

John P. Sullins: Aspects of Telerobotic Systems

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

It was not intentional. My PhD program focused on artificial intelligence, artificial life and consciousness. During my studies I was persuaded by the works of Rodney Brooks and others, who were arguing that embedding AI and robotic systems in real world situations is the only way to gain traction on the big issues troubling AI. So, I began studying autonomous robotics, evolutionary systems, and artificial life. Right away I began to be troubled by a number of ethical issues that harried this research and the military technological applications it was helping to create. Just before I finished my doctorate the events of September eleventh occurred closely followed by a great deal of interest and money being directed at military robotics. Instead of going into defence contract research, as a number of my peers were doing, I decided to go into academic philosophy as this seemed like the best angle from which to speak to the ethics of robotics. Like the rest of us, I have been swept up by historical events and I am doing my best to try to understand this dangerous new epoch we are moving into.

In your work you have engaged questions regarding ethics of artificial life, ethical aspects of autonomous robots and the question of artificial moral agency. Where do you see the main challenges in the foreseeable future in these fields?

In the near term the main issue is that we are creating task accomplishing agents, which are being deployed in very ethically charged situations, be they AI(Artificial Intelligence), ALife (Artificial Life), or robotic in nature.

In ALife work is proceeding on the creation of protocells, which will challenge our commonsense conception of life and may open the door to designer biological weapons that will make the weapons of today look like the horse does now to modern transportation technology.

Autonomous robotics has two main challenges, the most imminent challenge is their use in warfare, which we will talk more about later, but there is also the emergence of social robotics that will grow in importance over the coming decades. Social robots are machines designed as companions, helpers, and as sexual objects. I believe

that a more fully understood concept of artificial moral agency is vital to the proper design and use of these technologies. What worries me most is that in robotics we are rushing headlong into deploying them as surrogate soldiers and sex workers, two activities that are surrounded by constellations of tricky ethical problems that even human agents find immensely difficult to properly navigate. I wish we could have spent some additional time to work out the inevitable bugs with the design of artificial moral agents in more innocuous situations first. Unfortunately, it looks like we will not have that luxury and we are going to have to deal with the serious ethical impacts of robotics without delay.

Concerning the use of robots by the military, Ronald Arkin has worked on an ethical governor system for unmanned systems. Do you think similar developments will be used in other application areas of robots in society? Especially the impact of robots on health care and care for the elderly concerns ethically sensitive areas.

Yes, I do think that some sort of ethical governor or computational application of moral logic will be a necessity in nearly every application of robotics technology. All of one's personal interactions with other humans are shaped by one's own moral sentiments. It comes so natu-

rally to us that it is hard to notice sometimes unless someone transgresses some social norm and draws our attention to it. If we expect robots to succeed in close interactions with people we need to solve the problem Arkin has addressed with his work. Right now, our most successful industrial robots have to be carefully cordoned off from other human workers for safety reasons, so there is no pressing need for an ethical governor in these applications. But when it comes to replacing a human nurse with a robot, suddenly the machine is through into a situation where a rather dense set of moral situations develops continuously around the patients and caregivers. For instance, one might think that passing out medication could be easily automated by just modifying one of the existing mail delivery robots in use in offices around the world. But there is a significant difference in that a small error in mail delivery is just an inconvenience, whereas a mistake in medication could be lethal. Suppose we could make a fool proof delivery system and get around the last objection, even then we have a more subtle problem. Patients in a hospital or nursing home often tire of the prodding, poking, testing and constant regimen of medication. They can easily come to resist or even resent their caregivers. So, a machine dropped into this situation would have to be able to not only get the

right medication to the right patient but then will need to also engage the patient in a conversation to try to convince him or her that it is interested in the well being of the patient and wants only what is best for him or her, listen attentively and caringly to the patients concerns and then hopefully convince the patient to take the medication. We can see that this simple task is imbedded into a very complex and nuanced moral situation that will greatly task any known technology we have to implement general moral intelligence. Therefore I think the medical assistant sector of robotics will not reach its full potential until some sort of general moral reasoning system is developed.

A lot of the challenges concerning the use of robots in society seem to stem from the question of robot autonomy and especially from the question of robots possibly becoming moral agents. Where do you see the main challenges in this field?

This is a great question and I have much to say about it. I have a complete technical argument which can be found in the chapter I wrote on Artificial Moral Agency in Technoethics, in the Handbook of Research on Technoethics Volume one, edited by Rocci Luppincini and Rebecca Addell. But I will try to distil that argument here. The primary challenge is that no traditional

ethical theory has ever given serious concern to even non human moral agents, such as animals, much less artificial moral agents such as robots, ALife, or AI, so we are existing in a conceptual void and thus most traditional ethicists and theologians would find the concept unthinkable or even foolish. I think it is important to challenge this standard moral certainty that humans are the only thing that count as moral agents and instead entertain the notion that it is possible, and in fact desirable, to admit non-humans and even artefacts into the club of entities worthy of moral concern. If you will allow me to quote myself from the work I cited above, "...briefly put, if technoethics makes the claim that ethics is, or can be, a branch of technology, then it is possible to argue that technologies could be created that are autonomous technoethical agents, artificial agents that have moral worth and responsibilities – artificial moral agents."

Let me explain myself a bit more clearly. Every ethical theory presupposes that the agents in the proposed system are persons who have the capacity to reason about morality, cause and effect, and value. But I don't see the necessity in requiring personhood, wouldn't the capacity to reason on morality, cause and effect, and value, be enough for an entity to count as a moral agent? And further, you probably do not

even need that to count as an entity worthy of moral concern, a “moral patient” as these things are often referred to in the technical literature. So, for me a thing just needs to be novel and/or irreplaceable to be a moral patient, that would include lots of things such as animals, eco-systems, business systems, artwork, intellectual property, some software systems, etc. When it comes to moral agency the requirements are a little more restrictive. To be an artificial moral agent the system must display autonomy, intentionality, and responsibility. I know those words have different meaning for different people but by “autonomy” I do not mean possessing of complete capacity for free will but instead I just mean that the system is making decisions for itself. My requirements of intentionality are similar in that I simply mean that the system has to have some intention to shape or alter the situation it is in. And finally the system has to have some moral responsibility delegated to it. When all of these are in place in an artificial system it is indeed an artificial moral agent.

If we speak about a moral judgement made by a machine or artificial life-form, what would be the impact of this on society and human self-conception?

There are many examples of how it might turn out badly to be found

throughout science fiction. But I do not think any of those scenarios are going to fully realize themselves. I believe this could be a very positive experience if we do it correctly. Right now, the research in moral cognition suggests that human moral agents make their decisions based largely on emotion, guided by some general notions acquired from religion or the ethical norms of their culture, and then they construct from these influences their exhibited behaviour. Working on artificial moral agents will force us to build a system that can more rationally justify its actions. If we are successful, then our artificial moral agents might be able to teach us how to be more ethical ourselves. We are taking on a great responsibility, as the intelligent designers of these systems it is ultimately our responsibility to make sure they are fully functioning and capable moral agents. If we can't do that we shouldn't try to build them.

We are not guaranteed success in this endeavour, we might also build systems that are amoral and that actively work to change the way we perceive the world, thus stripping ourselves of the requirements of moral agency. This is what I am working to help us avoid.

You have argued that telerobotic systems change the way we perceive the situation we are in and that this factor and its effect on

warfare is insufficiently addressed. Where do you see the main ethical challenges of this effect and what could be done to solve or at least mitigate these problems?

The main issue is what I call telepistemological distancing: how does looking at the world through a robot colour one's beliefs about the world? A technology like a telero-botic drone is not epistemically passive as a traditional set of binoculars would be. The systems of which the drone and pilot are part of are active, with sensors and systems that look for, and pre-process, information for the human operators' consumption. These systems are tasked with finding enemy agents who are actively trying to deceive it in an environment filled with other friendly and/or neutral agents, this is hard enough for just general reconnaissance operations but when these systems are armed and targets are engaged this obviously becomes a monumental problem that will tax our telepistemological systems to the limit. It does not stop there, once the images enter into the mind of the operator or soldier, a myriad social, political, and ethical prejudgments may colour the image that has been perceived with further epistemic noise.

As we can see, there are two loci of epistemic noise; 1) the technological medium the message is contained in and 2) the preconditioning

of the agent receiving the message. So, if we are to solve or mitigate these problems they have to be approached from both of these directions. First, the technological medium must not obscure information needed to make proper ethical decisions. I am not convinced that the systems in use today do that so I feel we should back off in using armed drones. The preconditioning of the operator is a much harder problem. Today's soldiers are from the X-Box generation and as such come into the situation already quite desensitized to violence and not at all habituated to the high level of professionalism needed to follow the strict dictates of the various ROEs, LOW, or Just War theory. A recent report by the US Surgeon General where US Marines and Soldiers were interviewed after returning home from combat operations in the Middle East suggests that even highly trained soldiers have a very pragmatic attitude towards bending rules of engagement they may have been subject to. As it stands only officers receive any training in just war theory but drones are now regularly flown by non officers and even non military personnel such as the operations flown by the CIA in the US, so I am worried that the pilots themselves are not provided with the cognitive tools they need to make just decisions. To mitigate this we need better training and very close command and control maintained on these

technologies and we should think long and hard before giving covert air strike capabilities to agencies with little or no public accountability.

As far as CIA UAV operations are concerned, one can witness a continuous increase. As you mentioned there are various problems connected with them. To single out just one: do you think the problem with the accountability of the actions – i.e. the question of the locus of responsibility – could be solved in an adequate manner?

This is a very hard problem that puts a lot of stress on just war theory. A minimal criteria for a just action in war, is obviously that it be an action accomplished in the context of a war. If it is, then we can use just war theory and the law of war to try to make some sense of the action and determine if it is a legal and/or moral action. In situations where a teleroobot is used to project lethal force against a target, it is not clear whether the actions are acts of war or not. Typically, the missions that are flown by intelligence agencies like the CIA are flown over territory that is not part of the overall conflict. The “War on Terror” can spill out into shadowy government operators engaging an ill defined set of enemy combatants anywhere on the globe that they happen to be. When this new layer of difficulties is added to the others I have mentioned in this interview,

one is left with a very morally suspect situation. As an example we can look at the successful predator strike against Abu Ali al-Harithi in Yemen back in 2002. This was the first high profile terrorist target engaged successfully by intelligence operatives using this technology. This act was widely applauded in the US but was uncomfortably received elsewhere in the world, even by those other countries that are allied in the war on terror. Since this time the use of armed drones has become the method of choice in finding and eliminating suspected terrorists who seek sanctuary in countries like Pakistan, Yemen, Sudan, Palestine, etc. It is politically expedient because no human intelligence agency agents are at risk and the drone can loiter high and unseen for many hours waiting for the target to emerge. But this can cause wars such as these to turn the entire planet into a potential battlefield while putting civilians at risk who are completely unaware that they are anywhere near a potential fire-fight. While I can easily see the pragmatic reasons for conducting these strikes, there is no way they can be morally justified because you have a non military entity using lethal force that has caused the death and maiming of civilians from countries that are not at war with the aggressor. I am amazed that there has not been sharp criticism of this behaviour in international settings.

Negotiations and treaties will no doubt be needed to create specific rules of engagement and laws of war to cover this growing area of conflict. Yet, even if the major players can agree on rules of engagement and laws for the use of drones that does not necessarily mean the rules and laws obtained will be ethically justified. To do that we have to operate this technology in such a way that we respect the self determination of the countries they are operated in so that we do not spread the conflict to new territories, and we must use them with the double intention of hitting only confirmed military targets and in such a way that no civilians are intentionally or collaterally harmed. I would personally also suggest that these missions be flown by trained military personnel so that there is a clear chain of responsibility for any lethal force used. Without these precautions we will see more and more adventurous use of these weapons systems.

One of the problems you have identified in UAV piloting is, that there is a tendency for these to be controlled not only by trained pilots, typically officers with in-depth military training, but also by younger enlisted men. Also do you see the future possibility to contract UAV piloting to civil operators? What would be the main challenges in these cases and what kind of spe-

cial training would you think would be necessary for these UAV operators?

Yes, there is a wide variety of UAVs in operation today. Many of them do not require much training to use so we are seeing a trend emerging where there are piloted by younger war fighters. Personally, I prefer that we maintain the tradition of officer training for pilots but if that is impossible and we are going to continue to use enlisted persons, then these drone pilots must be adequately trained in the ethical challenges peculiar to these technologies so they can make the right decisions when faced by them in combat situations.

Since the larger and more complex aircraft like the Predator and Raptor, are typically piloted from locations many thousands of miles away, it is quite probable that civil contractors might be employed to fly these missions. That eventuality must be avoided, at least when it comes to the use of lethal force in combat missions. The world does not need a stealthy telerobotic mercenary air force. But, if we can avoid that, I do think there is a place for this technology to be used in a civil setting. For instance, just recently a Raptor drone was diverted from combat operations in Afghanistan and used to help locate survivors of the earthquake in Haiti. Certainly, that is a job that civil pilots

could do. Also, these machines are useful for scientific research, fire patrols, law enforcement, etc. All of which are missions that would be appropriate for civilians to accomplish. The ethical issues here are primarily those of privacy protection, expansion of the surveillance society, and accident prevention. With that in mind, I would hope that civil aviation authorities would work to regulate the potential abuses represented by these new systems.

Regarding the impact of telerobotic weapon systems on warfare, where do you see the main challenges in the field of just war theory and how should the armed forces respond to these challenges?

Just war theory is by no means uncontroversial but I use it since there are no rival theories that can do a better job than just war theory even with its flaws. It is, of course, preferable to resolve political differences through diplomacy and cultural exchange, but I do think that if conflict is inevitable, we must attempt to fight only just wars and propagate those wars in an ethical manner. If we can assume our war is just, then in order for a weapons system to be used ethically in that conflict, it must be rationally and consciously controlled towards just end results.

Telerobotic weapons systems impact our ability to fight just wars in

the following ways. First they seem to be contributing to what I call the normalization of warfare. Telerobots contribute to the acceptance of warfare as a normal part of everyday life. These systems can be controlled from across the globe so pilots living in Las Vegas can work a shift fighting the war in the Middle East and then drive home and spend time with the family. While this may seem like it is preferable, I think it subtly moves combat into a normal everyday activity in direct confrontation with just war theory that demands that warfare be a special circumstance that is propagated only in an effort to quickly return to peaceful relations. Also, telerobots contribute to the myth of surgical warfare and limit our ability to view one's enemies as fellow moral agents. That last bit is often hard for people to understand, but moral agents have to be given special regard even when they are your enemy. Just war attempts to seek a quick and efficient end to hostilities and return to a point where the enemy combatants can again respect one another's moral worth. For instance, look how many of the European belligerents in WWII are now closely allied with each other. The way one conducts hostilities must not be done in a way that prevents future cooperation. Telerobotic weapons seem to be doing just the opposite. The victims of these weapons have claimed that they are cowardly and that far from

being surgical, they create devastating civilian casualties. These allegations may or may not be true, but they are the image that much of the world has of those countries that are using these weapons fanning the flames of intergenerational hatred between cultures.

So what you are saying is, that the current method of using UAVs might actually endanger one of the principles of just war theory, the probability of obtaining a lasting peace (iustus finis), in other words the short term military achievements might curb the long term goals of peace?

Yes, that is exactly right. People who have had this technology used against them are unlikely to forgive or reconcile. When these technologies are used to strike in areas that are not combat zones they tend to fan the flames of future conflict even if they might have succeeded in eliminating a current threat. This can cause a state of perpetual warfare or greatly exacerbate one that is already well underway. For instance, we can see that the use of remote controlled bombs, missiles and drones by both sides of the conflict in Palestine are not ending the fight but are instead building that conflict to new highs of violence.

The armed forces should respond to this by understanding the long-

term political costs that come with short-term political expediency. Right now, a drone strike that causes civilian casualties hardly raises concern in the home audience. But in the rest of the world it is a source of great unease. It is also important to resist the temptation to normalize telerobotic combat operations. I would suggest backing off on using these weapons for delivery of lethal force and move back to reconnaissance missions. And yes, I do know that that will never happen, but at least we should use these weapons only under tight scrutiny, in declared combat zones, with the intent both to justly propagate the conflict and eliminate non combatant casualties.

One question connected to the normalization of warfare through telerobotics, is the so called shift-work fighting. Where do you see the main challenges in the blending of war and civilian life and how could this be countered?

I need to be careful here so that I am not misunderstood. I do understand that these technologies take the war fighters that would have had to risk their own lives in these missions out of danger and put in their place an easily replaceable machine. That is a moral good. But what I want to emphasize is that it is not an unequivocal good. Even if our people are not getting hurt, there will be real human agents on the other end of the cross hairs.

Making a shoot or don't shoot decision is one of the most profound a moral agent can be called on to make. It can not be done in an unthinking or business-as-usual way. When we blend war fighting with daily life we remove these decisions from the special moral territory they inhabit in just war theory and replace it with the much more casual and pragmatic world of daily life. Realistically I do not think there is anyway to counter this trend. It is politically expedient from the viewpoint of the commanders, it is preferable to the individual war fighters, and there does not seem to be any international will to challenge the countries that are using UAVs in this way. As the technology advances we will see more and more naval craft and armoured fighting vehicles operated telerobotically and semi autonomously as well. For instance, this is a major plank of the future warfare planning in America and quite a bit of money is being directed at making it a reality. It is my hope though, that these planners will take some of these critiques seriously and work to keep the operators of these future machines as well trained and professional as possible and that they operate them with no cognitive dissonance. By that I mean the operators should be well aware that they are operating lethal machinery in a war zone and that it is not just another day at the office.

I understand, that in your speech at the IEEE International Conference on Robotics and Automation 2009 in Kobe, you have also presented recommendations for the use of telerobotic weapon systems. What should be our top priority at the moment?

The Conference in Kobe was very interesting. Roboticists such as Ronald Arkin are working hard on designing systems that will act like "ethical governors" in the hope that future autonomous and semi autonomous military robots will be able to behave more ethically than humans do in combat situations. I believe the top priority right now should be to tackle this idea seriously so we can make sure that these ethical governors are more than just an idea but an actual functioning part of new systems. The main sticking point right now is that at least theoretically, a system with a functioning ethical governor would refuse orders that it deemed unethical, and this is proving to be a difficult technology to sell. If I can be permitted one more top priority it would be to investigate some of the claims I have made to provide more detailed information. Is telepistemological distancing real? Do drone pilots view the war as just a kind of super realistic video game? The military has the funds and personnel to carry out these studies and without this data we cannot rationally and consciously use

these weapons and therefore cannot use them ethically.

To mitigate the most detrimental negative effects of telepresence, there are five aspects one might consider:

- 1) Constant attention must be paid to the design of the remote sensing capabilities of the weapon system. Not only should target information be displayed but also information relevant to making ethical decisions must not be filtered out. Human agents must be easily identified as human and not objectified by the mediation of the sensors and their displays to the operator. If this is impossible, then the machine should not be operated as a weapon.
- 2) A moral agent must be in full control of the weapon at all times. This cannot be just limited to an abort button. Every aspect of the shoot or don't shoot decision must pass through a moral agent. Note, I am not ruling out the possibility that that agent may not be human. An artificial moral agent (AMA) would suffice. It is also important to note that AMAs that can intelligently make these decisions are a long ways off. Until then, if it is impossible to keep a human in the decision loop, then these machines must not be used as weapons.
- 3) Since the operator his or herself is a source of epistemic noise, it matters a great deal whether or

not that person has been fully trained in just war theory. Since only officers are currently trained in this, then only officers should be controlling armed telerobots. If this is impossible, then these machines should not be used as weapons.

- 4) These weapons must not be used in any way that normalizes or trivializes war or its consequences. Thus shift-work fighting should be avoided. Placing telerobotic weapons control centres near civilian populations must be avoided in that it is a legitimate military target and anyone near it is in danger from military or terrorist retaliation.
- 5) These weapons must never be used in such a way that will prolong or intensify the hatred induced by the conflict. They are used ethically if and only if they contribute to a quick return to peaceful relations.