

Armin Krishnan: Ethical and Legal Challenges

How and why did you get interested in the field of military robots?

I got interested in military robots more by accident than by design. I was originally specialized in political philosophy and I later became interested in the privatization of warfare, a tendency which seems to fundamentally weaken the institution of the modern nation state, as it is built on the idea of a monopolization of legitimate force within a territory and the suppression of illegitimate violence deployed beyond its borders. Of course, I came across Peter Singer's excellent book on Private Military Firms, which meant for me that I needed to find a slightly different research problem. After looking for some time intensively for a good and original angle, I ended up researching the transformation of defense and national security industries in terms of shifting from a manufacturing based business concept to a services based business concept. The introduction of high-tech weapons, sensors, and communications meant for the armed forces a greater reliance on contractors for a great variety of tasks, most of them, however, related to maintaining and operating technology and not com-

bat. This is not surprising, as mercenaries have always been a marginal phenomenon in military history, apart from some brief and exceptional periods where they prospered and where they could influence the outcome of major wars.

Anyway, when I was doing my research on military privatization and technology I figured that automation is one of biggest trends in the defense sector. Following the invasion in Afghanistan in late 2001 there has been a substantial increase in the use of military robots by the US military. Many defense projects started in the late 1990s, especially in the aerospace field, are relying on automation and robotics. They are aimed at developing systems that are either completely unmanned or are so automated that they require fewer crew members to operate a platform or system. I knew that there had been outlandish efforts by DARPA of building a robot army in the 1980s and that very little came out of it. This was the very stuff of the 1984 *Terminator* movie, which also highlighted public fears that machines could take over, or at least take away our jobs. So four or five years ago I was

observing a growth in the field of military robotics, but I was still very sceptical about the so-called *Revolution in Military Affairs* and military robots. These weapons and systems seemed only able to contribute very little to the military challenges at hand, namely dealing with internal conflicts characterized by guerrilla warfare and terrorism. On the other hand, I realized that it sometimes does not matter whether a particular weapon or technology is effective with regard to dealing with *present* challenges. The lure of new technology is so great that concerns about usefulness can be ignored and that a new weapon or technology will eventually find its own purpose and application. Automation and robotics has proved to be feasible and useful in many other societal contexts and industries. The armed forces cannot be artificially kept at a lower technological level and there are clearly military applications of robotics. I realized that it was only a matter of time before the military will take full advantage of new technologies such as robotics, no matter what. The next logical step was to consider the implications of having military robots fighting our wars. While precision weapons have helped to remove the human operator as far from danger as possible, wars fought by robots would actually mean that no human operators would need to be put at risk at all. This is indeed a very interesting problem from an ethical

perspective: what is the justification for using force and for killing other people, who we may regard as our enemies, if this could be done without putting any lives at risk and without sacrifice? Would this be a much more humane way of waging war, or its ultimate perversion? This question kept me thinking for a while and encouraged me to write a book on the topic of the legality and ethicality of autonomous weapons. Unfortunately, I still have not yet found the ultimate answer to this question. Maybe the answer will just lie in what society ultimately decides to do with a technology that is so powerful that it may deprive us of purpose and meaning in the long run, as more and more societal functions are getting automated.

In your recent book "Killer Robots: The Legality and Ethicality of Autonomous Weapons" you explore the ethical and legal challenges of the use of unmanned systems by the military. What would be your main findings?

The legal and ethical issues involved are very complex. I found that the existing legal and moral framework for war as defined by the laws of armed conflict and Just War Theory is utterly unprepared for dealing with many aspects of robotic warfare. I think it would be difficult to argue that robotic or autonomous weapons are already outlawed by international law. What

does international law actually require? It requires that noncombatants are protected and that force is used proportionately and only directed against legitimate targets. Current autonomous weapons are not capable of generally distinguishing between legitimate and illegitimate targets, but does this mean that the technology could not be used discriminatively at all, or that the technology will not improve to an extent that it is as good or even better in deciding which targets to attack than a human? Obviously not. How flawless would the technology be required to work, anyway?

Should we demand a hundred percent accuracy in targeting decisions, which would be absurd only looking at the most recent Western interventions in Kosovo, Afghanistan and Iraq, where large numbers of civilians died as a result of bad human decisions and flawed conventional weapons that are perfectly legal. Could not weapons that are more precise and intelligent than present ones represent a progress in terms of humanizing war? I don't think that there is at the moment any serious legal barrier for armed forces to introduce robotic weapons, even weapons that are highly automated and capable of making their own targeting decisions. It would depend on the particular case when they are used to determine whether this particular use violated interna-

tional law, or not. The development and possession of autonomous weapons is clearly not in principle illegal and more than 40 states are developing such weapons, indicating some confidence that legal issues and concerns could be resolved in some way. More interesting are ethical questions that go beyond the formal legality. For sure, legality is important, but it is not everything. Many things or behaviors that are legal are certainly not ethical.

So one could ask, if autonomous weapons can be legal would it also be ethical to use them in war, even if they were better at making targeting decisions than humans? While the legal debate on military robotics focuses mostly on existing or likely future technological capabilities, the ethical debate should focus on a very different issue, namely the question of fairness and ethical appropriateness. I am aware that "fairness" is not a requirement of the laws of armed conflict and it may seem odd to bring up that point at all. Political and military decision-makers who are primarily concerned about protecting the lives of soldiers they are responsible for clearly do not want a fair fight. This is a completely different matter for the soldiers who are tasked with fighting wars and who have to take lives when necessary. Unless somebody is a psychopath, killing without risk is

psychologically very difficult. Teleoperators of the armed *Predator* UAVs actually seem to suffer from higher levels of stress than jet pilots who fly combat missions. Remote controlling or rather supervising robotic weapons is not a job well suited for humans or a job soldiers would particularly like to do. So why not just leave tactical targeting decisions to an automated system (provided it is reliable enough) and avoid this psychological problem? This brings the problem of emotional disengagement from what is happening on the battlefield and the problem of moral responsibility, which I think is not the same as legal responsibility. Autonomous weapons are devices rather than tools. They are placed on the battlefield and do whatever they are supposed to do (if we are lucky). The soldiers who deploy these weapons are reduced to the role of managers of violence, who will find it difficult to ascribe individual moral responsibility to what these devices do on the battlefield. Even if the devices function perfectly and only kill combatants and only attack legitimate targets, we will not feel ethically very comfortable if the result is a one-sided massacre. Any attack by autonomous weapons that results in death could look like a massacre and could be ethically difficult to justify, even if the target somehow deserved it. No doubt, it will be ethically very challenging to find

acceptable roles and missions for military robots, especially for the more autonomous ones. In the worst case, warfare could indeed develop into something in which humans only figure as targets and victims and not as fighters and deciders. In the best case, military robotics could limit violence and fewer people will have to suffer from war and its consequences. In the long term, the use of robots and robotic devices by the military and society will most likely force us to rethink our relationship with the technology we use to achieve our ends. Robots are not ordinary tools, but they have the potential for exhibiting genuine agency and intelligence. At some point soon, society will need to consider the question of what are ethically acceptable uses of robots. Though “robot rights” still look like a fantasy, soldiers and other people working with robots are already responding emotionally to these machines. They bond with them and they sometimes attribute to the robots the ability to suffer. There could be surprising ethical implications and consequences for military uses of robots.

Do you think that using automated weapon systems under the premise of e.g. John Canning's concept (targeting the weapon systems used and not the soldier using it) or concepts like “mobility kill” or “mission kill” (where the primary goal is

to deny the enemy his mission, not to kill him) are ethically practicable ways to reduce the application of lethal force in armed conflicts?

John Canning was not a hundred percent happy with how I represented his argument in my book, so I will try to be more careful in my answer. First of all, I fully agree with John Canning that less than lethal weapons are preferable to lethal weapons and that weapons that target “things” are preferable to weapons that target humans. If it is possible to successfully carry out a military mission without using lethal force, then it should be done in this way.

In any case it is a very good idea to restrict the firepower that autonomous weapons would be allowed to control. The less firepower they control, the less damage they can cause when they malfunction or when they make bad targeting decisions. In an ideal case the weapon would only disarm or temporarily disable human enemies. If we could decide military conflicts in this manner, it would be certainly a great progress in terms of humanizing war. I have no problem with this ideal. Unfortunately, it will probably take a long time before we get anywhere close to this vision. Nonlethal weapons have matured over the last two decades, but they are still not yet considered to be generally a reasonable alternative to lethal

weapons in most situations. In conflict zones soldiers still prefer life ammunition to rubber bullets or TASERS since real bullets guarantee an effect and nonlethal weapons don't guarantee to stop an attacker. Pairing nonlethal weapons with robots offers a good compromise, as no lives would be at stake in case nonlethal weapons prove ineffective. On the other hand, it would mean to allow a robot targeting humans in general. It is not very likely that robots will be able to distinguish between a human who is a threat and a human who isn't. It is hard enough for a computer or robot to recognize a human shape – recognizing a human and that this human carries a weapon and is a threat is much more difficult. This means that many innocent civilians, who deserve not to be targeted at all, are likely to be targeted by such a robot. The effects of the nonlethal weapon would need to be very mild in order to make the general targeting of civilians permissible. There are still serious concerns about the long term health effects of the *Active Denial System*, for example.

To restrict autonomous weapons to targeting “things” would offer some way out of the legal dilemma of targeting innocent civilians, which is obviously illegal. If an autonomous weapon can reliably identify a tank or a fighter jet, then I would see no legal problem to allow the weapon to attack targets that are clearly

military. Then again it would depend on the specific situation and the overall likelihood that innocents could be hurt. Destroying military targets requires much more firepower than targeting individuals or civilian objects. More firepower always means greater risk of collateral damage. An ideal scenario for the use of such autonomous weapons would be their use against an armored column approaching through uninhabited terrain. That was a likely scenario for a Soviet attack in the 1980s, but it is a very unlikely scenario in today's world. The adversaries encountered by Western armed forces deployed in Iraq or in Afghanistan tend to use civilian trucks and cars, even horses, rather than tanks or fighter jets. A weapon designed to autonomously attack military "things" is not going to be of much use in such situations. Finally, John Canning proposed a "dial-a-autonomy" function that would allow the weapon to call for help from a human operator in case lethal force is needed. This is some sort of compromise for the dilemma of giving the robot lethal weapons and the ability to target humans with nonlethal weapons and of taking advantage of automation without violating international law. I do not know whether this approach will work in practice, but one can always be hopeful. Most likely weapons of a high autonomy will only be useful in high-intensity conflicts and they will have to con-

trol substantial firepower in order to be effective against military targets. Using autonomous weapons amongst civilians, even if they control only nonlethal weapons, does not seem right to me.

In your book you also put the focus on the historical developments of automated weapons. Where do you see the new dimension in modern unmanned systems as opposed to for example intelligent ammunitions like the cruise missile or older teleoperated weapon systems like the "Goliath" tracked mine during the Second World War.

The differences between remotely controlled or purely automated systems and current teleoperated systems like *Predator* are huge. The initial challenge in the development of robotics was to make automatons mechanically work. Automatons were already built in Ancient times, were considerably improved by the genius of Leonardo da Vinci, and were eventually perfected in the late 18th century. Automatons are extremely limited in what they can do and there were not many useful applications for them. Most of the time they were just used as toys or for entertainment. In terms of military application there was the development of the explosive "mine" that could trigger itself, which is nothing but a simple automaton. The torpedo and the "aerial torpedo" developed in the First World

War are also simple automatons that were launched in a certain direction with the hope of destroying something valuable. In principle, the German V1 and V2 do not differ that much from earlier and more primitive automated weapons. With the discovery of electricity and the invention of radio it became possible to remote control weapons, which is an improvement over purely automated weapons in so far as the human element in the weapons system could make the remote controlled weapon more versatile and more intelligent. For sure, remote controlled weapons were no great success during the Second World War and they were therefore largely overlooked by military historians.

A main problem was that the operator had to be in proximity to the weapon and that it was very easy to make the weapon ineffective by cutting the communications link between operator and weapon. Now we have TV control, satellite links and wireless networks that allow an operator to have sufficient situational awareness without any need of being close to the remotely controlled weapon. This works very well, for the moment at least, and this means that many armed forces are interested in acquiring teleoperated systems like *Predator* in greater numbers. The US operates already almost 200 of them. The UK operates two of the heavily armed

Reaper version of the *Predator* and has several similar types under development. The German Bundeswehr is determined to acquire armed UAVs and currently considers buying the *Predator*. Most of the more modern armed forces around the world are in the stage of introducing such weapons and, as pointed out before, the US already operates substantial numbers of them. The new dimension of *Predator* opposed to the V1 or *Goliath* is that it combines the strengths of human intelligence with an effective way of operating the weapon without any need of having the operator in close proximity. Technologically speaking the *Predator* is not a major breakthrough, but militarily its success clearly indicates that there are roles in which “robotic” systems can be highly effective and even can exceed the performance of manned systems. The military was never very enthusiastic about using automated and remote controlled system, apart from mine warfare, mainly because it seemed like a very ineffective and costly way for attacking the enemy. Soldiers and manned platforms just perform much better.

This conventional wisdom is now changing. The really big step would be the development of truly autonomous weapons that can make intelligent decisions by themselves and that do not require an operator in order to carry out

their missions. Technology is clearly moving in that direction. For some roles, such as battlespace surveillance, an operator is no longer necessary. A different matter is of course the use of lethal force. Computers are not yet intelligent enough that we could feel confident about sending an armed robot over the hill and hope that the robot will fight effectively on its own while obeying the conventions of war. Certainly, there is a lot of progress in artificial intelligence research, but it will take a long time before autonomous robots can be really useful and effective under the political, legal and ethical constraints under which modern armed forces have to operate. Again introducing autonomous weapons on a larger scale would require a record of success for autonomous weapons that proves the technology works and can be useful. Some cautious steps are taken in that direction by introducing armed sentry robots, which guard borders and other closed off areas. South Korea, for example, has introduced the Samsung Techwin SGR-1 stationary sentry robot, which can operate autonomously and controls lethal weapons. There are many similar systems that are field tested and these will establish a record of performance. If they perform well enough, armed forces and police organizations will be tempted to use them in offensive roles or within cities. If that happened, it

would have to be considered a major revolution or discontinuity in the history of warfare and some might argue even in the history of mankind, as Manuel DaLanda has claimed.

Do you think that there is a need for international legislation concerning the development and deployment of unmanned systems? And how could a legal framework of regulations for unmanned systems look like?

The first reflex to a new kind of weapon is to simply outlaw it. The possible consequences of robotic warfare could be similarly serious as those caused by the invention of the nuclear bomb. At that time (especially in the 1940s and 1950s) many scientists and philosophers lobbied for the abolition of nuclear weapons. As it turned out, the emerging nuclear powers were not prepared to do so. The world came several times close to total nuclear war, but we have eventually managed to live with nuclear weapons and there is reasonable hope that their numbers could be reduced to such an extent that nuclear war, if it should happen, would at least no longer threaten the survival of mankind. There are lots of lessons that can be learned from the history of nuclear weapons with respect to the rise of robotic warfare, which might have similar, if not greater repercussions for warfare.

I don't think it is possible to effectively outlaw autonomous weapons completely. The promises of this technology are too great to be ignored by those nations capable of developing and using this technology. Like nuclear weapons autonomous weapons might only indirectly affect the practice of war. Nations might decide to come to rely on robotic weapons for their defense. Many nations will stop having traditional air forces because they are expensive and the roles of manned aircraft can be taken over by land based systems and unmanned systems. I would expect the roles of unmanned systems to be first and foremost defensive. One reason for this is that the technology is not available to make them smart enough for many offensive tasks. The other reason is that genuinely offensive roles for autonomous weapons may not be ethically acceptable. A big question will be how autonomous should robotic systems be allowed to become and how to measure or define this autonomy. Many existing weapons can be turned into robots and their autonomy could be substantially increased by some software update. It might not be as difficult for armed forces to transition to a force structure that incorporates many robotic and automated systems. So it is quite likely that the numbers of unmanned systems will continue to grow and that they will replace lots of sol-

diers or take over many jobs that still require humans.

At the same time, armed conflicts that are limited internal conflicts will continue to be fought primarily by humans. They will likely remain small scale and low tech. Interstate conflict, should it still occur, will continue to become ever more high-tech and potentially more destructive. Hopefully, politics will become more skilled to avoid these conflicts. All of this has big consequences for the chances of regulating autonomous weapons and for the approaches that could be used. I think it would be most important to restrict autonomous weapons to purely defensive roles. They should only be used in situations and in circumstances when they are not likely to harm innocent civilians. As mentioned before, this makes them unsuitable for low-intensity conflicts. The second most important thing would be to restrict the proliferation of autonomous weapons. At the very least the technology should not become available to authoritarian regimes, which might use it against their own populations, and to non-state actors such as terrorists or private military companies. Finally, efforts should be made to prevent the creation of superintelligent computers that control weapons or other important functions of society and to prevent "doomsday systems" that can automatically retaliate against any attack. These are still

very hypothetical dangers, but it is probably not too soon to put regulatory measures in place, or at least not too soon for having a public and political debate on these dangers.

Nonproliferation of robotic technology to nonstate actors or authoritarian regimes, which I think definitively an essential goal, might be possible for dedicated military systems but seems to be something which might not be easily achieved in general, as already can be seen by the use of unmanned systems by the Hamas. In addition the spread of robot technology in the society in nonmilitary settings will certainly make components widely commercially available. How do you see the international community countering this threat?

Using a UAV for reconnaissance is not something really groundbreaking for Hamas, which is a large paramilitary organization with the necessary resources and political connections. Terrorists could have used remote-controlled model aircraft for terrorist attacks already more than thirty years ago. Apparently the Red Army Fraction wanted to kill the Bavarian politician Franz-Josef Strauß in 1977 with a model aircraft loaded with explosives. This is not a new idea. For sure the technology will become more widely available and maybe future terrorists will become more technically skilled. If somebody really wanted

to use model aircraft in that way or to build a simple UAV that is controlled by a GPS signal, it can clearly be done. It is hard to say why terrorists have not used such technology before. Robotic terrorism is still a hypothetical threat rather than a real threat. Once terrorists start using robotic devices for attacks it will certainly be possible to put effective countermeasures in place such as radio jammers. There is a danger that some of the commercial robotic devices that are already on the market or will be on the market soon could be converted into robotic weapons. Again that is possible, but terrorists would need to figure out effective ways of using such devices.

Generally speaking, terrorists tend to be very conservative in their methods and as long as their current methods and tactics “work” they have little reason to use new tactics that require more technical skills and more difficult logistics, unless those new tactics would be much more effective. I don’t think that would be already the case. At the same time, it would make sense for governments to require manufacturers of robotic devices to limit the autonomy and uses of these devices, so that they could not be converted easily into weapons. I think from a technical point of view that would be relatively easy to do. National legislation would suffice and it would probably not require

international agreements. To tackle the proliferation of military robotics technology to authoritarian regimes will be much more challenging. Cruise missile technology has proliferated quickly in the 1990s and more than 25 countries can build them. Countries like Russia, Ukraine, China, and Iran have proliferated cruise missile technology and there is little the West can do about it, as cruise missiles are not sufficiently covered by the Missile Technology Control Regime. What would be needed is something like a military robotics control regime and hopefully enough countries would sign up for it.

A lot of people see the problem of discrimination and proportionality as the most pressing challenges concerning the deployment of unmanned systems. Which are the issues you think need to be tackled right now in the field of law of armed combat?

I think most pressing would be to define autonomous weapons under international law and agree on permissible roles and functions for these weapons. What is a military robot or an “autonomous weapon” and under which circumstances should the armed forces be allowed to use them? It will be very difficult to get any international consensus on a definition, as there are different opinions on what a “robot” is or what constitutes “autonomy”. At the

same time, for any kind of international arms control treaty to work it has to be possible to monitor compliance to the treaty. Otherwise the treaty becomes irrelevant. For example, the Biological and Toxin Weapons Convention of 1972 outlawed biological weapons and any offensive biological weapons research, but included no possibility of monitoring compliance through on-site inspections. As a result, the Soviet Union violated the treaty on massive scale. If we want to constrain the uses and numbers of military robots effectively we really need a definition that allows determining whether or not a nation is in compliance with these rules. If we say teleoperated systems like *Predator* are legal, while autonomous weapons that can select and attack targets by themselves would be illegal, there is a major problem with regard to arms control verification. Arms controllers would most likely need to look very closely at the weapons systems, including the source code for its control system, in order to determine the actual autonomy of the weapon. A weapon like *Predator* could theoretically be transformed from a teleoperated system to an autonomous system through a software upgrade. This might not result in any visible change on the outside. The problem is that no nation would be likely to give arms controllers access to secret military technology. So how can we monitor

compliance? One possibility would be to set upper limits for all military robots of a certain size no matter whether they would be teleoperated or autonomous. This might be the most promising way to go about restricting military robots. Then again, it really depends on how one defines military robots. Under many definitions of robots a cruise missile would be considered a robot, especially as they could be equipped with a target recognition system and AI that allows the missile to select targets by itself. So there is a big question how inclusive or exclusive a definition of "military robot" should be. If it is too inclusive there will never be an international consensus, as nations will find it difficult to agree on limiting or abolishing weapons they already have. If the definition is too exclusive, it will be very easy for nations to circumvent any treaty by developing robotic weapons that would not fall under this definition and would thus be exempted from an arms control treaty.

Another way to go about arms control would be to avoid any broad definition of "military robot" or "autonomous weapon" and just address different types of robotic weapons in a whole series of different arms control agreements. For example, a treaty on armed unmanned aerial vehicles of a certain size, another treaty on armed unmanned land vehicles of a certain

size, and so on. This will be even more difficult or at least time consuming to negotiate, as different armed forces will have very different requirements and priorities with regard to acquiring and utilizing each of these unmanned systems categories. Once a workable approach is found in terms of definitions and classifications, it would be crucial to constrain the role of military robots to primarily defensive roles such as guard duty in closed off areas. Offensive robotic weapons such as *Predator* or cruise missiles that are currently teleoperated or programmed to attack a certain area/target, but that have the potential of becoming completely autonomous relatively soon, should be clearly limited in numbers, no matter whether or not they already have to be considered autonomous. At the moment, this is not urgent as there are technological constraints with respect to the overall number of teleoperated systems that can be operated at a given time. In the medium to long-term these constraints could be overcome and it would be important to have an arms control treaty on upper limits for the numbers of offensive unmanned systems that the major military powers would be allowed to have.

Apart from the Missile Technology Control Regime, there seem to be no clear international regulations concerning the use of unmanned systems. What is the relevance of

customary international law, like the Martens Clause, in this case?

Some academics take the position that “autonomous weapons” are already illegal under international law, even if they are not explicitly prohibited, as they go against the spirit of the conventions of war. For example, David Isenberg claims that there has to be a human in the loop in order for military robots to comply with customary international law. In other words, teleoperated weapons are OK, but autonomous weapons are illegal. This looks like a reasonable position to have, but again the devil is in the detail. What does it actually mean that a human is “in the loop” and how do we determine that a human was in the loop post facto?

I already mentioned this problem with respect to arms control. It is also a problem for monitoring the compliance to the *jus in bello*. As the number of unmanned systems grows, the ratio between teleoperators and unmanned systems will change with fewer and fewer humans operating more and more robots at a time. This means most of the time these unmanned systems will make decisions by themselves and humans will only intervene when there are problems. So one can claim that humans remain in the loop, but in reality the role of humans would be reduced to that of supervision and management. Be-

sides there is a military tradition of using self-triggering mines and autonomous weapons have many similarities with mines. Although anti-personnel land mines are outlawed, other types of mines such as sea mines or anti-vehicle mines are not outlawed. I think it is difficult to argue that autonomous weapons should be considered illegal weapons under customary international law. Nations have used remote-controlled and automated weapons before in war and that was never considered to be a war crime in itself.

The bigger issue than the question of the legality of the weapons themselves is their usage in specific circumstances. If a military robot is used for deliberately attacking civilians, it would be clearly a violation of the customs of war. In this case it does not matter that the weapon used was a robot rather than an assault rifle in the hands of a soldier. Using robots for violating human rights and the conventions of war does not change anything with regard to illegality of such practices. At the same time, using an autonomous weapon to attack targets that are not protected by the customs of war does not seem to be in itself to be illegal or run counter the conventions of war. Autonomous weapons would only be illegal if they were completely and inherently incapable of complying with the customs of war. Even then the decision about

the legality of autonomous weapons would be primarily a political decision rather than a legal decision. For example, nuclear weapons are clearly weapons that are not discriminative and that are disproportionate in their effects. They should be considered illegal under customary international law, but we are still far away from outlawing nuclear weapons. The established nuclear powers are still determined to keep sizeable arsenals and some states still seek to acquire them. One could argue that nuclear weapons are just the only exception from the rule because of their tremendous destructive capability that makes them ideal weapons for deterrence. Furthermore, despite the fact that nuclear weapons are not explicitly outlawed there is a big taboo on their use. Indeed, nuclear weapons have never been used since the Second World War. It is possible that in the long run autonomous weapons could go down a very similar path.

The technologically most advanced states are developing autonomous weapons in order to deter potential adversaries. But it is possible that a taboo against their actual usage in war might develop. In military conflicts where the stakes remain relatively low such as in internal wars a convention could develop not to use weapons with a high autonomy, while keeping autonomous weapons ready for possible high-intensity

conflicts against major military powers, which have fortunately become far less likely. This is of course just speculation.

Another aspect which has come up in the discussion of automated weapon systems is the locus of responsibility. Who is to be held responsible for whatever actions the weapons systems takes? This may not be a big issue for teleoperated systems but gets more significant the more humans are distanced from "the loop".

Are we talking about legal or moral responsibility? I think there is a difference. The legal responsibility for the use of an autonomous weapon would still need to be defined. Armed forces would need to come up with clear regulations that define autonomous weapons and that restrict their usage. Furthermore, there would need to be clear safety standards for the design of autonomous weapons. The manufacturer would also have to specify the exact limitations of the weapon. The legal responsibility could then be shared between a military commander, who made the decision to deploy an autonomous weapon on the battlefield and the manufacturer, which built the weapon. If something goes wrong one could check whether a commander adhered to the regulations when deploying the system and whether the system itself functioned in the way guaranteed by the

manufacturer. Of course, the technology in autonomous weapons is very complex and it will be technically challenging to make these weapons function in a very predictable fashion, which would be the key to any safety standard. If an autonomous weapon was not sufficiently reliable and predictable, it would be grossly negligent of a government to allow the deployment of such weapons in the first place. With respect to moral responsibility the matter is much more complicated. It would be difficult for individuals to accept any responsibility for actions that do not originate from themselves. There is a big danger that soldiers get morally “disengaged” and that they no longer feel guilty about the loss of life in war once robots decide whom to kill. As a result, more people could end up getting killed, which is a moral problem even if the people killed are perfectly legal targets under international law. The technology could affect our ability to feel compassion for our enemies. Killing has always been psychologically very difficult for the great majority of people and it would be better if it stayed that way. One way to tackle the problem would be to give the robot itself a conscience. However, what is currently discussed as a robot conscience is little more than a system of rules. These rules may work well from an ethical perspective, or they may not work well. In any case such a robot conscience is no substitute

for human compassion and ability to feel guilty about wrongdoings. We should be careful with taking that aspect of war away. In particular, there is the argument that bombers carrying nuclear weapons should continue to be manned, as humans will always be very reluctant to pull the trigger and will only do so in extreme circumstances. For a robot pulling the trigger is no problem, as it is just an algorithm that decides and as the robot will always remain ignorant of the moral consequences of that decision.

In addition to the common questions concerning autonomous unmanned systems and discrimination and proportionality you have also emphasized the problem of targeted killing. Indeed, the first weaponized UAVs have been used in exactly this type of operation, e.g. the killing of Abu Ali al-Harithi in Yemen in November 2002. How would you evaluate these operations from a legal perspective?

There are two aspects to targeted killings of terrorists. The first aspect is that lethal military force is used against civilians in circumstances that cannot be defined legally as a military conflict or war. This is in any case legally problematic no matter how targeted killings are carried out. In the past Special Forces have been used for targeted killings of terrorists. So the *Predator* strikes are in this respect not something

new. For example, there has been some debate on the legality of the use of ambushes by the British SAS aimed at killing IRA terrorists. If there was an immediate threat posed by a terrorist and if there were no other ways of arresting the terrorist or of otherwise neutralising the threat, it is legitimate and legal to use lethal force against them. The police are allowed to use lethal force in such circumstances and the military should be allowed to do the same in these circumstances. At the same time, one could question in the specific cases whether lethal action was really necessary. Was there really no way to apprehend certain terrorists and to put them to justice? I seriously doubt that was always the case when lethal action was used against terrorists.

This brings us to the second aspect of the question. I am concerned about using robotic weapons against terrorists mainly because it makes it so easy for the armed forces and intelligence services to kill particular individuals, who may be guilty of serious crimes or not. "Terrorist" is in itself a highly politicised term that has often been applied to any oppositionists and dissenters out of political convenience. Besides it is always difficult to evaluate the threat posed by an individual, who may be a "member" of a terrorist organization or may have contacts to "terrorists". If we define terrorism as war requiring a military response and if

we use robotic weapons to kill terrorists rather than apprehend them, we could see the emergence of a new type of warfare based on assassination of key individuals. Something like that has been tried out during the Vietnam War by the CIA and it was called *Phoenix Program*. The aim was to identify the Vietcong political infrastructure and take it out through arrest or lethal force. In this context 20,000 South Vietnamese were killed. Robotic warfare could take such an approach to a completely new level, especially, if such assassinations could be carried out covertly, for example through weaponized microrobots or highly precise lasers. This would be an extremely worrying future scenario and the West should stop using targeted killings as an approach to counterterrorism.

Where do you see the main challenges concerning unmanned systems in the foreseeable future?

I think the main challenges will be ethical and not technological or political. Technology advances at such a rapid pace that it is difficult to keep up with the many developments in the technology fields that are relevant for military robotics. It is extremely difficult to predict what will be possible in ten or 20 years from now. There will always be surprises in terms of breakthroughs that did not happen and breakthroughs that happened. The best

prediction is that technological progress will not stop and that many technological systems in place today will be replaced by much more capable ones in the future. Looking at what has been achieved in the area of military robotics in the last ten years alone gives a lot of confidence for saying that the military robots of the future will be much more capable than today's. Politics is much slower in responding to rapid technological progress and national armed forces have always tried to resist changes. Breaking with traditions and embracing something as revolutionary as robotics will take many years. On the other hand, military robotics is a revolution that has been already 30 years in the making. Sooner or later politics will push for this revolution to happen. Societies will get used to automation and they will get used to the idea of autonomous weapons. If one considers the speed with which modern societies got accustomed to mobile phones and the Internet, they will surely become similarly quickly accustomed to robotic devices in their everyday lives. It will take some time for the general public to accept the emerging practice of robotic warfare, but it will happen. A completely different matter is the ethical side of military robotics. There are no easy answers and it is not even likely that we will find them any time soon. The problem is that technology and politics will most likely outpace the development of

an ethic for robotic warfare or for automation in general. For me that is a big concern. I would hope that more public and academic debate will result in practical ethical solutions to the very complex ethical problem of robotic warfare.