mind extreme cold or heat. All in all, they’ll make ideal soldiers. But they won’t be more effective. The recently leaked Drone Papers suggest nearly nine out of ten people killed by drone strikes weren’t the intended target. This is when there’s still a human in the loop, making the final life-or-death decision.

The statistics will be much worse when we replace that human with a computer. Killer robots will also be more efficient at killing us. Terrorists and rogue nations are sure to use them against us. It’s clear if they’re not banned that there will be an arms race. It is not overblown to suggest that this will be the next great revolution in warfare after the invention of gunpowder and nuclear bombs. The history of warfare is largely one of who can more efficiently kill the other side. This has typically not been a good thing for humankind.

Objection 2. Robots will be more ethical

In the terror of battle, humans have committed many atrocities. And robots can be built to follow precise rules. However, it’s fanciful to imagine we know how to build ethical robots. AI researchers like myself have only just started to worry about how you could program a robot to behave ethically. It will takes us many decades to work this out. And even when we do, there’s no computer we know that can’t be hacked to behave in ways that we don’t desire. Robots today cannot make the distinctions that the international rules of war require: to distinguish between combatant and civilian, to act proportionally, and so on. Robot warfare is likely to be a lot more unpleasant than the war we fight today.

Objection 3. Robots can just fight robots

Replacing humans with robots in a dangerous place like the battlefield might seem like a good idea. However, it's also fanciful to suppose that we could just have robots fight robots. There's not some separate part of the world called "the battlefield." Wars are now fought in our towns and cities, with unfortunate civilians caught in the crossfire. The world is sadly witnessing this today in Syria and elsewhere. Our opponents today are typically terrorists and rogue nations. They are not going to sign up to a contest between robots. Indeed, there's an argument that the terror unleashed remotely by drones has likely aggravated the many conflicts in which we find ourselves today.

Objection 4. Such robots already exist and we need them

I am perfectly happy to concede that a technology like the autonomous Phalanx anti-missing system which sits on many naval ships is a good thing. You don't have time to get a human decision when defending yourself against an incoming supersonic missile. But the Phalanx is a defensive system. And my colleagues and I did not call for defensive systems to be banned. We only called for offensive autonomous systems to be banned. Like the Samsung sentry robot currently active in the DMZ between North and South Korea. This will kill any person who steps into the DMZ from four km away with deadly accuracy. There's no reason we can't ban a weapon system that already exists. Indeed, most bans, like those for chemical weapons or cluster munitions, have been for weapon systems that not only exist, but have been used in war.

Objection 5. Weapon bans don't work

History would contradict this argument. The 1998 UN Protocol on Blinding Lasers resulted in blinding lasers, designed to cause permanent blindness, being kept out of the battlefield. If you go to Syria today — or any of the other war zones of the world — you won't find this weapon, and not a single arms company anywhere in the world will sell it to you. You can't un-invent the technology that supports blinding lasers, but there's enough stigma associated with them that arms companies have stayed away.

I hope a similar stigma will be associated with autonomous weapons. We won't be able to un-invent the technology, but we can put enough stigma in place that robots aren't weaponized. Even a partially effective ban would be likely worth having. Anti-personnel mines still exist today despite the 1997 Ottawa Treaty. But 40 million such mines have been destroyed. This has made the world a safer place and resulted in many fewer children losing their life or a limb.

> We won't be able to un-invent the technology, but we can put enough stigma in place that robots aren't weaponized

AI and robotics can be used for many great purposes. Much the same technology will be needed in autonomous cars as in an autonomous drone. And autonomous cars are predicted

to save 30,000 deaths on the roads of the United States every year. They will make our roads, factories, mines and ports safer and more efficient. They will make our lives healthier, wealthier and happier. In the military setting, there are many good uses of AI. Robots can be used to clear minefields, bring supplies in through dangerous routes, and shift mountains of signal intelligence. But they shouldn't be used to kill.

We stand at a crossroads on this issue. I believe it needs to be seen as morally unacceptable for machines to decide who lives and who dies. In this way, we may be able to save ourselves and our children from this terrible future.

In July 2015, I helped organise an open letter to the UN calling for action that was signed by thousands of my colleagues, other AI researchers. The letter was released at the start of the main international AI conference. Sadly the concerns we raised in this letter have yet to be addressed. Indeed, they have only become more urgent.

Open letter from 2015 signed by thousands of AI researchers

> Autonomous weapons select and engage targets without human intervention. They might include, for example, armed quadcopters that can search for and eliminate people meeting certain pre-defined criteria, but do not include cruise missiles or remotely piloted drones for which humans make all targeting decisions. Artificial Intelligence (AI) technology has reached a point where the deployment of such systems is – practically if not legally – feasible within years, not decades, and the stakes are high: autonomous weapons have been described as the third revolution in warfare, after gunpowder and nuclear arms.
>
> Many arguments have been made for and against autonomous weapons, for example that replacing human soldiers by machines is good by reducing casualties for the owner but bad by thereby lowering the threshold for going to battle. The key question for humanity today is whether to start a global AI arms race or to prevent it from starting. If any major military power pushes ahead with AI weapon development, a global arms race is virtually inevitable, and the endpoint of this technological trajectory is obvious: autonomous weapons will become the Kalashnikovs of tomorrow.

> Unlike nuclear weapons, they require no costly or hard-to-obtain raw materials, so they will become ubiquitous and cheap for all significant military powers to mass-produce. It will only be a matter of time until they appear on the black market and in the hands of terrorists, dictators wishing to better control their populace, warlords wishing to perpetrate ethnic cleansing etc. Autonomous weapons are ideal for tasks such as assassinations, destabilising nations, subduing populations and selectively killing a particular ethnic group. We therefore believe that a military AI arms race would not be beneficial for humanity. There are many ways in which AI can make battlefields safer for humans, especially civilians, without creating new tools for killing people.
>
> Just as most chemists and biologists have no interest in building chemical or biological weapons, most AI researchers have no interest in building AI weapons – and do not want others to tarnish their field by doing so, potentially creating a major public backlash against AI that curtails its future societal benefits. Indeed, chemists and biologists have broadly supported international agreements that have successfully prohibited chemical and biological weapons, just as most physicists supported the treaties banning space-based nuclear weapons and blinding laser weapons.
>
> In summary, we believe that AI has great potential to benefit humanity in many ways, and that the goal of the field should be to do so. Starting a military AI arms race is a bad idea, and should be prevented by a ban on offensive autonomous weapons beyond meaningful human control.

I urge you to join the global campaign to make the world a better place by banning such weapons.



Toby Walsh
Toby Walsh is professor of Artificial Intelligence at the University of New South Wales, Sydney (Australia), and he is a visiting professor at the Technical University of Berlin (Germany). He is a member of the Australian Academy of Science and of the Association for the Progress of Artificial Intelligence. He is the author of a recently published book: *2062: The World that AI Made*, where he explores the impact that the AI might have on society, including the consequences on war.