

## Peter Asaro: Military Robots and Just War Theory

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

When I was writing my dissertation on the history of cybernetic brain models and their impact on philosophical theories of the mind, I became very interested in the materiality of computation and the embodiment of mind. From a technological perspective, materiality had a huge impact on the development of computers, and consequently on computational theories of mind, but this material history has been largely ignored, perhaps systematically to make computation seem more like pure mathematics.

During this time, I was asked to write a review of a book by Hans Moravec, about robots with human-level cognition, which made some pretty wild speculations based on the notion that cognition was a purely Platonic process that would someday escape its materiality. For instance, the idea that computational simulations might become just as good as real things if they were complicated enough, and contained enough detail and data. It seemed to me that this missed the

role of material processes in cognition and computation.

This led me to start thinking about explicitly material forms of artificial cognition, more specifically robots as computers with obvious input-output relations to the material world. Pretty soon I was making a documentary film about social and emotional robotics, *Love Machine* (2001), which explored how important embodiment is to emotions like love and fear, and how roboticists were seeking to model these and what it would mean to build a robot that could love a person.

Because of that film, a few years later I was invited to write a paper on "Robot Ethics." In researching that paper, I came across Colin Allen and Wendell Wallach's work on artificial moral agents, I was struck again by a sense that embodiment and materiality were not getting the attention they deserved in this emerging field. It seemed to me that the goal of robot ethics should not be to work out problems in ethics using computers, but to actually figure out ethical rules and policies for how to keep real robots from doing real harm to real people. The most obvious place where such

harm might occur, and thus ethical considerations should arise, also turns out to be the area of robotics research that is receiving by far the most funding: military applications. The more research I did on the state-of-the-art of military robotics, the more I realized that this was a social and political issue of great importance, as well as one of philosophical interest. So I pursued it.

*In the last couple of years, how did philosophy as a professional field adjust to the intensified development and deployment of artificial intelligence, robots in general and of unmanned systems by the military in particular? As a philosopher yourself, in your personal opinion, how should and how could philosophers contribute to the debates in this field?*

I would say that as a professional field, I am a bit disappointed that philosophy has not had a better organized response to the rise of technology in general, and the intensified development and deployment of AI and robots in particular. While there are some good people working on important issues in these areas, there are only a handful of groups trying to organize conferences, workshops and publications at the intersection of philosophy and real-world computing and engineering. Especially compared to other subfields like medical ethics, bio-ethics, neuro-ethics, or even

nano-ethics, where there seems to be more funding available, more organizations and institutes, and more influence on the actual policies in those areas. But information ethics has been getting traction, especially in the areas of information privacy and intellectual property, so perhaps robot ethics will start to catch up in the area of military robotics. It is still a small group of people working on this problem, and most of them seem to be on your interview list.

In my opinion, philosophers can make significant contributions to the debates on the use of military robotics. Philosophers are often accused of navel-gazing and irrelevance, whereas the development and use of lethal military robotics presents philosophically interesting problems with pressing real-world relevance. So this issue has the potential to make philosophy more relevant, but only if philosophers are willing to engage with the real-world complexity of the debate. And doing so can be fraught with its own moral and ethical issues – you have to consider if your own work could be used to justify and rationalize the development of some terrible new weapon. The theoretical work requires a great deal of intellectual integrity, and the policy work requires a great deal of moral sensitivity. I think these are the traits of a good philosopher.

*A lot of people think about military robots and unmanned systems merely in technological categories. Why do you think it is necessary to broaden the approach and to stress ethical and philosophical aspects if machines are to be developed and used in military contexts?*

Part of the reason that military robots snuck up on us so quickly, despite the warnings from science fiction, is that in many ways they are only small technological steps beyond military systems that we already know and accept in modern warfare. The initial strategy to call these systems into question is to argue that “autonomy” is a critical disjunction, a qualitative leap, in the evolution of military robots. But I do not think it is necessary to make that argument in order to question the morality of using robotics. In fact, my most recent article focuses on the morality of tele-operated robotics. Rather, I think we can look at the history of military strategy and technology, especially beginning in World War I and continuing through the Cold War and the Global War on Terror, and see how our generally accepted views of what is ethical in war have evolved along with new technologies. It is not a very flattering history, despite the fact that most officers, soldiers and engineers have made concerted efforts to make ethical choices along the way.

In my view, the critical ethical issues are systemic ones. We will not have more ethical wars just because we have more ethical soldiers, or more ethical robots. First of all, this is because there will always be a fundamental question of whether a war is just or not. The moral justification for developing and amassing military power will always depend upon the morality of the group of individuals who wield that power and how they choose to use it (Just War theorists call this *jus ad bellum*).

Second of all, warfare is a cultural practice. While it is cliché to say that warfare has been around as long as humans have (or even longer among other animals, perhaps), it is important to note that how wars are fought is built upon social, cultural and ethical norms that are very specific to a time and a culture. Over the last two centuries, warfare has become increasingly industrialized, subjected to scientific study, and made increasingly efficient. One result of those efforts is the incredibly sophisticated weapons systems that we now have. On the one hand, it is not necessary that efficiency should be the highest value – nations could have pursued honour, chivalry, valour, glory, or some other values as the highest, and then warfare would look different now. On the other hand, efficiency alone is not sufficient to win a war

or control a population because there is a huge socio-psychological element as well – which is why we have also seen militaries develop and deploy media and communication technologies, as well as rhetoric and propaganda, to shape people's perceptions and beliefs. Even if we believe Machiavelli when he advises his prince that it is better to be feared than loved, fear is still a psychological phenomenon, and even the most ruthless and technologically advanced tyranny could not maintain itself without sufficiently aligning the interests of the people with its own. There are numerous examples of great and mighty militaries that have successfully destroyed the military forces of their enemies, but ultimately failed to conquer a territory because they failed to win the "hearts and minds" of those who lived there. Which is just another way of saying that warfare is a cultural practice. Of course, there are also many examples of conquerors simply trying to eliminate the conquered peoples, and the efficiency of modern weapons makes genocide more technically feasible than it was historically. Robot armies could continue this trend to terrible new levels, allowing even smaller groups of people to dominate larger territories and populations, or commit genocides more quickly and with fewer human collaborators. Hannah Arendt argued that because of this, robot

armies are potentially more insidious than atomic weapons.

If we want to take a step back from history, and the question of why we have come to a place where we are building lethal military robots, we can ask how we should build such robots, or whether we should build them at all, or what we should be building instead. So from a strategic point of view, the US might undermine support for terrorists more efficiently through aid programs to places where terrorism thrives due to poverty, than they would by putting those funds towards demonstrating their military superiority. We can also ask what values a nation is projecting when they commit such vast amounts of time and resources to fighting a war by remote-control, or with autonomous robots. Having received your questions just after the 40<sup>th</sup> anniversary of the Apollo 11 moon landing, I am reminded that despite its being a remarkable event in human history, it only occurred because of the specific history of the Space Race as a competition between the ideologies of the Cold War. In that case, the US scored a symbolic victory in technological achievement by landing a man on the moon, but it was also about projecting values of ingenuity, technological sophistication and teamwork. The US also spent a vast amount of mental and monetary resources in achieving that goal. In the case of military robotics,

I think it is a philosophical question to ask what values are being promoted and projected by these technologies, and if those are the values society ought to be pursuing above others. If we want to project technological prowess and pragmatic ingenuity, this could also be done through developing technologies, public works, aid, and environmental projects that ameliorated the underlying social, political and resource problems.

*Contrary to most of the media coverage, the unmanned systems deployed by the military today are in general mostly tele-operated (though including some autonomous functions or potential) but not fully autonomous. In your last article for the IEEE Technology and Society magazine<sup>1</sup> you were specifically pointing out the importance of ethical considerations regarding these systems, which rely on human decision making and analyzed three different approaches to the design of these systems. Could you elaborate on that?*

In that paper I was approaching the ethics of tele-operated lethal military robots as a problem in engineering ethics. That is, I wanted to ask what it would mean to actually design such a system “ethically.” Mary Cummings, a former Navy combat pilot who now teaches interface design at MIT, has taken a similar approach. She calls her approach

“value-centered design” and the idea is to have engineers brainstorm about potential ethical or safety issues, establish sets of values that should be design goals (like limiting civilian deaths), and then to actually evaluate and compare the alternative system designs according to those values. Another view proposed by Ron Arkin (actually for autonomous robots but it could be applied to tele-operated robots as well) is that of the “ethical governor.” Basically, this is a system which follows a set of rules, like the Laws of Armed Conflict and Rules of Engagement, and stops the robot if it is about to commit a war crime or an atrocity. This approach assumes that you can develop a set of rules for the robot to follow which will guarantee it does nothing unethical on the battlefield.

The problem with both of these approaches is that they see values and ethical rules as black boxes. It is as if we can simply program all the ethical rules and make the robot follow them without considering the context in which ethical decisions are made. However, in real-world moral and ethical decision-making, humans *deliberate*. That is, they consider different perspectives and alternatives, and then decide what is right in a given situation. Am I really more ethical because my gun will not fire when I point it at innocent people, or am I just less likely to shoot them? I think that if we

really want to make robots (or any kind of technology) more ethical, we should enhance the ethical decision-making of the people that operate them. The paper then asks: What would it mean to build technologies that actually do that? I propose a “user-centered approach,” which seeks to understand how people actually make ethical decisions, as an information-processing problem. What kind of information do people actually use to make these lethal decisions? What roles do emotion, empathy, and stress play? We really do not understand these things very well yet, but I think the answers might surprise us, and might also lead to the design of technological systems which actually make it harder for people to use them unethically because they are better informed and more aware of the moral implications of their use of the system.

*So if I understand you correctly, instead of equipping the user with an artificial ethical governor, you would prefer to “equip” the user with ethical values and understanding and leave the actual decision-making in the human sphere. This would be similar to the “keep the human in the loop” approach, which has also been put forward by some people in the militaries. On the other hand, especially the amount of information to be processed in shorter and shorter time by the human operator/supervisor of mili-*

*tary systems is likely to increase beyond the capacity of the human physique, which might offer an advantage to systems without human participation. Do you think that this user-centered approach (and similar matters) could be regulated by international legislation, for example a ban on all armed autonomous systems without human integration of decision-making?*

The short answer is: Yes, we should seek an international ban on all autonomous lethal systems, and require all lethal systems to have significant human involvement in the use of lethal force. Just what “significant human involvement” might mean, and how to make that both technologically effective and politically acceptable to potential participants to a treaty is a matter for discussion. Sure, there are questions about how to implement and enforce such a treaty, but just having an international consensus that such systems are immoral and illegal would be a major step.

I think we should strive to keep the human in the loop both because this clarifies moral responsibility in war, and because humans are already very sophisticated ethical information processing systems. Information technologies are quite plastic and can be developed in a variety of ways depending on our goals and interests. What I am suggesting is that instead of trying to

formalize a set of rules for when it is OK for a robot to kill someone, and build that into a robot as a black-box module, that as an ethical engineer one might instead invest technological development resources into improving the lethal decision-making of humans.

I have heard various versions of the argument that there is too much information, or not enough time, for humans to make the necessary decisions involved, and so there is, or soon will be, a need to automate the process. For instance, those who supported the “Star Wars” Strategic Defense Initiative argued that human reaction times were not sufficient to react to a nuclear assault, and so the missile defense system and retaliation should be made fully automatic. But while our intuitions might be to accept this in a particular high-risk case, this is actually a misleading intuition. If that particular case is highly improbable, and there are many potential high-risk system malfunctions with having such an automated system, then the probability of catastrophe from malfunction could be much higher than from the case it is designed to defend against. I think we are better off keeping humans in the loop and accepting their potential fallibility, as opposed to turning our fate over to an automated system that may have potentially catastrophic failures.

The mistaken intuition comes from the fact that you can justify all sorts of things when the fate of the whole world (or all of humanity, or anything of infinite or absolute value) is at stake, even if the probabilities are vanishingly small compared to the risks you incur from the things you do to avoid it. There is much more to the debates about keeping humans in the nuclear loop, particularly in nuclear deterrence theory, and in training simulations where many people (not aware it is a simulation) do not “push the button” when ordered to. I bring up this example because the history of this kind of thinking continues to have a huge influence on military technology and policy well after the end of the Cold War. While in the case of nuclear war the decisions may result in the end of civilizations, in robotic war the decisions may only result in the end of tens or hundreds human lives at a time (unless you are worried about robots taking over). The stakes are smaller, but the issues are the same. The differences are that our intuitions get distorted at the extremes on the one hand, and on the other hand that because the decision to kill one person on a battlefield where so many already die so senselessly does not seem like much of a change, so we might be seduced into accepting autonomous lethal robots as just another technology of war. For robotic systems, our intuition might

be to accept autonomous lethal robots with some kind of built-in safety system, or even believe that they might be “better” than humans at some decision-making task. However, the real risks of building and deploying such systems, and their negative long-term effects on strategy and politics, are probably much higher than the safety gains in the hypothetical design cases, but we just do not have any easy way to measure and account for those systemic risks.

I rather like Arkin’s concept of the ethical governor for robots, actually, and think it is compatible with keeping humans in the loop. My disagreement is with his argument that such a system can outperform a human in general (though for any well-defined, formalized and operationalized case you can probably program a computer to do better if you work at it long enough) because the real world will always present novel situations that are unlike the cases the robot is designed to deal with. The basic idea for the ethical governor is for it to anticipate the consequences of the robot’s actions, and to override the planned actions of the robot whenever it detects that someone will be wrongly killed as a result. That could be used as a safety mechanism that prevents humans from making mistakes by providing a warning that requires an override. Moreover, when we look at the

current situation, and see that humans do far better than robots when it comes to ethical decision making, why are we investing in improving robot performance, rather than in further improving human performance?

Besides, if we really want to automate ethical decision-making, then we need to understand ethical decision-making, not just in theory but empirically. And so I argue that the first step in user-centered design is to understand the ethical problems the user faces, the cognitive processes they employ to solve those problems, and to find out what kind of information is useful and relevant, so that we can design systems that improve the ethical decision-making of the people who operate these lethal systems. I call this “modelling the moral user.” If part of the problem is that there is too much information, that just means that we need to use the technology to process, filter and organize that information into a form that is more useful to the user. If part of the problem is that users do not know how much to trust or rely upon certain pieces of information, then the system needs to make transparent how and when information was obtained and how reliable it is. These are questions that are important both philosophically, as matters of practical epistemology and ethics, and from an engineering perspective.



*In the last couple of years unmanned systems were deployed and used by the US Armed Forces in considerable numbers, e.g. in Afghanistan and Iraq, and are becoming a more and more common sight in and above the operational areas. With the ongoing developments, the ethical and legal debate on the deployment of robots as military (weapon) systems has intensified. From your point of view, what should be the main considerations regarding the Law of Armed Conflict and Just War Theory?*

There are several crucial areas of concern in the Pentagon's increased adoption of robotic technology. It is hard to say what the greatest concern is, but it is worth paying attention to how military robots are already contributing to new strategies.

We should be immediately concerned at the increasing use of armed UAVs within Pakistan over the past 12 months--a policy begun under President Bush and embraced by President Obama. This policy is born out of political expediency, as a military strategy for operations in a country which the US is not at war with, nor is there any declared war.

By stating that it is a matter of political expediency I mean that the fact that these robotic technologies exist provides a means for a kind of lethal

US military presence in Pakistan which would not be possible otherwise, without either the overt consent of the Pakistani government, expand the official war zone of the Afghan war to include parts of Pakistan, an act of war by the US against Pakistan's sovereignty, or the US risking the loss of pilots or commandos in covert raids (who would not be entitled to the rights of prisoners of war under the Geneva Conventions because they would not be participating in a war). There is a lack of political will within Pakistan to allow the US military to operate freely against the Taliban within its borders (though it was recently revealed that Pakistan does allow the US to operate a UAV launching base within its borders), just as there is a lack of political will in the US to destabilize Pakistan and take responsibility for the consequences. The UAVs provide a means to conduct covert raids with reduced risks, and while these raids are publicly criticized by officials of the Pakistani government, the situation seems to be tolerated as a sort of compromise solution. Despite the recent news that a US drone has assassinated the head of the Taliban in Pakistan, I am skeptical that these UAV "decapitation" raids will make a significant impact on the military or political problems that Pakistan faces, and may do more harm than good in terms of the long-term stability of Pakistan. This is a bad precedent for international conflicts insofar as it

appears to have resulted in numerous unnecessary civilian casualties outside of a declared war zone, and moreover it seems to legitimate a grey area of covert war fought by robots (thus allowing robots to circumvent international and local laws against extra-judicial and targeted killings and kidnappings much in the way on-line casinos circumvent laws against gambling through the physical separation of an agent and their actions). It is not surprising that these missions are under the operational control of the CIA (rather than the military), and that the CIA actually outsources the arming and launching of the UAVs in Pakistan to non-governmental mercenary forces such as Blackwater/XE. So while proponents of lethal robots are invoking Just War Theory and arguing that they can design these robots to conform to its standards, we see that the most frequent use of lethal robots today, in Pakistan, falls completely outside the requirements of Just War Theory because there is no war, and the military is not even pulling the trigger precisely because it is illegal for them to do so.

However, it should be noted that in Afghanistan the civilian casualties have been far greater in airstrikes from conventional aircraft and from commando raids, than from UAVs. I believe this is probably due to the fact that the Predator UAVs are only armed with Hellfire missiles, which are fairly accurate and rela-

tively small compared to the large guided bombs dropped by conventional aircraft (but are now carried by the recently deployed Reaper UAVs), and because there have been comparatively fewer armed UAV missions so far. Commando raids probably have higher civilian casualty rates in part because the commandos have a strong interest in self-preservation and are much more vulnerable than aircraft (manned or unmanned), and due to the particular circumstances in Afghanistan – where nearly every household keeps guns and often military assault rifles for home-defense, and the natural reaction to gunfire in the streets is to come out armed with the house-hold gun. When those circumstances are combined with Rules of Engagement that allow commandos to kill civilians presenting a threat by carrying guns, it is not surprising that many civilians who support, or at least have no interest in fighting against, the US forces wind up getting killed in such raids. So while on the one hand we might want to argue that UAVs could reduce civilian casualties in such raids, we could also ask the systemic question of whether such raids are an effective or ethical strategy at all or, as some have argued, are really a tactic posing as a strategy. The Dutch military forces in Afghanistan have developed a very different strategy based on a community-policing

model, rather than a surgical-strike model, though unfortunately it is not being used in all regions of the country.

Ultimately, the situations in both Afghanistan and Pakistan require political solutions, in which the military will play a role, but even the most sophisticated robotic technologies imaginable cannot improve the situation by military means alone. So I think it is also a philosophical question to ask whether military technologies are being used in ways that actually work against, or merely postpone, addressing and solving the underlying problems.

In the near term of the next decade, I think the primary concern will be the proliferation of these technologies to regional conflicts and non-government entities. UAVs are essentially remote-controlled airplanes, and the ability to obtain the basic technologies and arm them is within the grasp of many organizations, including terrorists and other non-state actors. This is also being coupled with a trend towards unconventional, asymmetric war, and organized violence and terrorism which we often call "war" but actually falls outside the purview of Just War Theory and international law. *Al Qaeda* may be waging a campaign of international violence with political aims, but they are not a nation fighting a war for political control of

a geographic territory. President Bush decided to call it a war and to use the military to fight *Al Qaeda*, and that decision has created other problems with treating members of *Al Qaeda* as prisoners of war, and putting them on trial for crimes, *etc.* So even if we have an international treaty that bans nation-states from building autonomous lethal robots, we will still face a challenge in preventing individuals and non-state organizations from building them. Of course, an international ban would dissuade the major military technology developers by vastly shrinking the potential economic market for those systems, which would greatly slow their current pace of development. Everyone would still be better off with such a ban, even if some systems still get built illegally. It will be much easier for small terrorist groups to obtain these technologies once they have been developed and deployed by militaries all over the world, than for them to try to develop these technologies themselves.

In the coming years we need to be vigilant of the Pentagon's efforts to make various robotic systems increasingly autonomous. Even autonomous self-driving cargo trucks have the potential to harm civilians, but obviously it is the armed systems that should be watched most closely. The current paradigm of development is to have

a single soldier or pilot controlling multiple robotic systems simultaneously through videogame-like interfaces. While this reduces personnel requirements, it also leads to information overload, confusion, mistakes, and a technological “fog of war.” This may actually increase the pressure to make robotic systems fully autonomous, with engineers arguing that robots will actually perform better than humans in high-stress lethal decision making.

In the long term we need to be very concerned about allowing robotic systems to make autonomous lethal decisions. While there are already systems like Phalanx and Patriot that do this in limited ways, they are often confused by real-world data. In two friendly-fire incidents in 2003, Patriot missile defense systems operating in an automatic mode mistook a British Tornado and an American F-18 as enemy missiles and shot them down. Of course, we can design clever control systems, and improved safeguards, and try to prevent such mistakes. But the world will always be more complex than engineers can anticipate, and this will be especially true when robots engage people face-to-face in counter-insurgency, urban warfare, and security and policing roles (domestic as well as military). To distinguish someone fearfully defending their family from someone who represents a genuinely organized military threat is incredibly

complicated – it depends on social, cultural and linguistic understanding that is not easily formalized as a set of rules, and is well beyond our technological capabilities for the foreseeable future. We need to be vigilant that such systems are not put in service without protest, and we should act now to establish international treaties to ensure that such systems are not developed further.

Interpreting and applying the Laws of Armed Conflict (LOAC) and developing Rules of Engagement (ROE) involve legal, political and military considerations. Because they have the potential to overwhelm individual ethical choices, or the ethical designs of robots, these interpretive processes ought to be open to critical investigation and public discussion. Arkin is confident that we can build the LOAC and ROE into the robots, but I think there are some problems with this. First, robots will not be able to do the interpretive work necessary to apply the rules to real-world situations. So what is really being put into the robots is an interpretation already made by system designers, built upon numerous assumptions and engineering considerations which may not work out in the real world. Second, sometimes the ROE are vague, confusing, or even inconsistent, and humans do not always understand when or how they

should be applied, so I cannot see how robots could do better.

Apart from the practical concerns of the technologies currently being developed, we should also be concerned about the shift in the philosophy of warfare they represent. The trend is to remove soldiers from the battle. While this is certainly good for their safety, it comes at a cost to the safety of others – in particular civilians on both sides of the conflict. The psychological distance created by remote-control or automated warfare serves to diminish the moral weight given to lethal decisions. It also serves to turn soldiers into civilians in that they start fighting wars from computer terminals in air-conditioned rooms miles away from the battle. As such it lends credence to terrorists who would claim civilians as legitimate targets. If you look at the wars that the US has been involved in over the last century, you see that as the military technology advances, the overall ratio of civilians to soldiers killed has also increased. And that is despite the wide-spread use of so-called “smart” weapons in Iraq. So while we are making war safer for soldiers, we are not really making it safer for civilians. We should be very concerned about the tendency of new military technologies to shift the risks from soldiers to civilians, as this can actually undermine the possibility of a “just war” even as the new technologies are being called “smart” or “ethical.”

*Concerning the ability of discrimination, it has been brought forward, that on the one hand artificial intelligence and sophisticated sensors could be more capable in performing this task than any human. And on the other hand that it would not even be necessary for autonomous systems to excel in the distinction of combatants/non-combatants but it would be sufficient if they equalled their human counterparts. Regarding Just War Theory, is this a maintainable argument and how would you review these and similar approaches?*

Discrimination is a crucial criterion for Just War Theory, and it has been argued that automated systems might perform better than humans at the discrimination task. I think the question is: If we accepted that automated systems could outperform humans, or if we were actually presented with evidence that some system could perform the discrimination task at or above human levels, is that a good argument for allowing them to make autonomous lethal decisions? The short answer is: No.

First, discrimination is necessary but not sufficient for ethical killing in war. The point of the discrimination criterion is that it is never acceptable to intentionally kill innocent civilians, or to kill people indiscriminately in war. This does not imply that it is always acceptable to kill

enemy combatants (except, it is argued, in “total war” though I do not accept that argument). The way it is usually construed, combatants have given up their right not to be killed by putting on a uniform. Even under this construal, it is immoral to unnecessarily kill enemy combatants. For instance, killing retreating soldiers, especially just before a clearly immanent final victory or surrender, is generally viewed as immoral, though it is legal under international law. According to a rights-based view of Just War Theory, it is necessary for enemy combatants to also present an actual threat in order to justify their being killed. This could be much more difficult for automated systems to determine, especially since enemy combatants might only pose a threat to the robot, and not to any humans – does that count as a sufficient threat to warrant killing them?

Second, the other major criterion for Just War Theory is proportionality – that the killing and violence committed is proportional to the injustice that it seeks to correct. Just War Theory allows the killing of just enough enemy soldiers in order to win the battle or the war. Proportionality also requires that the use of violence is calibrated to justice. For example, if you punch me in the arm I might be justified in punching you back, but not justified in killing you. Similarly, if one nation were to

repeatedly violate the fishing rules in the territorial waters of another nation, this would not justify a full-scale invasion, or the bombing of the offending nation’s capital city, though it might justify sinking an offending fishing vessel. In this sense, proportionality can be viewed as a retributive component of Just War Theory. Just War Theory also allows for the unintentional killing of innocent civilians, often called “collateral damage,” through the doctrine of double-effect. But the potential risk of killing civilians and the potential strategic value of the intended target, for example when considering whether to bomb a military installation with a school next to it, must both be taken into account in determining whether the risks and costs are justified. I do not believe that an automated system could be built that could make these kinds of determinations in a satisfactory way, because they depend upon moral values and strategic understandings that cannot be formalized. Of course, there are utilitarians and decision theorists who will argue that the values of innocent human lives, and the values of strategic military targets can be objectively established and quantified, but the methods they use essentially treat humans as oracles of value judgements – usually individual preferences or market-established values derived from aggregates of unquestioned individual valuations – rather than actually

provide an algorithm for establishing these values independently of humans. So again, I would not trust any automated algorithm for establishing values in novel situations.

*According to the criteria of Just War Theory, do you think there could be a substantial objection against a military operation because of unmanned systems/military robots being used in it, now or – thinking of the future potential of increasing autonomy of these systems – in a future conflict?*

Since I think that merely meeting the discrimination criterion of Just War Theory is not sufficient for meeting the other criteria, and I doubt that any fully automated system will ever meet the proportionality criteria, I think there are grounds for arguing against the use of systems that make fully automated lethal decisions in general.

Of course, I think we can make a substantial case for international bans on autonomous lethal robots, or other things like space-based weapons, regardless of whether they violate Just War Theory in principle. International treaties and bans depend more upon the involved parties seeing it as being in their mutual interest to impose binding rules on how warfare is conducted. The fundamental weakness of Just War Theory, as Walzer presents it, is that it cannot really be

used to argue definitively against any military technology, insofar as both sides consent to use the technology against each other. The Ottawa Treaty is a notable exception here, insofar as it bans anti-personnel landmines on the basis of their indiscriminate killing of civilians, even long after a war. Mostly that treaty succeeded because of international outrage over the killing and maiming of children by landmines, and the expense of cleaning up mine fields. Basically, politicians could look good and save money by banning a weapon with limited applications that does not really change the balance of military powers.

International treaties tend to be somewhat arbitrary in what they ban, from the perspective of Just War Theory. Blinding enemy combatants is a more proportional way to neutralize the threat they pose than killing them, yet blinding lasers are banned as “disproportionately harmful” weapons. Space-based weapons are not intrinsically unjust, but they represent a potential “tragedy of the commons” in that destroying just a few satellites could put enough debris in orbit to start a chain-reaction of collisions that would destroy most of the orbiting satellites and make it nearly impossible to launch any more into orbit in the future. So it really is in the long-term interest of all nations to ban space-based weapons. There is a

United Nations Committee On the Peaceful Uses of Outer Space (UNCOPUOS) in Vienna that has done some really good work forging international cooperation in space. They have been working for many years to convince the international community to ban space-based weapons, but it is curiously unfortunate that the US, which stands to lose the most strategically from space-based weapons because it has so many satellites in orbit, is the country that is blocking treaties to keep weapons out of space. Perhaps we could form a UN committee on the peaceful uses of robotics?

In your posing of the question, you seem to be asking about whether one could argue against the use of autonomous lethal systems in a particular military operation. The particular case is actually harder to argue than the general case. If military planners and strategists have chosen a specific target, and planned an operation, and plan on using autonomous lethal robots to execute the plan, then we might appear to have a case where these technologies seem acceptable. First of all, there is a significant amount of human decision-making already in the loop in such a case, especially in that there is a valid "target." Second, if it is the kind of mission where we would be deciding between firing a cruise missile to destroy a target, or sending

autonomous lethal robots to destroy the same target, that case is much trickier. Taking the case in isolation, the robots might spare more innocent civilians than a missile, or might collect some valuable intelligence from the target before destroying it. Viewing it in a broader systemic context can change things, however, as there will be new options made possible by the technology. So while there could be cases where an autonomous robot might offer a better option than some technology we already have, there may also be other new technologies that provide even better options. And we can always invent a hypothetical scenario in which a particular technology is the best possible option. But again, I think we need to be careful about how we define and think about autonomy and the level of control of the "humans-in-the-loop." If the humans using this option are willing to take responsibility for the robots completely destroying the target (as would be the case if they used a missile instead), and are in fact held responsible if the target turns out to be a school full of children with no military value, then the fact that they used robots instead of a missile makes little difference. The problem we must avoid is when the humans are not held responsible because they relied on the robot having a safety mechanism that was supposed to prevent it from



killing children. Our frameworks for ethical decision-making do not take into account how technologies change the options we have. The easiest solution to the problem is to make such autonomous systems illegal under international law.

---

<sup>1</sup> Peter M. Asaro, Modeling the Moral User in: *IEEE Technology and Society*, 28, 2009, p.20-24.