INFORMATION PAPER SERIES

Presented by the International Law Department, U.S. Naval War College

SERIES NUMBER 13-1 (2013)

Subject: Legal Implications of Autonomous Weapon Systems

Author: LTC Jeffrey S. Thurnher, U.S. Army, Faculty, U.S. Naval War College

Date: 15 May 2013

1. **Purpose**. To provide an overview of the legal issues and controversy surrounding the potential development and deployment of autonomous weapon systems.

2. **Background**. Autonomous weapon systems (AWS) are weapons that, once activated, can select and engage targets without further intervention by a human operator. Though fully autonomous weapons are not known to currently exist in any nation's arsenal, recent technological advances, particularly in artificial intelligence, demonstrate that they may be possible in the not too distant future. In fact, several experts predict AWS will become the norm on the battlefield within twenty years. Significant opposition to AWS has already formed, with some prominent groups calling for a complete ban on future research and development of these systems. Critics rely on legal, ethical, and moral arguments against AWS. This information paper will focus solely on the legal arguments and aspects of AWS. Undeniably, before AWS can be deployed on the battlefield, the systems must be capable of fully complying with the law of armed conflict (LOAC). Although some aspects of LOAC may prove particularly challenging for autonomous systems, contrary to the critics' assertions, the use of AWS would be lawful in certain circumstances.

3. **Current State of Technology**. Although the U.S. does not yet possess fully autonomous weapons, DoD began including automatic and autonomous features into weaponry years ago. For example, the Navy has long used close-in weapon systems designed to defeat incoming missile attacks by reacting automatically to the signatures of such threats. Many air defense weapon systems, like the Army's Patriot Missile system, have also been able to function for many years with various degrees of autonomy. In recent years, artificial intelligence researchers have demonstrated remarkable advances. Innovative, new computing approaches are more closely simulating human thought patterns and allowing wider separation between human operators and robotic systems. The U.S. and other nations have seized on these modernizations, and, in coming years, autonomous systems will likely be able to perform complex tasks without human operator involvement, such as taking off and landing on aircraft carriers, refueling in mid-air, and tracking enemy submarines for

months at a time. The forecast for continued improvements in autonomous capabilities seems optimistic, and future systems can be expected to possess almost unimaginable increases in capability.

4. **DoD Policy on AWS**. The DoD has thus far taken a cautious approach to the development of AWS. In November 2012, it released DoD Directive 3000.09, entitled "Autonomy in Weapon Systems," which established policies and guidelines for the development of autonomous functions in weapons. In particular, the directive mandates that any autonomous system must be designed so that the decision to use lethal force retains appropriate levels of human judgment. Overall, the directive does not advocate for the development of fully autonomous weapons. In fact, the talking points memorandum associated with the directive's release indicates that the U.S. currently has no plans to develop lethal AWS, other than human-supervised systems used defensively. Nonetheless, the directive outlines a required approval chain and details the necessary legal reviews to be conducted before AWS may be designed. It further mandates the enactment of a series of precautionary measures, such as safety measures and anti-tampering mechanisms, designed to keep AWS from striking unintended targets. The directive's provisions, however, do not apply to cyber systems or any type of mines.

5. **Critics of AWS**. Despite the issuance of the DoD Directive, strong opposition to AWS exists. Fearful that AWS might not be able to comport with international law, several nongovernmental organizations have called for a preemptive international treaty banning their development and use. Chief among these groups is Human Rights Watch, an influential human rights organization, who issued a critique of AWS in a November 2012 report entitled "Losing Humanity: The Case against Killer Robots." In April 2013, a coalition of NGOs, including Human Rights Watch, formed the Campaign to Stop Killer Robots. The group, led by a well-known Nobel Laureate, has begun lobbying governments and citizens in the UK and around Europe to support a ban on AWS development. Also in April 2013, a United Nations Special Rapporteur issued a report to the UN Human Rights Council recommending a suspension of all AWS research and development until nations can agree on a legal and regulatory framework for their use. Such prominent criticism has garnered significant worldwide media attention.

6. **Contentious Legal Aspects of AWS**. It is incontrovertible that the law of armed conflict applies to AWS and their use. Of course, the fact that autonomous weapon systems will locate and attack persons and objects without human interaction raises unique issues. Several aspects of LOAC are the focus of AWS critics and need to be examined closely:

a. **Distinction**. Distinction, one of the most fundamental principles of LOAC, obliges a combatant to distinguish between combatants and civilians, as well as between military and civilian objects. The principle is intended to protect civilians by directing military attacks against only military objectives, and it unequivocally applies to AWS. When analyzing whether an autonomous weapon's use complies with the principle of distinction, the surrounding context and environment are critically important. Circumstances may exist in which AWS would only need a low level ability to distinguish in order to comply with the rule. Examples might include conflicts that occur in remote areas, such as underwater, deserts, or places like the DMZ in Korea. In less clear-cut environments, the demands on AWS to distinguish are much higher. For instance, on cluttered battlefields or in urban areas, AWS may need to be equipped with robust sensor packages and advanced recognition software. Even with such cutting-edge

capabilities, there could be complex situations where AWS are simply unable to fulfill this requirement and, therefore, could not lawfully be used. Ultimately, for the use of AWS to be lawful, the systems would be expected to reasonably distinguish between combatants and civilians given the particular environment and circumstances of the battlefield ruling at the time.

b. **Proportionality**. Proportionality requires combatants to examine whether the expected collateral damage from an attack would be excessive in relation to the anticipated military gain. This complex analysis has traditionally involved a human judgment call evaluated on the basis of reasonableness. Many critics of autonomous weapons have questioned whether the systems can fulfill this requirement. In particular, determining the military advantage of a target may present a significant challenge for AWS. The military advantage of a particular target is extremely contextual, and its value can change rapidly based upon developments on the battlefield. Human operators may be able to develop sliding scale-type mechanisms which regularly update and provide the AWS with the relative military advantage value of a given target. Operators might also help fulfill this principle by detailing strict rules of engagement and establishing other controls, such as geographic or time limits on use. Nevertheless, these complicated issues would need to be resolved, if the future use of AWS is to comport with the principle of proportionality.

c. **Subjectivity**. Subjectivity plays a significant role in various facets of LOAC. For example when dealing with proportionality, a combatant must subjectively decide whether the expected collateral harm is excessive in relation to the anticipated military advantage. It is doubtful that AWS will be able to make these subjective determinations themselves in the foreseeable future, even with the most optimistic projections for artificial intelligence advancements. Many opponents of AWS contend that the systems are unlawful because they are unable to make subjective targeting judgments. This view is somewhat misguided, however, by failing to fully appreciate how the AWS targeting process will actually occur. To comply with the law, humans will need to inject themselves at various points into the process and make the necessary subjective determinations. Sometimes these judgment calls will be made before the AWS are even launched. In providing attack criteria or thresholds, the human operator is framing the environment for which AWS will operate. The human is essentially providing the subjective answers in advance. Then, with that guidance embedded into the software, AWS will be tasked with making objective calculations about how to perform on the battlefield. This analysis represents a new way of looking at the LOAC subjectivity requirements and will invariably be controversial.

d. **Accountability**. In many ways, AWS represent a greater separation of humans from the battlefield. Therefore, significant questions arise when one looks to assess accountability for battlefield conduct. Opponents of AWS argue that the removal of humans from the final targeting decisions prevents the proper assignment of legal responsibility. This position fails to take into account the full involvement of human commanders in the overarching targeting process. Contrary to the critics concerns, humans can be held legally responsible for the results of AWS attacks, even when they are not controlling the specific actions of the system. Some accountability issues are relatively straightforward. Naturally, an individual who intentionally programs AWS to engage in actions that amount to a war crime would certainly be liable. Likewise, an individual would be

responsible for using a system in an unlawful manner, such as deploying AWS that are incapable of distinguishing into areas where civilians are expected to be located. Other AWS accountability issues, however, are more complex. For example, critics argue that some autonomous weapons and their embedded software will be too complicated for a human commander to reasonably understand. In such cases, it may be difficult to hold the commander responsible for the system's actions and an accountability gap might appear. Given these challenges, accountability for AWS may represent a significant hurdle.

7. **Conclusion**. Autonomous technology is poised to revolutionize warfare. AWS offer numerous advantages in terms of protecting one's own force and may even ultimately provide better protections for civilians. Before they can be deployed, however, the unique legal concerns with AWS must be fully addressed.

**The views expressed are those solely of the author and should not be understood as necessarily representing those of the U.S. Department of Defense or any other government entity.