

Robert Sparrow: The Ethical Challenges of Military Robots

How and why did you become interested in the field of military ethics and, in particular, the field of military robots?

I've been interested in military ethics ever since I first started studying philosophy at the age of 17. I've always thought that questions of political philosophy are the most urgent philosophical questions because they relate to the way we should live alongside each other. Questions of military ethics – or at least of Just War theory – have been some of the most controversial political questions in Australia, given the Australian government's tendency to follow the United States into various wars around the globe despite the absence, in most cases, of any direct threat to Australia. So I have always been interested in Just War theory insofar as it provided me with the tools to think about the justification of these wars.

I became interested in the ethics of military robotics via a more round-about route. I originally started writing about ethical issues to do with (hypothetical) artificial intelligences as an exercise in applying some novel arguments in moral psychology. Similarly, I wrote a paper about

the ethics of manufacturing robot pets such as Sony's Aibo in order to explore some issues in virtue ethics and the ethics of representation. However, in the course of writing about robot pets I began reading up on contemporary robotics and became aware of just how much robotics research was funded by the military. So I wrote my paper, "Killer Robots", partly – like the earlier papers – as a way of investigating the relationship between moral responsibility and embodiment, but also because I thought there was a real danger that the development of military robots might blur the responsibility for killing to the point where no one could be held responsible for particular deaths. Since then, of course, with the development and apparent success of Global Hawk and Predator, robotic weapons have really taken off (pardon the pun!) so that issues that even 10 years ago looked like science fiction are now urgent policy questions. Consequently, my current research is much more focused on responding to what we know about how these weapons are used today.

The United States' Army's Future Combat System is probably the

most ambitious project for fielding a hybrid force of soldiers and unmanned systems to date. From a general perspective, what are your thoughts on the development and deployment of unmanned systems by the military?

In a way, I think the current enthusiasm for military robotics is a reflection of the success of anti-war movements in making it more difficult for governments to sustain public support for war once soldiers start coming home in body bags. I suspect that governments and generals look at unmanned systems and see the possibility of being able to conduct wars abroad over long periods without needing to worry about losing political support at home. So the desire to send robots to fight is a perverse consequence of the triumph of humanist values. The extent to which this development has occurred at the cost of concern for the lives of the citizens of the countries in which these wars are fought is an indication of the limited nature of that triumph.

At the same time, of course, it's entirely appropriate and indeed admirable that the people in charge of weapons research and procurement should be concerned to preserve the lives of the men and women that governments send into combat. Unmanned systems clearly have a valuable role to play in this regard and it would be a mistake to

downplay this. It is difficult to see how there could be anything wrong with the use of robots to neutralise IEDs or clear minefields, for instance.

I also think there is a certain "gee whiz" around robot weapons that is responsible for much of the enthusiasm for them at the moment. Certainly, it's easier to get the public excited about a military robot than about human beings fulfilling similar roles. And I suspect this is even true *within* some parts of the military-industrial complex. Defence ministers want to be able to claim that their country has the most "advanced" weapons, even where the new weapons don't perform that differently from the old. Spending money on military equipment puts more money in the pockets of the corporations that provide campaign funding than does spending money on personnel, which works to the advantage of the robots. It's also worth remembering that there is often an enormous gap between what arms manufacturers claim a system will be capable of when it is commissioned and what they actually deliver. This is especially the case with robots. The PowerPoint presentations and promotional videos in which the systems function flawlessly are often a far cry from the reality of how they work in chaotic environments. However, it is surprising how influential the PowerPoint presentations seem to

be when it comes to determining which systems are funded.

Finally, even if systems do function reliably, it is possible they will be much less useful than their designers intend. One suspects that, in the not-too-distant future, there will be a re-evaluation of the usefulness of military robots, with people realising they are a good solution only in a very limited range of circumstances. To a person with a hammer, everything looks like a nail, so when militaries possess unmanned systems they will tend to want to use them. Yet there is more to war than blowing people up. It's pretty clear that the Predator is precisely the *wrong* weapon to use to try to "win" the war in Afghanistan, for instance. Insofar as anyone has any idea about what it would mean to win this war, it would involve winning the "hearts and minds" of Afghans to the West's cause and creating conditions that might allow Afghans to govern themselves and to live free of poverty and fear. No amount of destroying "high-value targets" from 16,000 feet will accomplish this. Indeed, it seems probable that the civilian casualties associated with Predator strikes radically decrease popular support in Afghanistan for Western goals there. As David Kilcullen and Andrew McDonald Exum pointed out in a recent *New York Times* opinion piece, missile strikes from Predator are a tactic

substituting for a strategy. There are features of unmanned systems that encourage this – the "gee whiz" nature of what they can do and the fact that they don't place warfighters' lives in jeopardy.

What would you say are currently the most important ethical issues regarding the deployment and development of military robots?

Last time I counted, I had identified at least 23 distinct ethical issues to do with the use of robotic weapons – so we could talk about the ethics for a long time ... To my mind, the *most* important issue is the ethics of what Yale philosopher, Paul Kahn, has described as "riskless warfare". If you watch footage of UAVs in action it looks a lot like shooting fish in a barrel. The operators observe people in Iraq or Afghanistan, make a decision that they are the enemy, and then "boom" – they die. The operators are never in any danger, need no (physical) courage, and kill at the push of a button. It is hard *not* to wonder about the ethics of killing in these circumstances. What makes the particular men and women in the sights of the Predator legitimate targets and others not? Traditionally, one could say that enemy combatants were legitimate targets of our troops because they were a threat to them. Even enemy soldiers who were sleeping might wake up the next morning and set about

attacking you. Yet once you take all of our troops out of the firing line and replace them with robots remotely operated from thousands of kilometres away, then it is far from clear that enemy combatants pose any threat to our warfighters at all. Armed members of the Taliban might *want* to kill us but that may not distinguish them from their non-combatant supporters.

Kahn has suggested that when the enemy no longer poses any threat, we need to move from “war” to “policing”, with the justification for targeting particular individuals shifting from the distinction between combatants and non-combatants to the question of whether particular individuals are involved in war crimes at the time. I’m not sure the notion of “threat” does all the work Kahn’s argument requires, because, as the legitimacy of targeting sleeping combatants suggests, even in ordinary warfare the enemy is often only a hypothetical or counterfactual threat. Nonetheless, there does seem to be *something* different about the case in which the enemy has only the desire and not the capacity to threaten us and some of my current research is directed to trying to sort out just what the difference is.

After that, there are obvious concerns about whether unmanned systems might lower the threshold of conflict by encouraging govern-

ments to think that they can go to war without taking casualties, or by making accidental conflict more likely. There are also some interesting questions about what happens to military culture and the “warrior virtues” when warfighters no longer need to be brave or physically fit. Finally, there is an important and challenging set of issues that are likely to arise as more and more decision making responsibility about targeting and weapon release is handed over to the robots. At the moment, systems rely upon having human beings “in the loop” but this is unlikely to remain the case for too much longer; in the longer term, systems that can operate without a human controller will be more deadly and survivable than those that rely upon a link to a human controller. Eventually we will see an “arms race to autonomy” wherein control of the weapons will be handed over to on-board expert systems or artificial intelligences. A whole other set of ethical issues will arise at that point.

In passing, I might mention that one of the objections people raise most often about robot weapons – that they make it easier to kill, by allowing “killing at a distance” – seems to me to be rather weak. Crossbows allow people to kill at a distance and cruise missiles allow them to kill without ever laying eyes on their target. Operating a weapon by remote control doesn’t seem to add

anything new to this. Indeed, one might think that the operators of UAVs will be *more* reluctant to kill than bombardiers or artillery gunners because they typically see what happens to the target when they attack it.

You mentioned earlier that it is hard to see anything wrong with the use of robots for tasks like mine clearing or IED disposal. In your 2008 article, "Building a Better WarBot. Ethical Issues in the Design of Unmanned Systems for Military Applications", you go further than that and suggest that it is not just ethical to use robots but ethically mandated to do so if possible. Are there other scenarios in which you think the use of robots is morally required? Also, in that paper, you point towards the often neglected effects the use of teleoperated robots has on their operators. Is this something which should be considered more in the discussion of ethical challenges of military robots?

There is some truth to the thought, "why send a person, when a robot can do it?" Commanders should be trying to protect the lives of those they command. Thus, if a robot can do the job instead of a human being, without generating other ethical issues, then, yes, it would be wrong not to use the robot.

Of course, there are two important caveats in what I've just said.

Firstly, the robot must be capable of succeeding in the mission – and, as I've said, I think there are fewer military applications where robots are serious competitors with human warfighters than people perhaps recognise.

Secondly, there must not be other countervailing ethical considerations that argue against the use of the robot. In particular, attacks on enemy targets by robotic systems must meet the tests of discrimination and proportionality within *jus in bello*. As long as there remains a human being "in the loop", making the decision about weapon release, this need not present any special difficulty, so the use of teleoperated weapons such as the Predator will often be ethically mandated if the alternative is to put a human being in danger in order to achieve the same tasks. Suppression of enemy air defences is another case, often mentioned in the literature, where it may be wrong not to use a robot. If fully autonomous weapons systems are involved, however, the balance of considerations is likely to change significantly. Except in very specific circumstances, such as counter-fire roles, wherein it is possible to delineate targets in such a way as to exclude the possibility of killing non-combatants, these weapons are unlikely to be capable of the necessary discrimination. Moreover, with both sorts of weapons there may be other ethical issues to

take into account, which might make it more ethical to send a human warfighter.

It is also worth keeping in mind that the ethics of using a weapon once it exists and the ethics of developing it may be very different. We may have good reasons *not* to develop weapons that it might be ethical to use – for instance, if the development of the weapon would make war more likely.

Regarding the operators, yes, I very much believe that people should be paying more attention to the effects that operating these weapons will have – indeed, are already having – on their operators and to the ethical issues arising from them. Remotely operating a weapon like the Predator places the operator in a unique position, both “in” and outside of the battlespace. Their point of view and capacity for military action may be in Afghanistan, while they themselves are in Nevada. After they fire their weapons, by pressing a few controls, they “see” the bloody results of their actions. Yet they have access to few of the informal mechanisms arising out of deployment in a foreign theatre that may help warfighters process the experiences they have been through. I have heard anecdotal reports from several sources that suggest the rates of post-traumatic stress disorder in the operators of the Predator are extremely high – and it certainly

wouldn't surprise me if this was the case.

Gustav Däniker coined the term “miles protector” in 1992 after the Gulf War and summed up the new tasks of the future soldier in the slogan “protect, aid, rescue” (“Schützen, Helfen, Retten”). On the other hand there are arguments for the soldier to return to the role of the warfighter, often called the “core task” of soldiers. Do you think the shift from “war” to “policing” will have a significant impact on the self-image of soldiers and could you elaborate on your research in this matter?

I don't think increased use of robots will lead to a shift from “war” to “policing”. Rather, I am arguing that the appropriate model to use in order to think about the justification of killing people who are no threat to you is “policing”. Police need to take much more care to protect the lives of bystanders and have far fewer moral privileges in relation to killing than do soldiers during wartime. So, yes, *were* soldiers to start to take on this role, this would require a significant shift in their self-image. However, as I said, the argument I'm making about robots concerns the question of when, if ever, killing people via a robot is justified – not how often this is likely to happen. Unfortunately, I think it is much more likely that armed forces will use robots to kill people when they

shouldn't than it is that they will change the nature of the missions they engage in because they have these new tools.

Robots are of limited use in the sorts of peace-keeping and peace-enforcement missions that Däniker had in mind when he coined the phrases you mention. However, they do clearly have their place. Avoiding casualties may be especially important when governments cannot rely upon public support for taking them because the national interest is not at stake. Mine-clearing and bomb disposal are often an important way of winning the support of local populations – and robots can play a role here. The sort of surveillance that UAVs can provide is clearly a vital asset if one's goal is to prevent conflict and keep enemies apart. To the extent that armed UAVs can attack targets more precisely, with less risk of unintended deaths, they may also contribute to the success of peace enforcement missions. However, ultimately success in these sorts of deployments will depend upon talking with local people and on building trust and relationships on the ground. Robots have nothing to contribute to this goal and may even get in the way of achieving it – if, for instance, commanders' access to intelligence from UAVs prevents them from seeking human intelligence, or if the robots function in

practice to isolate and alienate troops from the local population.

On the other hand, I do think the use of robotic weapons has the potential to radically unsettle the self-image of soldiers ... if not along the lines you suggest. For instance, there is no need for war-fighters to be courageous – at least in the sense of possessing physical courage – if they will be operating weapons thousands of miles away; nor need they be especially fit or even able-bodied. There can be no argument that women should not take on “combat” roles operating robots, as physical strength is irrelevant in these roles, as is vulnerability to sexual assault (I'm not saying these were ever good arguments – just that it is especially obvious that they have absolutely no validity in this circumstance). It is hard to see how notions of “comradeship” apply when troops involved in the same battle – or even in the same unit – may be in completely different locations. It is not clear that one can really display mercy by means of a robot: one might refrain from slaughtering the enemy but this in itself is not sufficient to demonstrate the virtue of mercy. Indeed, there are whole sets of virtues and character traits currently associated with being a good “warrior” that may be completely unnecessary – or even impossible to cultivate – if one's role is operating a robot.

Of course, it has always only been a minority of those serving in the armed forces who needed to be brave, resolute, physically fit, etcetera, and we are a long way yet from being able to replace significant numbers of frontline troops with robots. Yet it is clear that there is a real tension between the dynamics driving the introduction of unmanned systems and the traditional function and self-image of soldiers. Eventually, I suspect, this will cause real problems for military organisations in terms of their internal cultures and capacity to recruit.

Since the St Petersburg Declaration of 1868 there have been various initiatives to restrict the use of weapons which cause unnecessary suffering. Do you think there is a need for additional international legislation to regulate the development and deployment of robots by the military? If so, what could be brought forward in favour of such legislation?

I definitely think we should be working towards an international framework for regulating the development and deployment of military robots – although perhaps not for the reason you suggest nor by the means you suggest.

I haven't seen any reason yet to believe that the use of robots will cause unnecessary suffering in the way that, for instance, nerve gas or

dum dum bullets arguably do. Nor will robots necessarily kill any more people than the weapons and systems they will replace.

The reason to be worried about the development of more and more sophisticated robotic weapons is that these systems may significantly lower the threshold of conflict and increase the risk of accidental war. The fact that governments can attack targets at long distances with robotic weapons without risking casualties may mean that they are more likely to initiate military action, which will tend to generate more wars. I think we have already seen this effect in action with the use of the Predator in Pakistan and northern Africa. If robotic weapons begin to be deployed in roles with "strategic" implications – for instance, if the major powers start to place long-range and heavily armed uninhabited aerial vehicles or unmanned submersibles on permanent patrol just outside the airspace or territorial waters of their strategic rivals – then this will significantly decrease the threshold of conflict and increase the risk of accidental war. If fully autonomous weapons systems enter into widespread use then this will put a trigger for war into the hands of machines, which might also increase the risk of accidental war.

So, yes, there are very good reasons to want to regulate the development of these weapons. However,

for pragmatic reasons to do with the likelihood of reaching agreement, I think it might be better to approach this as a traditional case for arms control, with bilateral or regional agreements being a priority, perhaps with the ultimate goal of eventually extending these more widely. It is hard to see the United States or Israel, which have a clear lead in the race to develop robotic weapons, accepting restrictions on the systems until it is in their interests to do so. Yet if their strategic competitors become capable of deploying weapons that might pose a similar level of threat to them then they might be willing to consider arms control. Concerns about the threshold of conflict and risk of accidental war are familiar reasons to place limits on the number and nature of weapons that nations can field. As I argue in a recent paper, "Predators or Ploughshares?", in *IEEE Technology and Society Magazine*, a proper arms control regime for robotic weapons would need to govern: the range of these weapons; the number, yield, and range of the munitions they carry; their loiter time; and their capacity for "autonomous" action. If we could achieve one or more bilateral agreements along these lines it might then be possible to extend them to a more comprehensive set of restrictions on robotic weapons, perhaps even in the form of international law. I suspect we are a long way from that prospect at this point in time.

When it comes to the attribution of responsibility for the actions of military robots you have suggested an analogy between robots and child soldiers. Could you elaborate on this?

It is important to clarify that I was writing about cases in which it might be plausible to think that the robot "itself" made the decision to kill someone. There are actually three different scenarios we need to consider when thinking about the responsibility for killing when robots are involved.

The first is when the "robot" is a remote control or teleoperated device, as is the case with Predator and other UAVs today. In this case, it is really the human being that kills, using the device, and the responsibility rests with the person doing the killing.

The second is where the robot is not controlled by a human being but reacts to circumstances "automatically" as it were, as it would if it were controlled by clockwork or by a computer. In this case, the appropriate model upon which to conceptualise responsibility is the landmine. While there is a sense in which we might say that a landmine "chose" to explode at some particular moment, we don't think that there is any sense in which the moral responsibility for the death that results rests with the mine.

Instead, it rests with the person who placed the mine there, or who ordered it to be placed there, or who designed it, etcetera. This model remains appropriate to robots, even if the robot contains a very sophisticated onboard computer capable of reacting to its environment and tracking and attacking various targets, etcetera, – as long as there is no question that the robot is a machine lacking consciousness and volition. When computers are involved it may be difficult to identify which person or persons are responsible for the “actions” of the machine. However, it is clear both that the question of responsibility will be no different in kind to others that arise in war due to the role of large organisations and complex systems and that the appropriate solution will usually be to *assign* responsibility to some person.

A third scenario will arise if robots ever come to have sufficient capacity for autonomous action that we start to feel uncomfortable with holding human beings responsible for their actions. That is, if we ever reach the point where we want to say that the robot itself made the decision to kill someone. It’s clear that none of the current generation of military robots come anywhere near to possessing this capacity – whether they ever will depends upon the progress of research into genuine artificial intelligence.

It was this third scenario that I was investigating in my article on “Killer Robots”. I was interested in whether it will *ever* be possible to hold even genuine artificial intelligences *morally* responsible for what they do, given the difficulties involved in applying some of our other concepts, which are connected to responsibility, to machines – concepts such as suffering, remorse, or punishment. It seems as though there is a “gap” in the spectrum of degrees of autonomy and responsibility, wherein certain sorts of creatures – including, possibly, robots – may be sufficiently autonomous that we admit they are the origin of their actions, but not to the extent that we can hold them morally responsible for their actions. When we are dealing with entities that fall into this gap then we rightly feel uncomfortable with holding someone else responsible for their actions, yet it is hard to see what the alternative might be – unless it is to admit that no one is responsible. The latter option is not something we should accept when it comes to the ethics of war.

The use of child soldiers was the best model I could come up with to help think about this scenario. With child soldiers, you can’t really hold them morally responsible for what they do, however, nor would it be fair to hold their commanding officer morally responsible for what they do, if he or she was ordered to send

them into battle. Even the person who conscripts them seems to be responsible for *that* rather than for what the children do in battle. One – though not necessarily the most important – of the reasons why using child soldiers in warfare is unethical, then, is that they may cause deaths for which no one may properly be held responsible. I think there is a similar danger if we ever reach the point where we would be willing to say that robots were really making the decision as to who should live or die ...

Though it is still disputed whether there will be ever something like a genuine artificial moral agent, it seems clear that artificial intelligence in military robots will continually improve and the roles of military robots will expand in future armed conflicts. So if robots gradually enter this third scenario – being sufficiently autonomous that they are the origin of their actions but not such that we can hold them morally responsible for their actions – how could this be integrated in the existing ethics of war? And is “keeping the human in the loop” – which the military always insist they will do, whenever these weapons are mentioned – a serious and plausible possibility?

The answers to your two questions are closely connected. Let me begin with your second question because it is, perhaps, slightly easier to an-

swer and because the answer to this question has important implications for the answer to your first question.

We could *insist* upon keeping human beings in the loop wherever robots are used but this could only be sustained at a high cost to the utility of these systems – and for that reason I think it is unlikely to happen, despite what military sources say today. The communications infrastructure necessary to keep a human being in the loop is an obvious weak point in unmanned systems. In the longer term, the tempo of battle will become too fast for human beings to compete with robots. For both these reasons, the military is eventually likely to want to field systems that are capable of operating in “fully autonomous” mode: if an arms race to build robotic weapons should develop, then nations may have little choice *but* to field autonomous weapons. Moreover, there are some potential roles for unmanned systems, such as long-range anti-submarine warfare or “stealthed” precision air strikes, where it simply will not be possible to put a human being in the loop. Yet, again, these are applications that nations in pursuit of military supremacy – or even parity – can ill afford to ignore. It is therefore a politically expedient fiction, which the military are promulgating, to insist that there will always be a human in the loop. What’s more, I

think the better military analysts know this!

The answer to your second question is therefore both “yes” and “no”. Keeping human beings in the loop *is* plausible in the sense that we could do it and – I will argue in a minute – we may have good reasons to do it. However it is *not* a serious possibility in the sense that it is not likely to happen without a concerted effort being made to achieve it.

To turn now to your first question. As far as integrating autonomous weapons systems into the ethics of war goes, I believe this will be very difficult – as my comparison with child soldiers suggests. The obvious solution, which is, I believe, the one that militaries will eventually come to adopt, is to *assign* responsibility for the consequences of the use of autonomous weapons to the person who orders their use; we might think of this as insisting that the commander has “strict liability” for any deaths that result. However, the question then arises as to whether or not this is fair to the military officers involved? Commanders are currently held responsible for the activities of the troops they command but this responsibility is mitigated if it can be shown that individuals disobeyed their orders and the commander took all feasible steps to try to prevent this. Where this occurs, the moral re-

sponsibility for the troops’ actions devolves to the troops themselves. It is this last step that will be impossible if it is machines that have “chosen” to kill without being ordered to do so, which is why we may need to insist upon the strict liability of the commander. However, this means there is a risk the commander will be held responsible for actions they could not have reasonably foreseen or prevented. I must admit I also worry about the other possibility – that no one will be held responsible.

If we do begin using autonomous weapons systems with something approaching genuine artificial intelligence in wartime, then we must insist that a human being be held responsible for the consequences of the operations of these weapons at all times – this will involve imposing strict liability. The alternative would be to reject the use of these systems and to insist upon keeping a human being in the loop. However, as I’ve said, there are many dynamics working against this outcome.

I should mention that another alternative that has received a significant amount of attention in the literature and the media recently – that we should “program ethics” into the weapon – is to my mind an obvious non-starter. Ron Arkin at Georgia Tech has recently published a book advocating this. However, with all respect to Ron, who

was extremely kind to me when I visited him at Georgia Tech, this is a project that could only seem plausible as long as we entertained a particularly narrow and mechanical view of ethics.

It *will* undoubtedly be possible to improve the capacity of robots to discriminate between different categories of targets. Moreover, there are, perhaps, some categories of targets that it will almost always be ethical to attack. John Canning, at the US Naval Surface Warfare Centre, is very keen on the idea that autonomous weapons systems might be programmed to attack only those holding weapons or even to attack only the weapon system, thereby disarming the enemy.

However, even if it is possible to build such systems there is a real possibility of deadly error. The proper application of the principles of discrimination and proportionality, which largely determine the ethics of using lethal force in wartime, is extremely context dependent. Even if the potential target is an enemy Main Battle Tank – which you'd normally think it would be okay to attack – whether or not this is ethical in any particular case will depend on context: whether the enemy has surrendered, or is so badly damaged as to no longer pose a threat, or has recently started towing a bus full of school children. More generally, assessments of

when someone or something is a legitimate military target will often depend on judgements about the intentions of the enemy, which in turn need be informed by knowledge of history and politics. Robots don't have anywhere near the capacity to recognise the relevant circumstances, let alone come to the appropriate conclusions about them – and there is no sign that they are likely to have these for the foreseeable future. So even the idea that we could rely upon these systems to be capable of discrimination seems to me a fantasy.

When it comes to the idea that they could actually reason or behave ethically, we are even more firmly in the realm of science fiction. Acting ethically requires a sensitivity to the entire range of human experience. It simply isn't possible to "algorithmise" this – or at least no philosopher in human history has been able to come up with a formula that will determine what is ethical. I would be very surprised if any engineer or computer scientist managed to do so!

You mentioned at the outset that your early research was about non-military robots. Before we finish, can we talk about that for a moment? Do you have any thoughts on the use of robots more generally, their impact on society, and their possible influence on interpersonal relations? I know that people are

talking about a future for robots in the entertainment and sex industries and that you have written about the ethics of using robots in aged care settings. Should we be looking forward to the development of robot pets and companions?

I think it's highly improbable that robots will have much influence on society or interpersonal relations for the foreseeable future – mostly because I think it is unlikely that robots will prove to be useful in our day-to-day lives anytime soon. Since the 1950s at least, people have been talking about how we would soon have robots living and working alongside us. I am still waiting for my robot butler!

There are some pretty straightforward reasons for the absence of any useful robots outside of very specific domains, although they are often ignored in media discussions of the topic. Humans are complex and unpredictable creatures, which makes us hard for robots to deal with. In order for robots to be able to perform useful roles around the home or in the community, they would need to be large, which means they will be heavy and therefore dangerous, and extremely sophisticated, which means they will be expensive and difficult to maintain. For all these reasons, robots and humans don't mix well and in domains where robots do play a significant role,

such as manufacturing, this has been made possible by keeping robots and people apart.

Bizarrely, war turns out to be a relatively friendly environment for robots. Killing someone, by pointing and firing a weapon at them, is a much easier task for a robot than helping them is. War is also a domain in which it is plausible to think one might be able to reliably separate those humans we don't want to place at risk of injury from the robots that might injure them through the simple expedience of ordering the human beings to stay clear of the robots. This also has the virtue of protecting the robots. Budgets for "defence" spending being what they are, military robots can be very expensive and still profitable to sell and manufacture. "Domestic" robots would have to compete with underpaid human carers and servants, which makes it much tougher to make them commercially viable. There is, admittedly, more room for the development of more-and-more sophisticated robotic toys, including sex toys, but I think we are a long way from the point where these will start replacing relations between people or between people and their (real) pets.

None of this is to say that I don't think there are ethical issues associated with the attempt to design robots for these roles. Designing robots so that people mistake them

for sentient creatures involves deception, which may be problematic. Thinking it would be appropriate to place robots in caring roles in aged care settings – or even to use them to replace human workers, such as cleaners, who may be some of the few people that lonely older people have daily contact with – seems to me to involve a profound lack of empathy and respect for older people.

I *am* looking forward to seeing more robots. Robots are cool! I think the engineering challenges are fascinating, as is what we learn about the problems animals and other organisms have solved in order to live in the world. However, we should remember that engineers want to – and should be funded to – build robots because of the challenges involved and that often the things they are required to say nowadays to secure that funding involve them moving a long way outside of their expertise. As soon as people start talking about real-world applications for robots, the most important things to consider are facts about people, societies, politics, economics, etcetera. These are the things that will determine

whether or how robots will enter society. Indeed, it has always been the case that when people appear to be talking about robots, what they are mostly talking about is human beings – our values, our hopes and fears, what we think are the most pressing problems we face, and what sort of world we want to live in. This is one of the reasons why I chuckle whenever I hear anyone talking about Asimov’s “three laws of robotics” as though these were a serious resource to draw upon when thinking about how to build ethical robots. Asimov was writing about people, not robots! The robots were just devices to use to tell stories about what it meant to be human.

The fact that human beings build – and talk about – robots to satisfy and amuse other human beings means that the most important truths about robots are truths about human beings. When it comes to talking about the future of robotics, then, you would often do just as well – or even better – talking to a philosopher or other humanities scholars rather than to an engineer or roboticist.

References

- Sparrow, R. 2004. “The Turing Triage Test.” *Ethics and Information Technology* 6(4): 203-213.
- Sparrow, R. 2002. “The March of the Robot Dogs.” *Ethics and Information Technology* 4(4): 305-318.
- Sparrow, R. 2007. “Killer Robots.” *Journal of Applied Philosophy* 24(1): 62-77.

- Kilcullen, David, and Andrew McDonald Exum. 2009. "Death from above, Outrage Down Below." *New York Times*, May 17, WK13.
- Kahn, Paul W. 2002. "The Paradox of Riskless Warfare." *Philosophy & Public Policy Quarterly* 22(3): 2-8.
- Sparrow, R. 2009. "Building a Better WarBot : Ethical issues in the design of unmanned systems for military applications". *Science and Engineering Ethics* 15(2):169–187.
- Daniker, Gustav. 1995. *The Guardian Soldier. On the Nature and Use of Future Armed Forces*. UNIDIR Research Paper No. 36. New York and Geneva: United Nations Institute for Disarmament Research.
- St Petersburg Declaration 1868. Declaration Renouncing the Use, in Time of War, of Explosive Projectiles Under 400 Grammes Weight. Saint Petersburg, 29 November /11 December 1868. Available at <http://www.icrc.org/IHL.NSF/FULL/130?OpenDocument>.
- Sparrow, R. 2009. "Predators or Plowshares? Arms Control of Robotic Weapons". *IEEE Technology and Society* 28(1): 25-29.
- Arkin, Ronald C. 2009. *Governing Lethal Behavior in Autonomous Systems*. Boca Raton, FL: Chapman and Hall Imprint, Taylor and Francis Group.
- Canning, John S. 2009. " You've Just Been Disarmed. Have a Nice Day!". *IEEE Technology and Society* 28(1): 12-15.
- Sparrow, R., and Sparrow, L. 2006. "In the hands of machines? The future of aged care." *Minds and Machines* 16:141-161.