



CHALLENGES OF AUTONOMOUS WEAPONS

While autonomous weapons are not new, few of the ethical, legal or operational implications have been clearly identified and solved. In this section, we look at some of the legal aspects and the tactical implications that flow from them. The overriding need to limit collateral damage and avoid killing innocents means that the man-in-the-loop is vital, but is his increasing distance from the battlefield a disadvantage? UCAVs are a case in point and we look at them and the weapons they will carry in the future. We shall return to the subject in future issues.

Grounds for Discrimination: Autonomous Robot Weapons

by Professor Noel Sharkey

Noel Sharkey is Professor of AI and Robotics and Professor of Public Engagement EPSRC Senior Media Fellow at the University of Sheffield. Here he examines various legal and humanitarian aspects of using autonomous weapons and the issues that these raise.

The development most likely to revolutionise war in the 21st Century is the unmanned system. It began with Tesla's efforts to develop remote-controlled torpedoes in the late 19th Century,¹ but it is only in the last decade that remote controlled (also known as tele-operated) mobile machines have become commonplace in conflict zones. The next big step is to take the human out of the loop with autonomous robot weapons, bringing together the latest advances in technology, navigation and artificial intelligence. The US has published 25-year plans for unmanned aircraft,² ground vehicles,³ and naval vehicles,⁴ and a roadmap up until 2032.⁵ It is vitally important that the excitement about the new technology and possibilities of risk-free war do not mask the ethical questions that they raise. We must ensure that the evolution of unmanned weapons conforms to basic humanitarian law.

Protection of Innocents

One of the cornerstones of humanitarian law⁶ – the laws of armed conflict and the law of war, as established in the Geneva and Hague conventions and the various treaties and protocols – is the protection of innocents. This is part of the justice in the conduct of a war, *jus in bello*, and is often expressed as the principle of *discrimination* – only combatants/warriors are legitimate targets of attack. All others, including children, civilians, service workers and retirees, should be immune from attack. It is not just a matter of uniform; soldiers who are wounded, have surrendered or are mentally ill are also immune.⁷

In modern warfare it is difficult to fully protect non-combatants. For example, in attacking a warship, some non-combatants such as chaplains and medical staff may be unavoidably killed. It is also difficult when large explosives are used near civilian populations, or when missiles get misdirected. But the laws of war have a way of handling the unintentional killing of innocents. Thomas Aquinas, in the 13th Century, developed the doctrine of Double Effect. Put crudely, it is OK to kill innocents during a conflict providing that (i) you did not intend to do so, or (ii) that killing the innocents was not a means to winning, or (iii) the importance to the defence of your nation is proportionally greater than a number of civilian deaths.

We may be about to unleash new weapons that could violate all of these principles

The modern equivalent is the Principle of Proportionality which, "... requires that the anticipated loss of life and damage to property incidental to attacks must not be excessive in relation to the concrete and direct military advantage expected to be gained".⁸ But we may be about to unleash new weapons that could violate all of these principles.

Lethal Autonomous Robots

Between four and six thousand robots are currently operating on the ground in Iraq and Afghanistan. These are mainly deployed in dull, dirty or dangerous tasks such as disrupting or exploding improvised explosive devices and surveillance in dangerous areas such as caves. There are only three armed Talon SWORDS robots made by Foster-Miller, although more



The BLU-108 parachutes to near the ground and releases four Skeet warheads [Textron Defense Systems]

are expected soon. Most of the armed robots are in the sky – semi-autonomous Unmanned Combat Air Vehicles such as the MQ1-Predator that flew some 400,000 mission hours up to the end of 2006 and has flown significantly more since, and the more powerful MQ-9 Reapers with a payload of 14 Hellfire missiles. These can navigate and search out targets but, like the ground robots, it is a remote operator, this time thousands of miles away in the Nevada desert, who makes the final decision about when to apply lethal force.

There is now massive spending and plans are well under way to take the human out of the loop so that robots can operate autonomously to locate their own targets and destroy them without human intervention.⁹ This is high on the military agenda of all the US forces: "The Navy and Marine Corps should aggressively exploit the considerable warfighting benefits offered by autonomous vehicles (AVs) by acquiring operational experience with current systems and using lessons learned from that experience to develop future AV technologies, operational requirements, and systems concepts."¹⁰ There is now a number of autonomous ground vehicles such as DARPA's 'Unmanned Ground Combat Vehicle and Perceptor Integration System' otherwise known as the Crusher.¹¹ And BAE systems recently reported that they have "... completed a flying trial which, for the first time, demonstrated the coordinated control of multiple UAVs autonomously completing a series of tasks".¹²

The move to autonomy is clearly required to fulfil the current US military plans. Tele-operated systems are more expensive to manufacture and require many support personnel to run them. One of the main goals of the Future Combat Systems project is to use robots as a force multiplier so that one soldier on the battlefield can be a nexus for initiating a large-scale robot attack from the ground and the air. Clearly one soldier cannot operate several robots alone and it takes the soldier away from operational duties.

Discrimination

The ethical problem is that no autonomous robots or artificial intelligence systems have the necessary skills to discriminate between combatants and innocents. Allowing them to make decisions about who to kill would fall foul of the fundamental ethical precepts of a just war under *jus in bello* as enshrined in the Geneva and Hague conventions and the various protocols set up to protect civilians, wounded soldiers, the sick, the mentally ill and captives. There are no visual or sensing systems up to that challenge.

A computer can compute any given procedure that can be written down in a programming language. We could, for example, give the robot computer an instruction such as, "If civilian, do not shoot". This would be fine if, and only if, there was some way of giving the computer a clear definition of what a civilian is. We certainly cannot get one from the Laws



The Skeet warhead released by BLU-108 [Textron Defense Systems]

of War that could provide a machine with the necessary information. The 1944 Geneva Convention requires the use of common sense, while the 1977 Protocol 1 essentially defines a civilian in the negative sense as someone who is not a combatant:

- A civilian is any person who does not belong to one of the categories of persons referred to in Article 4 A (1), (2), (3) and (6) of the Third Convention and in Article 43 of this Protocol. In case of doubt whether a person is a civilian, that person shall be considered to be a civilian.
- The civilian population comprises all persons who are civilians.
- The presence within the civilian population of individuals who do not come within the definition of civilians does not deprive the population of its civilian character.¹³

And even if there was a clear computational definition of a civilian, we would still need all of the relevant information to be made available from the sensing apparatus. All that is available to robots are sensors such as cameras, infrared sensors, sonars, lasers, temperature sensors and ladars etc. These may be able to tell us that something is a human, but they could not tell us much else. In the labs there are systems that can tell someone's facial expression or that can recognise faces, but they do not work on real-time moving people.

In a conventional war where all of the combatants wore the same clearly marked uniforms (or better yet, radio frequency tags) the problems might not be much different from those faced for conventional methods of bombardment. But the whole point of using robot weapons is to help in warfare against insurgents, and in these cases sensors would not help in discrimination. This would have to be based on situational

awareness and on having a theory of mind, i.e. understanding someone else's intentions and predicting their likely behaviour in a particular situation. Humans understand one another in a way that machines cannot and we don't fully understand how. Cues can be very subtle and there is an infinite number of circumstances where lethal force is inappropriate. Just think of children being forced to carry empty rifles or insurgents burying their dead.

There is also the Principle of Proportionality and again there is no sensing or computational capability that would allow a robot such a determination, and nor is there any known metric to objectively measure needless, superfluous or disproportionate suffering.¹⁴ They require human judgement. Yes, humans do make errors and can behave unethically, but they can be held accountable. Who is to be held responsible for the lethal mishaps of a robot? Certainly not the machine itself. There is a long causal chain associated with robots: the manufacturer, the programmer, the designer, the department of defence, the generals or admirals in charge of the operation and the operator.

International Guidelines

There are no current international guidelines for, or even discussions about, the uses of autonomous robots in warfare. These are needed urgently. If there was a political will to use them then there would be no legal basis on which to complain.¹⁵ This is especially the case if they could be released somewhere where there is a fairly high probability that they will kill a considerably greater number of enemy combatants (uniformed and non-uniformed) than innocents – i.e. the civilian death toll was not disproportionate to the military advantage.

The ethical problem is that no autonomous robots or artificial intelligence systems have the necessary skills to discriminate between combatants and innocents

In this way autonomous robots would be legally similar to submunitions such as the BLU-108 developed by Textron Defense Systems.¹⁶ The BLU-108 parachutes to near the ground where an altitude sensor triggers a rocket that spins it upwards. It then releases four Skeet warheads at right angles to one another. Each has a dual-mode active and passive sensor system: the passive infrared sensor detects hot targets such as vehicles, while the active laser sensor provides target profiling. They can hit hard targets with penetrators or destroy soft targets by fragmentation.

But the BLU-108 is not like other bombs because it has a method of target discrimination. If it had been developed in the 1940s or 1950s there is no doubt that it would have been classified as a robot and even now it is debatably a form of robot. The Skeet warheads have autonomous operation and use sensors to target their weapons. The sensors provide discrimination between hot and cold bodies of a certain height but, like autonomous robots, they cannot discriminate between legitimate targets and civilians. If BLU-108s were dropped on a civilian area they would destroy buses, cars and lorries. Like conventional bombs, discrimination between innocents and combatants requires accurate human targeting judgements. It is this and only this that keeps the BLU-108 within humanitarian law.¹⁷

But the whole point of using robot weapons is to help in warfare against insurgents, and in these cases sensors would not help in discrimination

Future Use

To use robot technology over the next 25 years in warfare would at best be like using the BLU-108 submunition – i.e. can sense a target but cannot discriminate innocent from combatant. But the big difference with the types of autonomous robots currently being planned and developed for aerial and ground warfare is that they are not perimeter-limited like the Skeet. The BLU-108 has a footprint of 820ft all around. By way of contrast, mobile autonomous robots are limited only by the amount of fuel or battery power that they can carry. They can potentially travel long distances and move out of line-of-sight communication.

Imagine the potential devastation of heavily armed robots in a deep mission out of radio communication. The only humane recourse of action is to severely restrict or ban the deployment of these new weapons until there have been international discussions about how they might pass an ‘innocents discrimination test’. At the very least there should be discussion about how to limit the range and action of autonomous robot weapons before the inevitable proliferation. ■

Systems/ Land Warfare and Munitions, 3090 Pentagon, Washington DC 20301-3090

- ⁴ *The Navy Unmanned Undersea Vehicle (UUV) Master Plan*, Department of the Navy, USA, 9 November 2004
- ⁵ *Unmanned Systems Roadmap 2007–2032*, US Department of Defense, 10 December 2007
- ⁶ For a more detailed discussion of humanitarian law see Schmitt, M.N. (The Principle of Discrimination in 21st Century Warfare, *Yale Human Rights and Development Law Journal* 143, 1999
- ⁷ But see also Ford, John S., The Morality of Obliteration Bombing, *Theological Studies*, pages 261 – 309, 1944
- ⁸ Petraeus D.H. and Amos, J.F., *Counterinsurgency*, Headquarters of the Army, Field Manual FM 3-24 MCWP 3-33.5, Section 7-30
- ⁹ Sharkey, N., Cassandra of False Prophet of Doom: AI robot and war, *IEEE Intelligent Systems*, Volume 23 No 4, 14 –17, 2008, July – August Issue
- ¹⁰ Committee on Autonomous Vehicles in Support of Naval Operations National Research Council (2005) *Autonomous Vehicles in Support of Naval Operations*, Washington DC, The National Academies Press
- ¹¹ Pentagon’s ‘Crusher’ Robot Vehicle Nearly Ready to Go, *Fox News*, 27 February 2008
- ¹² United Press International, BAE Systems Tech Boosts Robot UAVs IQ, *Industry Briefing*, 26 February 2008
- ¹³ Protocol 1 Additional to the Geneva Conventions, 1977 (Article 50)
- ¹⁴ Bugsplat software and its successors have been used to help calculate the correct bomb to use to destroy a target and calculate the impact. A human is there to decide and it is unclear how successful this approach has been in limiting civilian casualties
- ¹⁵ But it seems that, regardless of treaties and agreements, any weapon that has been developed may be used if the survival of a state is in question. The International Court of Justice (ICJ) (1996) *Nuclear Weapons Advisory Opinion* decided that it could not definitively conclude that in every circumstance the threat or use of nuclear weapons was axiomatically contrary to international law, see Stephens, D. and Lewis, M.W.) *The Law of Armed Conflict – a Contemporary Critique*, *Melbourne Journal of International Law* 6, 2005
- ¹⁶ Thanks to Richard Moyes of Landmine Action for pointing me to the BLU-108 and to Marian Westerberg and Robert Buckley from Textron Defense Systems for their careful reading and comments on my description
- ¹⁷ A key feature of the BLU-108 is that it has built-in redundant self-destruct logic modes that largely leave battlefields clean of unexploded warheads and thus keeps it out of the 2008 treaty banning cluster munitions

NOTES

- ¹ Sharkey, N. and Sharkey, A. (in press) The Electro-mechanical Robot Before the Computer, *Journal of Mechanical Engineering Science*
- ² *Unmanned Aircraft Systems Roadmap 2005–2030*, Office of the US Secretary of Defense, 2005
- ³ *Joint Robotics Program Master Plan FY2005*, LSD (AT&L) Defense