Chapter 11
Military Uses of AI

In this chapter we discuss the controversial topic of autonomous weapons systems. We define terms, present elements and perspectives related to the debate surrounding autonomous weapons systems (AWS) and present the main arguments for and against the use of such systems.

War is inherently controversial. Many military uses of AI and robotics are likewise contentious. Perhaps the most controversial aspect of this topic is the development and use of lethal autonomous weapons systems capable of autonomously making life and death decisions regarding human targets. Some argue that cruise missiles are a form of lethal autonomous weapons system. Systems such as the Patriot missile system, AEGIS naval weapons system, Phalanx weapons system, and Israeli Harpy weapons system are examples of lethal autonomous weapons systems currently in use. The Patriot system, AEGIS, and the Phalanx system are generally considered defensive weapons (Fig. 11.1).

The Harpy is an offensive fire-and-forget weapon that targets enemy air warfare radar systems. It is worth noting that the term military robot includes many non-lethal applications. For example, autonomous robots may be used for mine clearing, explosive ordnance disposal, command and control, reconnaissance, intelligence, mobile network nodes, rescue missions, supply and resupply missions, and support operations. Debates against military robots may vary with respect to the role of the robot. The aim of this chapter is to present an objective account of the most commonly presented arguments for and against the use of AWS in war.

Fig. 11.1   MIM-104 Patriot (*Source* Darkone)

11.1   Definitions

We begin by defining some common terms.

Autonomous   In AI and robotics autonomy simply means the ability to function without a human operator for a protracted period of time (Bekey 2005). Robots may have autonomy over the immediate decision that they make but generally do not have autonomy over their choice of goals. There is some controversy as to what "autonomous" means for lethal weapons systems.

Lethal and Harmful autonomy   A weapon can be said to be "autonomous" in the "critical functions of targeting" if it can do one or more of following without a human operator. If the weapon can decide what classes of object it will engage, then it would be autonomous in terms of defining its targets. No current AWS has this capability. If a weapon can use sensors to select a target without a human operator, it can be said to have autonomy in the selection function of targeting. Many existing weapons can select targets without a human operator. If a weapon can fire on a target without a human operator, it can be said to have autonomy in the engage function of targeting. Many existing weapons can engage already selected targets autonomously. For example, the Patriot anti-missile system can select targets autonomously but by design requires a human operator to hit a confirm button to launch a missile. Once the missile is launched, it can hit its

target without a human operator. Given the speeds involved, human control of a Patriot missile is not be possible.

Non-Lethal Autonomy   An AWS may have "autonomy" in many other functions. It might be able to take-off and land autonomously and it might be able to navigate autonomously. However, this non-lethal "autonomy" is generally not regarded as morally controversial.

Killer robots   Autonomous weapons are often called "killer robots" in mass media reports. Some object to the use of the term. Lokhorst and van den Hoven describe the phrase as an "insidious rhetorical trick" (Lokhorst and Van Den Hoven 2012). However, this is favoured by the "Campaign to Stop Killer Robots".[1] This is umbrella group of human rights organisations seeking an international ban on lethal autonomous weapons systems.

11.2 The Use of Autonomous Weapons Systems

Arguments can and are made against the use of an autonomous weapons system. Generally these arguments focus on the following issues.

11.2.1 Discrimination

Proponents typically concede that machines cannot, in general, discriminate as well as humans. However in some particular cases they can discriminate better than humans. For example, Identification Friend or Foe (IFF) technology sends a challenge message to an unindentified object in the sky which the object must answer or risk being shot down. Typically in air war, contested airspace is known to civilian air traffic control and neutral aircraft will not enter it. However, in 2014, a civilian airliner, Malaysia Airlines flight MH 17 en route from Amsterdam to Kuala Lumpur was shot down by a Russian Surface to Air Missile (SAM) operated by Russian secessionists in the Eastern Ukraine. This SAM system was human-operated and not equipped with IFF. Some have observed that a more advanced system would have known the target was a civilian airliner.

Proponents also note that vision systems are continuously improving. Advancing technology has dramatically improved the ability of vision, auditory, LIDAR and infra-red systems which are quickly reaching parity with humans in terms of object discrimination. A possible ethical dilemma may be approaching if an autonomous system demonstrates clear superiority to humans in terms of targeting. We may be ethically obligated to consider their use. Even so, it remains difficult for machines to distinguish between different types of behavior such as acting peaceful or fighting in a conflict.

---

[1] http://stopkillerrobots.org.

Opponents note that it has been frequently claimed that AWS cannot discriminate between combatants and non-combatants. Noel Sharkey, a leading campaigner against AWS, has wondered if a robot could discriminate between a child holding an ice-cream cone and a young adult holding a gun (Sharkey 2010).

11.2.2 Proportionality

Proponents, on the other hand, state that "excessive" is a relative concept which is not well-defined in International Humanitarian Law (IHL). Enemark makes the point that politicians generally do not advertise their proportionality calculations (Enemark 2013). Situations in which intelligence reveals the location of a high value target demand a decision.

Opponents claim that AWS cannot calculate proportionality (Braun and Brunstetter 2013). Proportionality is the ability to decide how much collateral damage is acceptable when attacking a military target. The standard is that "collateral damage" must not be "excessive" compared to the concrete military advantage gained. Proportionality calculations typically attempt to estimate the number of civilians that may be killed versus the military necessity of the target. Generating such calculations often involves input from a variety of experts including lawyers. It is difficult to imagine how an AWS could successfully complete such a calculation.

11.2.3 Responsibility

Opponents of AWS argue that machines cannot be held morally responsible. They then argue that this is a reason to ban AWS. It is indeed hard to imagine how a machine can be assigned moral responsibility. However those defending the use of AWS are inclined to assign moral responsibility for the actions of the machine to those that design, build and configure it. Thus those humans deploying an AWS can be held responsible for its actions (Arkin 2008). This raises the "problem of many hands" (Thompson 1980) in which the involvement of many agents in a bad outcome makes it unclear where responsibility lies. Clearly if an incident were to occur an investigation would result to determine fault. Legally it is easier to hold the collective entity responsible. There is a concept of "strict liability" in law that could be used to assign responsibility to the state that operates the weapon in an AWS regulation. We discuss liability with the more specific example of an mistargeting by an autonomous weapon in Sect. 5.2.

Opponents also argue that, unlike a human, an AWS cannot be held responsible for its actions or decisions. While machines can be grounded for performance errors there is no true way to punish these systems in a metaphysical sense. Moreover, it may not be just to punish the commanders of these systems if they utilize automatic targeting.

11.3 Regulations Governing an AWS

States at the UN agree that an AWS must be used in compliance with existing IHL. The Convention on Certain Conventional Weapons is generally considered the appropriate forum to discuss AWS regulations. These regulating bodies require meaningful human control over the AWS. In particular, this means the following:

1. an AWS must be able to distinguish between combatants and non-combatants;
2. an AWS must be able to calculate proportionality;
3. an AWS must comply with the principle of command responsibility.

11.4 Ethical Arguments for and Against AI for Military Purposes

11.4.1 Arguments in Favour

In IHL the doctrine of military necessity permits belligerents to do harm during the conduct of a war. Moreover, just war theory states that, although war is terrible, there are situations in which not conducting a war may be an ethically and morally worse option (International Committee of the Red Cross 2015). For example, war may be justifiable to prevent atrocities. The purpose of just war theory is to create criteria that ensures that war is morally justifiable. Just war theory includes criteria for (1) going to war (jus ad bellum) and (2) conducting war (jus in bello). The criteria for going to war include: just cause, comparative justice, competent authority, right intention, probability of success, last resort, and proportionality. The criteria for conducting war include: distinction, proportionality, military necessity, fair treatment of prisoners of war, and not using means and methods of warfare that are prohibited. Examples of prohibited means of warfare include chemical and biological weapons. Examples of prohibited means include mass rape and forcing prisoners of war to fight against their own side. The overall intent of IHL is to protect the rights of the victims of war. This entails rules that minimise civilian harm.

With respect to the use of AI and robots in warfare, some have argued that AMS may reduce civilian casualties (Arkin 2010). Unlike humans, artificially intelligent robots lack emotions and thus acts of vengeance and emotion-driven atrocities are less likely to occur at the hands of a robot. In fact, it may be the case that robots can be constructed to obey the rules of engagement, disobeying commands to violate civilian and enemy combatant rights (Arkin 2008). If nothing else, units being observed by a military robot may be less inclined to commit such atrocities. If, in fact, robots can be used to prevent atrocities and ensure the minimisation of civilian casualties, then military leaders may have an ethical obligation to use such robots. For, to not use such robots, condemns a greater number of civilians to die in a morally justified

war. Moreover, an AWS may be capable of non-lethal offensive action where human units must use lethal force.

Other argue that AI and military robots are necessary for defensive purposes. Some research has shown that in certain circumstances, such as aerial combat, autonomous systems have clear advantages over human systems (Ernest et al. 2016). Hence, sending humans to fight an AWS is unlikely to succeed and may result in substantial casualties. In this situation, leaders have an ethical obligation to reduce their own casualties even if this means developing AWS for their own purposes.

11.4.2 Arguments Against

It has been claimed that the advent of artificial intelligence technologies for military use could lead to an arms race between nations. Vladimir Putin, the President of the Russian Federation, said in 2017 that "the nation that becomes the leader in AI will rule the world." (James 2017). China has similarly increased spending on AI (Herman 2018) and the United States has long made the development of AI for defense purposes a priority (Department of Defense 2012). Experts generally agree that AWS will generate a clear and important military advantage (Adams 2001). Relatedly, some argue that use of an AWS is unfair in that such weapons do not result in equal risk to all combatants.

Researchers also note that the possession and use of autonomous weapons systems may actually instigate wars because the human cost of war is reduced. There is some evidence for this claim based on targeting killings in Iraq by the United States. The transition in Iraq from human piloted missions to unmanned aerial vehicles resulted in a dramatic increase in the number of targeting missions (Singer 2009). This evidence, although important, should not discount the political and technological factors that may also have contributed to the increase in targeted killings.

Perhaps the most philosophically interesting argument levelled against the use of AWS is the dignity argument claiming that "death by algorithm" is the ultimate indignity. In its more complex forms, the argument holds that there is a fundamental human right not to be killed by a machine. From this perspective, human dignity, which is even more fundamental than the right to life, demands that a decision to take human life requires consideration of the circumstances by a human being (Arkin et al. 2012; Heyns 2016). A related claim is that meaningful human control of an autonomous weapon requires that a human must approve the target and be engaged at the moment of combat.

11.5 Conclusion

To conclude we have sought to present, relevant definitions associated with autonomous weapons systems, the ideas behind the regulations that govern these sys-

tems, and the arguments for and against their use. An AWS cannot be lawfully used for genocide and massacre of civilians because existing humanitarian law already prohibits such acts. It is important to note that AWS are already regulated and must be used in accordance with existing international law. It is important to keep in mind that these systems are changing. As they do, they may raise new and important ethical issues that should be discussed within and between nations. Further, in much the same way that commanders are responsible for the actions of their soldiers, commanders are also responsible for the actions of their AWS.

Discussion sQuestions:

- Is the use of AWS in military conflicts justified? Explain.
- What limits or conditions would you set before an AWS could be used? List a set of conditions.
- Should there be new IHL for AWS? Discuss.

Further Reading:

- Ronald Arkin. *Governing lethal behavior in autonomous robots*. Chapman and Hall/CRC, 2009. ISBN 978-1420085945. URL http://www.worldcat.org/oclc/933597288
- Peter Warren Singer. *Wired for war: The robotics revolution and conflict in the twenty-first century*. Penguin, 2009. ISBN 1594201986. URL http://www.worldcat.org/oclc/958145424
- Paul Scharre. *Army of none: Autonomous weapons and the future of war*. WW Norton & Company, 2018.