

# THE END OF WAR

---

There is one job likely to disappear through automation by 2062 that I and many others especially fear. This is the job of fighting wars. Indeed, this replacement has already started to happen. An arms race has begun to develop robots that can replace humans in the battlefield. The media like to call them “killer robots”. But the technical term is “lethal autonomous weapons” or LAWS.

The problem with calling them killer robots is that this conjures up a picture of Terminator and technologies that are a long way away. But it is not Terminator that worries me (or indeed thousands of my colleagues working in Artificial Intelligence). It is much simpler technologies that are, depending on your perspective, at best or at worst less than a decade away. It is stupid AI that I fear currently. We’ll be giving machines that are not sufficiently capable the right to make life or death decisions.

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, 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<sup>39</sup>. And once we build such simple autonomous weapons, there will be an arms race to develop more and more sophisticated versions.

The world will be a much worse place if, in twenty years time, such lethal autonomous weapons are common place and there are no laws about LAWS. This will be a terrible development in warfare. But it is not inevitable. In fact, we get to choose in the next few years whether we go down this particular road.

## THE LURE OF KILLER ROBOTS

The attractions of autonomous weapons to the military are obvious. The

weakest link in a Predator drone is the radio link back to base. Indeed, drones have been sabotaged by jamming their radio link. So if you can have the drone fly, track and target itself, you have a much more robust weapon.

A fully autonomous drone also lets you dispense with a lot of expensive drone pilots. The United States Air Force could be renamed the United States Drone Force. It has more drone pilots than pilots of any other type of plane. By 2062, it won't be just more drone pilots than pilots of any other type of plane but more drone pilots than all other pilots put together. And whilst those drone pilots aren't risking their lives on combat missions, they suffer post traumatic stress disorder at similar rates to the rest of the air force's pilots.

Autonomous weapons offer many other operational advantages. They don't need to be fed or paid. They will fight 24/7. They will have super-human accuracy and reflexes. They will not need evacuating from the battlefield. They will obey every order to the letter. They will not commit atrocities<sup>40</sup> or violate International humanitarian law. They would be perfect soldiers, sailors and pilots.

Strategically, autonomous weapons are also a military dream. They let a military scale their operations unhindered by manpower constraints. One programmer can command hundreds, even thousands 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.

In September 2017, Vladimir Putin was reported to have said that whoever leads in Artificial Intelligence will rule the world. He predicted that future wars will be fought by drones, and when one side's drones are destroyed by drones of another, it will have no other choice but to surrender.

There are, however, many reasons why this military dream will have become a nightmare by 2062.

## THE MORALITY OF KILLING MACHINES

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. Certainly today, machines have no emotions, compassion or empathy. Are machines then fit to decide who lives and who dies?

War is a terrible thing. Lives are lost. People maimed violently and horribly. Civilians bombed. Populations terrorised. We act in war in ways that are not permissible in times of peace. In part, we permit this because the soldiers doing these acts are putting their own lives at stake. You are permitted to kill your enemy for, at that moment, it is your life or theirs.

Because war is a terrible thing, it should not in my view be an easy thing. It should not be something that we fight easily and "cleanly". If history has taught us one thing, the promise of clean wars is and will likely remain an illusion. War needs to be matter a last resort. Politicians need to justify why our sons and daughters are returning home in body bags.

The history of military technology is largely of killing becoming more remote. At first, we fought with our bare hands. Gunpowder let us step back and shoot from a distance. Aeroplanes let us attack from above. And technologies like drones now let us kill people remotely, no longer risking our own lives in the process.

Autonomous weapons are the ultimate step in disengaging us from the act of war. Machines not people will now do the killing all on their own. No humans involved. This changes the nature of warfare fundamentally. And with those changes many of the moral excuses we have for warfare start to disappear.

## WEAPONS OF MASS DESTRUCTION

Beyond the moral arguments, there are many technical and legal reasons to be concerned about killer robots. In my view, one of the strongest reasons for a ban is that they will revolutionise warfare. In fact, it has been called the third revolution in warfare.

The first revolution was the invention of gun powder by the Chinese. The second was the invention of nuclear weapons by the United States. Lethal autonomous weapons will be the third revolution. Each was a step change in the speed and efficiency with which we could kill.

Autonomous weapons will be weapons of mass 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 or even thousands of weapons. Like every other weapon of mass destruction before it<sup>41</sup>, like chemical weapons, biological weapons, and nuclear weapons, we will need to ban such 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 will require none of this. Take a small drone. Program this with a neural network that will identify, track and target any Caucasian face. Such face recognition software can be found in many smartphones today. Now attach a few grammes of high explosive to the drone. By putting together some existing technologies, you have a simple, cheap but nevertheless very lethal autonomous weapon.

Drive a truck with ten thousand of these drones into New York City, and you can mount an attack to rival 9/11. You don't even need to assume that such weapons are very accurate. Suppose your drone only works 1 in 10 times<sup>42</sup>. You would still be able to kill one thousand people in a matter of minutes. Supposing fifty percent accuracy, and you are up to five thousand people in minutes.

Building a weapon like this is much easier than building an autonomous car. 99.99% reliability might be unacceptable for an autonomous car but might

be more than adequate for those wanting to use a killer drone. As many car manufacturers plan to be selling fully autonomous cars by 2025, it is not unreasonable to expect arms manufacturers to be selling such killer drones in a few years time.

## WEAPONS OF TERROR

Autonomous weapons like this will be 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.

There are some who claim that robots can be more ethical than human soldiers. It is, in my view, the most interesting and challenging argument for autonomous weapons. But it ignores that we don't know today how to build autonomous weapons that will follow international humanitarian law.

The rules of war require you to target combatants and not civilians, to act proportional to the threat, to recognise and respect when a combatant is surrendering, or when they are injured and can no longer fight. We don't know yet how to build autonomous weapons that can make such distinctions.

By 2062, I expect that we will have worked out how to build ethical robots. Our lives will be full of autonomous devices which will need to be acting ethical. So it is likely we will have lethal autonomous weapons that could follow international humanitarian law.

However, we won't be able to stop such weapons from being hacked to behave in unethical ways. If you can get physical access to a computer system then you can almost surely hack it. And there are plenty of bad actors out there who will over-ride any safeguards that might be put in place.

Ironically, a number of countries like the United Kingdom oppose a ban on lethal autonomous weapons precisely because they violate international

humanitarian law. No new legislation is needed, they argue, to deal with such weapons.

History disagrees with such arguments. Chemical weapons violate international humanitarian law, in particular the 1925 Geneva Protocol. But in 1993, the Chemical Weapons Convention came into force to regulate them more strongly. The Convention was signed and ratified by the United Kingdom.

The Chemical Weapons Convention strengthened international law to prohibit use of any chemicals in warfare. It set up the Organisation for the Prohibition of Chemical Weapons (OPCW), an intergovernmental body based in The Hague, to monitor the development, production, stockpiling and use of chemical weapons. Today, over 90% of the world's declared stockpile of chemical weapons had now been destroyed. Weapon bans can have positive impacts on our safety and security.

## WEAPONS OF ERROR

In addition to being *weapons of terror*, autonomous weapons will be ***weapons of error***. From a technical perspective, the last place you would want to put a robot is in the battlefield. There's a good reason robots turned up first in places like car factories.

In a factory, you can control the environment. You get to decide where everything and everybody goes. You can even put the robot in a cage to protect by-standers. The battlefield is a very different environment, full of uncertainty and confusion. Not the place that you want to put a robot with deadly potential.

In November 2016, an investigation by *The Intercept* of U.S. military operations against the Taliban and al Qaeda in the Hindu Kush revealed that nearly 9 out of 10 people who died in drone strikes were not the intended targets. Remember this is while we still have a human in the loop, with situational awareness superior today to any machine. And that human is making the final life or death decision. As a technologist, if you asked me to replace the human drone pilot by a machine, I would aspire to make 9 out of 10 errors. I fear we would be making many more.

The potential for error is compounded by the much greater speeds of these weapons. Even with a human on the loop, machines may act too fast for humans to step in and prevent error. And systems of such weapons may behave in unexpected ways. Like on the stock market, they may get into unexpected feedback loops. But unlike the stock market, the results will be deadly and won't be something we can simply unwind. We may even end up fighting "flash wars" that we didn't intend to fight.

Errors create an additional problem. This is the "accountability gap". Who is going to be held responsible when lethal autonomous weapons make mistakes? Who will be court-martialled? Who will be prosecuted in the Hague? The gap is especially large when the weapon uses machine learning to identify and track targets. Now, the manufacturer didn't program the weapon's actual behaviours. They were learnt.

Worse still, the military will be tempted to leave such learning on in the battlefield. If they don't, an adversary will quickly find camouflage and other ways to confuse what is a fixed program. So they will want the autonomous weapon to continue to learn, to adapt to whatever the enemy does. Just like a human soldier would. But if an autonomous weapon is learning, a determined adversary will look for ways to train the weapon to neutralise its threat. They may even be able to train it to turn back on its handler. Who then will be responsible for its errors?

## GEO-POLITICAL STABILITY

At the strategic level, lethal autonomous weapons also pose new threats that might destabilise current stand-offs like that between North and South Korea. A swarm of small stealthy and autonomous drones will be very difficult to defend against today. This may tempt one side to launch a surprise attack. And the fear of such a surprise attack may lower the barriers to the use of greater, even nuclear, force.

Lethal autonomous weapons therefore threaten to upset the current balance of military power. You would no longer need to be an economic super-power to maintain a large and deadly army. It would only take a modest bank balance to have a powerful army of lethal autonomous



weapons. They will be the Kalashnikov's of the future. Unlike nuclear weapons, there will be cheap and easy to produce. And they will turn up on the arms black markets of the world.

This doesn't mean that lethal autonomous weapons can't be banned. Chemical weapons are cheap and easy to produce but have been banned. And we don't need to develop autonomous weapons as a deterrent against those who might ignore a ban. We don't develop chemical weapons to deter those who might sometimes use chemical weapons. We already have plenty of deterrents, military, economic and diplomatic with which to deter those who choose to ignore international treaties.

## CALL TO ARMS

In July 2015, I was sufficiently alarmed with developments in this area that I asked one thousand colleagues, researchers working in AI and Robotics, to sign an open letter calling upon the United Nations to ban offensive autonomous weapons. The letter was released at the start of the main international AI conference, the International Joint Conference on Artificial Intelligence<sup>43</sup>.

One thousand signatures seemed like a nice round number and a substantial fraction of the AI community. To put the number in perspective, the conference itself was expecting around a thousand delegates. By the end of the first day, the number of signatures had doubled from one thousand to two thousand. And it climbed rapidly throughout the week long conference.

The letter got a lot of press in part because it contained the names of some well known people like Stephen Hawking, Elon Musk and Noam Chomsky. But more important in my view is that the letter was signed by many leading researchers in AI and robotics. They came from universities around the world, as well as from commercial labs like Google's DeepMind, Facebook's AI Research Lab, and the Allen Institute for AI. These are the people who arguably best understand the technologies and the limitations.

The United Nations did pay attention to this warning. The letter helped push along informal discussions. And just over one year later in December 2016, at the main disarmament conference, the United Nations decided to



move forwards with formal discussions on the topic. Lethal autonomous weapons are now being considered by a Group of Governmental Experts (GGE), a group mandated by the UN General Assembly to address the issue.

If nations can reach consensus, my hope is that the GGE will put forwards a ban under the umbrella of the Convention on Certain Conventional Weapons. The full title of the Convention is actually “The Convention on Prohibitions or Restrictions on the Use of Certain Conventional Weapons Which May Be Deemed to Be Excessively Injurious or to Have Indiscriminate Effects”. But to diplomats, it is known simply as the CCW. The CCW is the open ended treaty used previously to ban land-mines, booby traps, incendiary weapons, and blinding lasers.

## ARMS RACE

In the open letter, we warned that there would be an arms race to develop more and more capable autonomous weapons. Sadly, that arms race has begun. The Pentagon has allocated \$18 billion in its current budget for the development of new types of weapons, many of them autonomous. Other countries like the United Kingdom, Russia, China and Israel also have sophisticated programmes under way to develop autonomous weapons.

Pick any sphere of the battle, in the air, on the land, on the sea, or under the sea, and there are autonomous weapons under development by militaries around the world. You can even argue that there is at least one autonomous weapon that is operational. This is Samsung’s SGR-A1 Sentry Guard Robot which guards the Demilitarised Zone (DMZ) between North and South Korea.

Now, there is no good reason to step into the DMZ. It is the most highly mined part of the world. But if the mines don’t kill you, Samsung’s robot will. It can automatically identify, target and shoot anyone who steps into no man’s land with its autonomous machine gun. It has deadly accuracy from kilometres away.

There are other weapons in operation that might be considered autonomous. We can exclude mines and other simple technologies as they make no decisions about targeting. But a weapon like the Phalanx anti-

missile system, sitting on ships of the Royal Australian and other navies, acts autonomously. This protects the ship against incoming supersonic missiles with a radar controlled gun. There is no time for a human to react when a missile comes over the horizon. The anti-missile system needs to identify, track and target autonomously.

There is little to be worried about such weapons. They are defensive. They have a very constrained window of operations. They protect the airspace around a naval ship in times of battle. They only target objects moving at supersonic speed. They actually save human lives. Most people, myself included, have little objections to such limited uses of autonomy.

On the other hand, an autonomous drone that loiters over the battlefield for days at a time is much more troubling. The scope of its operations would be much greater, both in time and space. A convoy appears below and it must decide by itself if this is a military convoy, an aid convoy or a wedding party. Machines today cannot make these distinctions reliably.

Like with the development of nuclear weapons, the world is locked into an undesirable course of action. We don't want a world with killer robots, but if our enemies have them, opponents of a ban argue that we had better have some ourselves. And so an arms race has started to develop weapons that we'd rather not have.

Even worse, it is not that true that we need autonomous weapons to defend ourselves against those who might attack us with autonomous weapons. The U.S. is, for example, currently exploring much simpler technologies like nets, and birds of prey to defend against remote controlled drones.

## OBJECTIONS TO A BAN

There are several arguments put forwards against a ban on killer robots. In my view, none of these objections stand up to close examination. One of the most serious objections is that robots will behave more ethically than human soldiers. But, as I argued earlier, we don't know yet how to build ethical robots. And we don't know if they will ever have the compassion and empathy required to behave ethically. Even supposing that we could build robots to behave ethically, we don't know how to build robots that

can't be hacked to behave unethically.

Another objection to a ban is that robots will get our soldiers out of harms way. Indeed, some critics go as far as to argue that we are therefore morally obliged to use them. Perhaps the most troublesome part of this argument is that it completely ignores those facing the killer robots. Killer robots will increase the speed with which we can kill the other side. Killer robots will lower the barriers to war. This ultimately could result in more deaths not less. We cannot care only about our own casualties.

A third objection to a ban is that it is impossible to define autonomous weapon. How can we ban something that we cannot even define? I would agree completely it is difficult to define autonomy. In Artificial Intelligence we're used to this. Most AI researchers have given up trying to define what Artificial Intelligence is. We just get on with building machines that are increasingly capable. I would expect any ban would not define autonomous weapon. It would simply identify that there is a line in the sand that should not be crossed.

Various weapon systems would be clearly one side of the line. A fully autonomous drone that loiters for days over the battlefield would likely be considered to be on the banned side of the line. But international consensus might be that a defensive system like the autonomous Phalanx anti-missile system on naval ships today is on the non-banned side of the line. As new technologies arrive, consensus will emerge as to where they sit with respect to the line.

A fourth objection is that new military technologies have only made the world a safer and less violent place. We should therefore be embracing autonomous weapons. Arguments like those put forwards by Pinker in "The Angels of Our Nature"<sup>44</sup> are often invoked. Pinker makes a convincing argument that the world today is a less violent place, and has less genocide than at any previous point in history.

There is, however, nothing that Pinker says which contradicts the need for a ban. The destructive impact of new technologies has only been curbed by the adoption of international humanitarian law, and new weapon treaties. Indeed, it was the bombs dropped on Venice from balloons by Austria forces

in 1849 - by most accounts the first aerial bombing campaign - that led to the Hague Peace Convention of 1899 banning aerial bombing. Like with other new technologies, a new law is needed to limit the use of killer robots.

A fifth objection is that, unlike other technologies that have been successfully banned like blinding lasers, we are talking about a very broad capability that could be added to almost any existing weapon. And many weapons today already have some limited forms of autonomy. It would be like trying to ban the use of electricity. Worse still, it will be impossible to check if a semi-autonomous weapon has had a software upgrade to make it fully autonomous.

This argument misunderstands how arms treaties work. There is no inspection regime for blinding lasers. There is no police force to ensure arms companies don't build anti-personnel mines. If violations occur, NGOs like Human Rights Watch document them. And headlines appear around the world condemning the acts. Resolutions will be made on the floor of the United Nation. And a distant threat remains of a court in the Hague.

This seems to be enough to ensure arms treaties are violated rarely. It ensures that arms companies don't sell banned weapons, that they aren't found on black markets, and that they don't fall into the hands of terrorists. We could hope for something similar with autonomous weapons.

## FALTERING STEPS

The UN decided at the end of 2016 that the Group of Governmental Experts on autonomous weapons would meet twice in 2017, first in August and again in November just before the annual CCW conference. Unfortunately, despite diplomats acknowledging the urgency of making progress on the issue, the August meeting was cancelled.

The United Nations has adopted new accountancy rules which require every meeting to pay for itself. And several countries, most notably Brazil, are behind in paying their dues. As far as I know, Brazil doesn't have anything special against the discussion on killer robots. It just hasn't paid any of its dues for several years.

To the United Nation's shame, they declined to accept a charitable donation to pay for the August meeting in any case. They argued that the UN only takes money from governments. This seems to forget the billion dollars that Ted Turner gave them. As a consequence, for want of less than a quarter of a million dollars, the issue was left undiscussed.

To shine a spotlight on these delays, I decided to act. At that time, only one company had come out against autonomous weapons, the Canadian company Clearpath Robotics. So I organised founders of over one hundred robotics and AI companies to sign a second open letter calling for the CCW to take action on killer robots.

We again released the letter at the opening of the main international AI conference, the International Joint Conference on Artificial Intelligence which was being held in Melbourne. By chance, the conference began on the very day in August 2017 that the first meeting of the Group of Governmental Experts was supposed to start but had been cancelled.

This second letter was signed by Demis Hassabis and Mustafa Suleyman, two of the founders of DeepMind, as well as many other well known people in AI and Robotics. Other signatures include Geoffrey Hinton and Yoshua Bengio, two of the fathers of deep learning as well as founders of AI companies, and Elon Musk in his capacity as founder of the AI company OpenAI.

Like the first letter, this new letter made news around the world. It demonstrated that industry, in addition to academia, supports the idea of regulating these technologies. It also introduced a phrase much repeated by the press. *"Once this Pandora's box is opened, it will be hard to close."*

At the end of 2017, I and the 137 founders of AI and Robotics companies that signed the letter were voted runner up for "Person of the Year" in the Arms Control Association's annual competition for the year's most influential contribution to disarmament. The rightful winners were the diplomats who negotiated the UN treaty prohibiting nuclear weapons. But it was pleasing to see the issue of autonomous weapons taken so seriously.<sup>45</sup>

Even arms companies can see benefit to a ban. BAE Systems is one of the

largest exporters of arms and a company prototyping the next generation of autonomous systems. At the World Economic Forum in 2016, the Chairman of BAE Systems Sir John Carr argued that fully autonomous weapons would not be able to follow the laws of war. He therefore called upon government to regulate them.

## MOUNTING PRESSURE

Twenty-two countries have so far called for the United Nations to ban lethal autonomous weapons. These are Algeria, Argentina, Bolivia, Brazil, Chile, Costa Rica, Cuba, Ecuador, Egypt, Ghana, Guatemala, Holy See, Iraq, Mexico, Nicaragua, Pakistan, Panama, Peru, State of Palestine, Uganda, Venezuela and Zimbabwe.

There is thus still some distance to go before it is a majority opinion within the United Nations, let alone the consensus. The countries so far in support are those most likely to be on the receiving end of such terrible weapons. There is, however, a growing consensus on the need for “meaningful human control” over any individual attack. This would require the technology to be predictable, the user to have relevant information, and the potential for timely human judgement and intervention.

Other countries are starting to face pressure to act. In November 2017, just before the Group of Governmental Experts on autonomous weapons met for the first time at the United Nations, the Prime Minister of Australia received a letter calling for Australia to become the next country to call for a pre-emptive ban. The letter was signed by over 100 AI and Robotics researchers in Australian universities.

In the interest of full disclosure, I wrote and organized this letter. The Prime Minister of Canada received a similar call signed by over 200 Canadian AI researchers organised by my colleague Professor Ian Kerr who holds the Canada Research Chair in Ethics, Law and Technology at the University of Ottawa.

The Australian letter, like the Canadian letter, argues that lethal autonomous weapons lacking meaningful human control sit on the wrong side of a clear moral line. It asks our government to announce their support

for the call to ban such weapons. “In this way, our government can reclaim its position of moral leadership on the world stage as demonstrated previously in other areas like the non-proliferation of nuclear weapons,” the letter says.

With Australia’s recent election to the UN’s Human Rights Council, the issue of lethal autonomous weapons is even more pressing for Australia to address. Autonomous weapons are fundamentally a human rights issue. The Special Rapporteur to the Human Rights Council Professor Christof Heyns was the first to call upon the UN to address the issue of autonomous weapons, arguing in 2013 that machines should not have life and death powers over humans<sup>46</sup>.

The AI and Robotics communities have sent a clear and consistent message over the past couple of years about autonomous weapons. We warned of an arms race. We can see that arms race beginning. We warned of the considerable technical, legal and moral risks of introducing autonomous weapons into the battlefield. Like with climate change, you can hear a few dissenting scientific voices. We need a moratorium not a ban say some. But the overwhelming majority warn of the considerable dangers.

## ALTERNATIVES TO A BAN

The United Kingdom’s position is that fully autonomous weapons would violate existing international humanitarian law, that the UK would never develop such weapons, and that no new treaties are needed to deal with this issue. There is some truth to the first claim. However, we have no guarantees to the second claim. In the past, the UK has secretly developed chemical and biological weapons. And history contradicts the third claim. New technologies have required strengthening international humanitarian law through the last century.

The primary alternative that the UK proposes to a ban are so called Article 36 weapon reviews. Article 36 of the 1977 Additional Protocol I of the Geneva Conventions requires states to review new weapons, means and methods of warfare to ensure they comply with international humanitarian law. The UK conducts such reviews for any new weapon system.



Article 36 reviews are an unsatisfactory alternative for multiple reasons. First, there is no accepted standard for weapons reviews. How can we ensure that Russia, to pick a country not completely at random, is as tough in reviewing new weapon systems as the UK? Second, there is no example of any weapon system that has ever failed an Article 36 review. This does not bode well for Article 36 reviews keeping any technology out of the battlefield in the future. And third, only a few countries are currently undertaking Article 36 reviews, and those that do have no obligation to publish their results.

## AVOIDING THIS FUTURE

We stand at a crossroads on this issue. We can choose to do nothing. Let arms companies develop and sell lethal autonomous weapons. This will take us to a very unpleasant place. Or we can speak up and hopefully get the United Nations to take action.

The academic community has sent a clear message as to their view. So, too, has industry. And in my experience speaking about the topic around the world, most of the public also strongly support a ban. A 2017 IPSOS survey of people in 23 countries found that, in most countries, a majority of respondents opposed fully autonomous weapons.

With most previous weapons, we had to witness their use before we took action. We had to witness the terrible effects of chemical weapons in the First World War to take action and bring in the 1925 Geneva Protocol. We had to witness the horror of Hiroshima, Nagasaki and live through the several near misses during the the Cold War before we banned nuclear weapons.

My fear then is that we will have to witness the terrifying impact of autonomous weapons before we have the courage and conviction to ban them. We have only one case, blinding lasers, where a ban was introduced pre-emptively. Whatever happens, by 2062 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 from taking this terrible path.