

Viewpoints

Lethal Autonomous Weapons Systems – Warfare’s Best Humanitarian Hope?

By Wing Commander Jim Beldon

Introduction

The past decade has witnessed a revolution in the use of remotely operated systems by the UK’s Armed Forces. Nowhere has this been more evident – or controversial – than in the air domain. Debate over the nomenclature of such systems – known variously as ‘Unmanned Air Vehicles (UAVs); ‘Uninhabited Air Systems (UASs); ‘Remotely Piloted Air Systems (RPAS)’ and the plethora of hybrids that these and other terms have spawned – reflects the ideological battle that continues to rage over the nature of such systems and the extent to which meaningful human control prevails over them. The term ‘drone’ has become the popular, yet currently misleading, term for such systems, which has been exploited by opponents to propagate the false notion that the RAF (through its use of the armed MQ-9 Reaper RPAS) is engaged in unethical and inhumane killing by autonomous machines beyond human control. In this characterization, the anti-drone lobby has been wholly wrong, as the Minister for the Armed Forces, Penny Mordaunt MP, recently addressed during a House of Commons adjournment debate on ‘Drones in Conflict’:

I will briefly provide a bit of clarity and on the record bust some of the myths that surround the term “drone”, which conjures up images of machines free from human oversight and able to operate with complete autonomy. That is the stuff of science fiction movies, not the reality. Although drones do not operate with an individual in the cockpit, the fact is that a trained professional human being is in control of the system at all times. The difference is that they operate remotely from the vehicle. The term “drone” also overlooks the fact that the aircraft itself is part of a much larger system composed of other vital components such as the ground stations, networks and, most importantly,

the personnel. ...The Government have no intention of developing systems that operate without this all-important human hand in the weapon command and control chain.¹

But, as this paper explores, there is reason to question whether the UK’s position on Lethal Autonomous Weapons Systems (LAWS) – which applies across all environments – is ultimately sustainable or even desirable, should such systems become viable.² The notion that a human should always be the ultimate decision maker in the delivery of lethal force is certainly correct now, but it is debatable whether this can or should remain so in the future.

At the MOD’s first RPAS-focused media event held at RAF Waddington in December 2013, the then Defence Secretary, Philip Hammond MP, observed that ‘Much of the criticism of unmanned aerial systems is based on misunderstanding. This event provides a great opportunity to better inform people about these life-saving assets and their variety of purposes.’³ In that aim, the event was successful, and so, albeit to a limited degree, has been the MOD’s subsequent communications effort in countering the ‘Killer Drones’ narrative concerning the use of Reaper, a cause somewhat hindered by the system’s unedifying name – an issue which the Prime Minister recently sought to address by giving the RAF’s next generation of armed RPAS the name ‘Protector’.⁴ Encouragingly, the idea of remotely piloted drones as ‘life savers’ or ‘protectors’ has gained some traction in academic circles, with Dr David Whetham of King’s College London⁵ and Professor Caroline Kennedy of the University of Hull⁶ arguing in favour of their use in UN Peacekeeping and Enforcement operations. In addressing the bad press surrounding drones (largely as a result of US ‘targeted killings’), Professor Kennedy argues that:

Drones, even armed drones, can be used in a virtuous manner to protect civilians in line with a UN mandate, just as they can be used in a manner which is perceived as immoral or unethical. ...armed drones are not innately evil or immoral weapons and if used in a manner which deters and prevents acts of genocide and human rights violations then they would likely be welcomed by a public under siege and in need of protection.⁷

The ‘virtuous’ life-saving drone, in the sense intended by Professor Kennedy, is one that is under permanent human control, reliant on the virtues of its human controllers. But is it conceivable that drones could themselves exhibit sufficient humanitarian virtue to make life taking decisions without explicit human involvement? So far, in promoting its case that the rules⁸ governing the employment of lethal force by RPAS are identical to those involving traditional aircraft, the MOD has ruled out developing autonomous weapons systems, stating that only expert military personnel make decisions involving the employment of lethal force and that it is neither possible nor desirable for such decisions ever to be delegated to non-human entities. However, the technology enabling ever higher degrees of automation is evolving rapidly, and so are the arguments. Edges are blurring between machine-made and man-made decisions. Accordingly, although the progression towards autonomy in weapons systems presents some useful opportunities, the legal, ethical and presentational challenges that accompany such

advances are already causing the UK's and other states' policymakers headaches, not least because of public unease over the use of 'drones'.

Perhaps unsurprisingly, owing to antipathy towards drone strikes, vastly more has been written arguing against their use and the development of autonomous weapons systems than has been written in favour.⁹ A notable exception is the US computer science and roboticist, Professor Ronald C. Arkin, whose research in the field of autonomy in military systems has added an important perspective to the potential military and ethical benefits that may result from the development of appropriately intelligent LAWS. His 2010 paper on 'Ethical Autonomy in Unmanned Systems'¹⁰ drew together research from a wide variety of sources and over many decades to demonstrate the failure of fighting men and women to behave ethically in war; indeed, nearly two-thirds of his paper was devoted to ethical, moral and psychological human failings in combat – a disquieting yet revealing insight for the military professional. To a large extent, therefore, Arkin's argument in favour of the development of ethical LAWS relies principally on the demonstrated incapability of humans to perform ethically in war rather than on the as yet unproven virtues of ethically endowed machines. In the intervening five years since Arkin's paper was first published, the ethical and legal arguments over LAWS have gained momentum, courting the attention of some eminent figures, such as Professor Stephen Hawking, and Noam Chomsky, who have called for a pre-emptive ban on LAWS.¹¹ Whilst this author disagrees with the position taken by Hawking et al, their intervention nonetheless highlights that not only has debate on LAWS gained momentum since Arkin's 2010 paper, but so has the technology.

Although a precise timetable cannot yet be given, there is reason to consider that evolution – or perhaps revolution – in the artificial intelligence and robotics fields will ultimately fulfil the dream (or nightmare) of drones endowed with the ability to form reasoned judgements and then decide and act on them without human input.¹² When technology spawns such capabilities, the only remaining impediments to their weaponisation would be international law and decision makers' ethics. Whilst no state (or any other body) has declared outright that a non-human entity could (or should) be empowered to decide on the employment of lethal force, it would be naïve to think that research and development is not already underway in certain states, including the USA, and that such technology is truthfully and universally considered to be undesirable. Furthermore, from a practical perspective, it is likely to become increasingly difficult to determine how the boundary between meaningful human control and machine autonomy can be universally defined and agreed, when the ideal would be to achieve a perfect unity between human and machine.

Already, designers of military and commercial equipment of all sorts seek to lever the mutual advantages of human and machine to achieve optimum synergy for the overall system. As machines become ever more intelligent and capable, it is likely that some functions currently performed best by humans will ultimately be better performed by machines, thereby releasing humans to exploit their consequently freed capacity to perform extant or new

functions for which their aptitude remains supreme. Indeed, there is fundamentally nothing new in this context: for example, aircraft autopilot systems perform certain functions better than their human counterparts, and history shows us that the exploitation of human and machine synergy has been in constant evolution since Palaeolithic man first hewed a cobble into a hand-axe. So far, it is arguable that this evolution has been constrained to the physical rather than conceptual domain, but is it really inconceivable that artificial intelligence should not supersede human decision making, including those decisions involving lethality, if the relevant technology proves itself to be more competent than human beings in making such decisions? Those, like Sharkey, who oppose the development of LAWS, argue that machines lack the sophisticated intelligence and psychology to understand higher intent or interpret human actions, intentions and emotions in the way that humans do.¹³ But such a standpoint reflects technology as it is now, not as it might be in the future. Furthermore, such a viewpoint fails to acknowledge that all humans are different and are liable to reach different conclusions when faced with identical inputs based on a whole range of subjective factors (including fear, selfishness, fatigue and ideology) – which Arkin explored extensively in his 2010 paper.¹⁴ It is simply misleading, as Sharkey implies, to assume that all humans are of equal virtue, intelligence and character, or that none is susceptible to the debilitating effect on ethical conduct that exposure to combat can promote. 'Designing out' the frailties that lead human combatants to act unethically and illegally should be a primary aim in the development of LAWS, with a commensurate uplift in the ethical conduct of warfare.

The UK Government argues that we might be a very long way off from witnessing the requisite advance in artificial intelligence that could enable such a possibility, but we cannot be sure. It must therefore be questioned whether the UK's policy on the use of remotely operated military systems, at the heart of which has been enshrined the primacy of human decision making, is sustainable indefinitely, or whether, in fact, the development and employment of genuinely autonomous weapons systems are inevitable and, perhaps, even ethically desirable.

Although there is no internationally agreed definition of what constitutes a 'LAWS',¹⁵ it may be understood that in order to be described as truly or fully 'autonomous', rather than simply 'automated', a system must be capable of independently 'interpreting higher level intent'¹⁶ and direction,¹⁷ analysing its physical and operational context in order to make decisions and act independently from further human influence; in the case of fully autonomous weapons systems, these include decisions to employ lethal force. The UK remains sceptical of the feasibility of such systems and categorically states that it 'does not possess fully autonomous weapon systems and currently has no intention of developing them. Such systems are not yet in existence and are not likely to be for many years, if at all.'¹⁸ Indeed, despite the very wide spectrum of opinions on the legalities and ethics of LAWS, there is a general consensus that none are in existence yet. According to Human Rights Watch (HRW), a member of the Campaign to Stop Killer Robots (CSKR), a 'civil society' organisation comprising a number of NGOs,¹⁹ 'Fully autonomous weapons... do not yet exist.'²⁰ Furthermore, it is generally accepted that in-service weapon systems exhibiting a

high degree of automation, such as loitering munitions, the Phalanx close-in anti-shipping-missile system and Israel's 'Iron Dome' anti-rocket surface-to-air missile system, fail to meet the definition of 'full autonomy', because humans programme them to respond within precisely defined parameters to pre-defined conditions. In Phalanx's case, when commanded to automatic mode, it automatically detects and engages sea-skimming supersonic anti-shipping missiles (which humans lack the necessary response time to counter adequately) according to very tightly controlled parameters. Because systems like Phalanx behave in accordance with the explicit programming instructions of humans in reaction to precisely pre-defined circumstances, they are usually defined as 'automated' rather than 'autonomous', although some refer to them as 'partially autonomous'. Noel Sharkey has described the reasoning process of such systems as ultimately rooted in the simple computer programming language of 'the humble IF/THEN statement'.²¹ Whilst Sharkey would argue that such systems are to be considered as 'autonomous', they do not fulfil the requirements set out above of being able to interpret higher level intent and analyse their context beyond the narrow scope of an 'IF/THEN' decision process. They are not endowed with the requisite initiative to respond to factors that lie outside those defined in their programmes. In sum, with apologies to Descartes, such systems do not 'think', therefore they are not [autonomous].



Automated, not autonomous: Phalanx Close-in Weapon System
(Credit: U.S. Navy)

So much for the current state of play. The future viability of LAWS is where opinion begins to diverge comprehensively. Contrary to the UK's position that autonomous weapons systems are 'not likely to be [in existence] for many years, if at all',²² HRW argues that 'weapons technology is moving rapidly toward greater autonomy' paving the way for weapons with the 'power to determine when to take human life'.²³ Despite the wide spectrum of views on the subject, notably the contested term 'greater autonomy', the international community is addressing the legal issues concerning LAWS through the auspices of the UN Office in Geneva's Convention on Certain Conventional Weapons (CCW) annual 'Meeting of Experts on Lethal Autonomous Weapons Systems', the most recent event having taken place between 13 and 17 April 2015. Because it is a diplomatic forum, the UK's lead department for LAWS is the Foreign and Commonwealth Office (FCO), supported by the MOD. In addition to state and UN representation, other participants in the forum include the Campaign to Stop Killer Robots and the International Committee for Robot Arms Control (ICRAC). So far there has been little tangible progress towards achieving international agreement on LAWS, even over the definition of the term. All parties agree, however, that contemporary technology is incapable

of producing systems with the required artificial intelligence to meet the broadly agreed understanding of what a truly autonomous system is, i.e. although a degree of autonomy can be achieved through the automation of certain functions of a weapons system, they are as yet incapable of exercising reasoning and judgement to the same sophisticated level as a human being. In these regards, humans continue to outperform machines and, in the view of the ICRC, a supersession by machines is ‘unlikely to be possible in the foreseeable future.’²⁴ Consequently, although highly automated systems have been demonstrated to perform well in highly predictable circumstances, so far not even the most complex ‘autonomous’ system has yet exhibited the power of judgement necessary to adapt satisfactorily to complex, dynamic and unexpected circumstances; moreover, as a consequence, when faced with the unpredictable, state-of-the-art ‘autonomous’ machines can behave unpredictably.

For those who fear imminent World domination by Terminator-esque killer robots, it should be reassuring to learn quite how relatively under-developed even the most advanced ‘partially autonomous’ systems are at present. Sharkey argues, with some justification, that ‘The autonomous robots being discussed for military applications are closer in operation to your washing machine than to a science fiction Terminator.’²⁵ It is notable, for example, that one of the most significant milestones so far reached autonomously by an unmanned system was the recent achievement by the US Navy’s X-47B Unmanned Combat Air System Demonstrator (UCAS-D)



The US Navy’s X-47B UCAS-D achieved autonomous in-flight refuelling on 22 April 2015 – a milestone for aviation, but no indicator that machines are capable of out-reasoning humans. (Credit: U.S. Navy)

of in-flight refuelling. Important though this milestone undoubtedly was in terms of extending the range and endurance of unmanned systems and in demonstrating the high technical merit of the machine in performing the delicate manoeuvres inherent in in-flight refuelling, it hardly marked a decisive breakthrough in the race to achieve machine supremacy over human judgement. Indeed, this success served as much to highlight the limits of artificial intelligence as it pointed to its potential.

Hence, in light of the pronounced limitations of current autonomous technology, the debate over LAWS has principally circulated around the issue of whether to introduce a pre-emptive ban on such systems, with groups such as ICRAC claiming that ‘The delegation of violence to a machine – whether lethal or less lethal – is a violation of human dignity.’²⁶ The UK rejects the premise of this argument, stating that it would never delegate the decision to employ lethal force to a machine and that IHL already prohibits their development. As the FCO has

stated, 'Whilst technological advances will likely increase the level of automation in some systems, just as in non-military equipment, the MOD has no intention to develop systems that operate without human intervention in the weapon command and control chain.'²⁷ The UK considers that its stance accords precisely with extant International Humanitarian Law (IHL), which it believes already effectively bans all states from introducing fully autonomous systems. Specifically, Article 36 of Additional Protocol 1 to the 1949 Geneva Conventions obliges states 'to determine whether [a weapon's] employment would, in some or all circumstances, be prohibited by this Protocol or by any other rule of international law applicable to the High Contracting Party.'²⁸ In its interpretation of Article 36, the UK contends that a fully autonomous system would never be capable of meeting the principles of humanity, proportionality and distinction in the targeting process and, therefore, IHL signatory states are compelled to limit weapons systems to those which operate under 'meaningful human control'. Article 1 of the UN's Universal Declaration of Human Rights states that 'All human beings are born free and equal in dignity and rights. They are endowed with reason and conscience and should act towards one another in a spirit of brotherhood.'²⁹ Consequently, under the current provisions of IHL, it can be argued that the principle of humanity is inseparable from the human species; ipso facto, no other living or artificial creation has the right to judge matters involving humanity. But in arguing that humans alone have the right to make decisions that have humanitarian implications, there is an inherent presupposition that either humans are (and always will) remain inherently superior to artificial creations in making judgements based on humanitarian principles, or that human mistakes or misdeeds will remain more admissible than machines' potential inerrancy. The first presupposition is open to conjecture, but in this author's opinion is unlikely to withstand the test of time; the second, ironically, seems almost certainly inconsistent with humanitarian objectives. So far in history, humans have failed consistently to live up to humanity's loftier ideals. Indeed, as Arkin argued in his 2010 paper, '... it seems unrealistic to expect normal human beings by their very nature to adhere to the Laws of Warfare when confronted with the horror of the battlefield, even when trained.'³⁰

To err is, indeed, human, as humanity's sad history of war and its associated crimes have lamentably demonstrated. But to forgive mankind en masse for its propensity for making bad decisions would be an error in itself if artificial intelligence is developed that is better equipped than humans to make better humanitarian decisions. To argue that decisions to employ lethal force should always be made by humans is to argue that ISIL's murderous reign of terror is more acceptable than, in another context, the sparing of a non-combatant by a machine whose 'mind' is unfettered by fatigue, fear, hatred or perverted ideology.

The UK's position is that it cannot envisage a point at which machines will be capable of exercising the principle of humanity enshrined in the Laws of Armed Conflict. Even defence companies exploring the potential of autonomy seem at pains to highlight the involvement of human decision making. BAe Systems, whose Taranis project seeks to employ facets of autonomous behaviour, is scrupulously coherent with this principle, emblazoning its Taranis web page with the emboldened statement 'CONTROLLED BY A HUMAN OPERATOR.'³¹

But proponents of a pre-emptive specific ban on LAWS contend that high levels of automation and autonomy materially influence human operators’ decisions in any case: in effect, they argue, the information presented by the system railroad the operator into taking a particular course of action. Furthermore, proponents of a bespoke ban argue that, without one, there is a danger of a new arms race, lowering the threshold on the use of force and the dilution of discrimination in its application. There are, of course, many shades of opinion on the subject, but none is as well-defined as the UK’s policy. The USA, which is the only state other than the UK to have publicly announced its policy on autonomous and semi-autonomous weapons, has provided some detail on its approach to LAWS, but it is ultimately more ambivalent than the UK regarding its interpretation of weapons: ‘Autonomous and semi-autonomous weapon systems shall be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force.’³² Quite what is meant by the intention to ‘allow appropriate levels of human judgment’ remains unclear, but senior US officials seem much more at ease with the concept that fully autonomous weapons systems will supersede some manned and remotely operated systems. US Secretary of the Navy Ray Mabus recently declared, ‘I’m for a full-up penetrating strike fighter. ...[UCLASS]³³ ought to be the bridge to a full-up strike fighter – an autonomous strike fighter – that [operates] in contested environments.’³⁴ By ‘contested environment’, it is reasonable to assume that Secretary Mabus means one in which not only can the opposition be expected to employ kinetic measures to defeat friendly systems, but one that is contested in electromagnetic terms too, i.e. an environment in which the ability to control a system via satellite link (or any other reliant on the electromagnetic spectrum) is disrupted. It might further be inferred, therefore, that human operator intervention would be severely limited, if not negated entirely, under such operational conditions. Hence, autonomy – i.e. self-reliance and the ability to think – would be vital facets of such a system. US Admiral Darrah went further in a recent interview:

“What we’re doing today is deterministic autonomy...it’s not autonomous” because boundaries and parameters are pre-set for the aircraft, he said. The admiral said the navy would continue using deterministic autonomy until artificial intelligence is capable of operating within the same rule set as humans... The navy is also studying autonomy as it relates to the system’s weapons, the admiral said, including the extent to which a weapon could someday make a targeting determination downrange.³⁵

Admiral Darrah’s statement indicates that the US is not only contemplating, but planning for, the use of LAWS, despite their apparent incompatibility with International Humanitarian Law – at least as the UK interprets it. According to Sharkey, ‘decision making robots... have appeared in all of the US military’s roadmaps since 2004.’³⁶ At least the US is relatively candid about its approach to such systems – other states (including China, Russia and France) are substantially more guarded on their own definitions of LAWS and their interpretations of International Humanitarian Law as applied to automatic and autonomous weapons systems.

From this temporal vantage point, it is uncertain if or when fully autonomous weapons systems will become viable. But let us, for a moment, at least assume that they will become feasible at some future point, either through evolution or revolution in the artificial intelligence and robotics domains. Technical viability will ultimately challenge legality. Any international accord that either confirms that IHL already effectively bans LAWS or introduces a bespoke pre-emptive ban on such systems may deter or delay the development of such systems, but it is unlikely that legislation could be anything more than a speed bump on the road to some form of military employment. Already, it is unclear how most states interpret IHL with respect to LAWS, and it is probable that some would actively pursue such systems as soon as technology facilitates them – the USA seems already to be marching down this path.

Weapons innovation nearly always usurps extant legislation – how, for example, could nuclear weapons ever meet the conditions of proportionality and humanity that IHL enshrines, yet despite this contradiction they continue to form a vital component of several states' military inventories? Moreover, paradoxically perhaps, nuclear weapons are generally (though not universally) considered to have exerted a positive effect on the relative peacefulness of the post-Second World War era. So, despite their seeming incompatibility with IHL, it may be considered that nuclear weapons have (so far, at least) made a positive contribution to peace and, therefore, have reduced the scale of human suffering through war. Consequently, for many states, despite their potentially apocalyptic consequences, nuclear weapons are considered to be peace-positive. Should technology permit, those states that judge LAWS to offer military advantage are likely to argue that such systems are more capable than humans in exercising the lofty principles of human reasoning and judgement than humans themselves because they would not be susceptible to the deleterious effects of anger, fatigue, fear, greed, hatred and pain to which humans are subject. Should LAWS develop to a point where they are capable of practising the highest levels of judgement and reason, unfettered by human frailties, it might reasonably be argued that they would be better equipped than humans to decide on matters concerning the use of lethal force – and to do so consistently. It should need no reminder that each and every crime against humanity has been committed by a human. Is it not humanity's humanitarian responsibility to make LAWS that are more virtuous than humans themselves?

Given the current limitations of artificial intelligence, humans remain best equipped to decide when to employ lethal force. But we already exist in a hybrid world where humans and machines co-exist, exploiting the synergy between the calculative accuracy of machines with human flexibility to deal with multifarious and unpredictable planning conundrums. High automation and partial autonomy have a place in our lives and professions now – the RAF Voyager incident of 2014 served to highlight the life-saving benefits of high automation in aircraft safety systems when humans get it wrong. But neither technology nor humanity are yet at a point where life-taking decisions can be delegated to machines. Nevertheless, whatever the status of IHL, it will take just a few LAWS genies to be released from their technological,

legal and ethical lanterns to revolutionise warfare. Although commendable, the UK's present position on LAWS looks vulnerable to an unpredictable and innovative future. Whether through technological evolution or revolution, it would be unwise to conclude that international law in any form will ultimately prevent the creation of systems displaying a degree of autonomy that draws into significant question the viability and appropriateness of 'meaningful human involvement' in decisions involving the employment of lethal force. Paradoxically, the machines may ultimately be more humane than humans; given humanity's track record, this does not appear to be an impossible or, indeed, undesirable aspiration. Hence, if or when technology matures to the point where machines can be endowed with the ideals of human virtue and the intelligence to interpret their context and higher intent accurately, it is surely advantageous, from both the perspectives of military advantage and ethics, to allow such machines to make lethal decisions. Therefore, rather than seeking to ban such technology or unilaterally withdraw from the development of such systems, it would be better for states to agree to a humanitarian code to which LAWS should adhere – IHL, which humans have proven lamentably incapable of observing, already provides a suitable framework.

Notes

¹ House of Commons Daily Hansard, "Drones in Conflict," www.parliament.uk, October 13, 2015: Columns 292-293, <http://www.publications.parliament.uk/pa/cm201516/cmhansrd/cm151013/debtext/151013-0004.htm#15101371000002>.

² This article considers the use of LAWS within the framework of combat and the Laws of Armed Conflict. It does not address the separate, but often conflated, debate surrounding the use of drones in targeted killings in states which are neither at war with the prosecuting state nor have given permission for such strikes to take place on their sovereign territory. For further insight into this issue, the author recommends reading: Noel Sharkey, "Saying 'No' to Lethal Autonomous Targeting," *Journal of Military Ethics* 9:4 (2010): 370-383.

³ UK Ministry of Defence, "Unmanned aerial systems on show," *UK Ministry of Defence*, last modified December 18, 2013, <https://www.gov.uk/government/news/unmanned-aerial-systems-on-show--7>.

⁴ BBC News, "UK drone fleet to double in fight against IS, says PM," BBC, last modified October 4, 2015, <http://www.bbc.co.uk/news/uk-politics-34436917>.

⁵ David Whetham, "Drones to protect," "Virtuous drones?," *The International Journal of Human Rights*, 19:2 (2015): 199-210.

⁶ Caroline Kennedy and James I. Rogers, "Virtuous drones?," *The International Journal of Human Rights*, 19:2 (2015): 211-227.

⁷ *ibid.*, 222.

⁸ International Humanitarian Law (or the Laws of Armed Conflict) and Rules of Engagement.

⁹ British attitudes towards drones are complex and dependent upon the context in which they are used. Opinion polling has shown support of up to 75% for drone strikes that would kill a known terrorist if no innocent civilians killed at the same time. See: YouGov, "British Attitudes Towards Drones," YouGov, last modified April 3, 2013, <https://yougov.co.uk/news/2013/04/03/british-attitudes-drones-and-targeted-killing/>.

¹⁰ Ronald C. Arkin, "The Case for Ethical Autonomy in Unmanned Systems," *Journal of Military Ethics* 9 (4) (2010): 332-341.

¹¹ Doug Bolton, "Stephen Hawking, Noam Chomsky and thousands of others sign open letter calling for a ban on 'killer robots,'" *The Independent*, last modified July 27, 2015, <http://www.independent.co.uk/life-style-7-4/stephen-hawking-noam-chomsky-and-thousands-of-others-sign-open-letter-calling-for-a-ban-on-killer-10420169.html>.

¹² Ben Goertzel, "Human-level artificial general intelligence and the possibility of a technological singularity. A reaction to Ray Kurzweil's *The Singularity Is Near*, and McDermott's critique of Kurzweil," *Artificial Intelligence* 171 (2007): 1161–1173.

¹³ Sharkey, "Saying 'No!' to Lethal Autonomous Targeting," *Journal of Military Ethics* 9:4 (2010): 370.

¹⁴ Arkin, "The Case for Ethical Autonomy in Unmanned Systems."

¹⁵ For Conference purposes, the UN and International Committee of the Red Cross used the following working definition of Autonomous Weapons System: 'Weapons that can independently select and attack targets, i.e. with autonomy in the 'critical functions' of acquiring, tracking, selecting and attacking targets'. (International Committee of the Red Cross, "Report of the ICRC Expert Meeting on 'Autonomous weapon systems,'" May 9, 2014, <https://www.icrc.org/eng/assets/files/2014/expert-meeting-autonomous-weapons-icrc-report-2014-05-09.pdf>).

¹⁶ Correct interpretation of 'higher level intent' relies on that intent being clearly and unambiguously expressed. The perils of poorly expressed commander's intent were, for example, famously illustrated at the Battle of Balaclava when Lords Lucan and Cardigan misinterpreted their Commander Lord Raglan's ambiguous intent, and led the Light Brigade to disaster. Even if LAWS are developed that can match humans' ability to interpret commanders' intent, that interpretation will, to a very large degree, remain reliant on the unambiguous articulation of that intent.

¹⁷ UK MOD Development, Concepts and Doctrine Centre, *JDN 2/11 The UK Approach to Unmanned Aircraft Systems* (UK: MOD, 30 March 2011), 2-3. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/33711/20110505JDN_211_UAS_v2U.pdf (accessed 2 June 2015). The JDN 2/11 definition of an autonomous system is: 'An autonomous system is capable of understanding higher level intent and direction. From this understanding and its perception of its environment, such a system is able to take appropriate action to bring about a desired state. It is capable of deciding a course of action, from a number of alternatives, without depending on human oversight and control, although these may still be present. Although the overall activity of an autonomous unmanned aircraft will be predictable, individual actions may not be.'

¹⁸ Alistair Burt (UK Under Secretary of State for Foreign Affairs), Letter to Kate Allen, Executive Director, Amnesty International UK (June 27, 2013), <http://www.una.org.uk/sites/default/files/Killer%20robots%20-%20reply%20from%20FCO%20-%2027%20June%202013.pdf>.

¹⁹ Human Rights Watch, "Shaking the Foundations – The Human Rights Implications of Killer Robots," Human Rights Watch (2014), http://www.hrw.org/sites/default/files/reports/arms0514_ForUpload_0.pdf.

²⁰ The Campaign to Stop Killer Robots comprises: Human Rights Watch; Article 36; Association for Aid and Relief Japan; International Committee for Robot Arms Control; Mines Action Canada; Nobel Women's Initiative; PAX (formerly known as IKV Pax Christi); Pugwash Conferences on Science & World Affairs; and the Women's International League for Peace and Freedom.

²¹ Noel Sharkey, "Saying 'No!' to Lethal Autonomous Targeting," *Journal of Military Ethics* 9:4 (2010): 377.

²² Alistair Burt, Letter to Kate Allen.

²³ Human Rights Watch, "Shaking the Foundations."

²⁴ International Committee of the Red Cross, "Report of the ICRC Expert Meeting on 'Autonomous weapon systems.'"

²⁵ Sharkey, "Saying 'No!' to Lethal Autonomous Targeting," 379-80.

²⁶ ICRC, "ICRC closing statement to the 2015 UN CCW Expert Meeting," ICRC, 17 Apr 15, <http://icrac.net/2015/04/icrac-closing-statement-to-the-2015-un-ccw-expert-meeting/>.

²⁷ Alistair Burt, Letter to Kate Allen.

²⁸ United Nations, *Protocol additional to the Geneva Conventions of 12 August 1949, and relating to the protection of victims of international armed conflicts (Protocol 1)* (New York: United Nations, June 8, 1977), 21. <https://treaties.un.org/doc/Publication/UNTS/Volume%201125/volume-1125-I-17512-English.pdf>.

²⁹ United Nations, "The Universal Declaration of Human Rights," United Nations, December 1948, <http://www.un.org/en/documents/udhr/index.shtml#a1>.

³⁰ Arkin, "The Case for Ethical Autonomy in Unmanned Systems."

³¹ BAE Systems, "TARANIS," BAE Systems, <http://www.baesystems.com/en/product/taranis> (accessed October 2, 2015).

³² US Department of Defence, *Department of Defense Directive: Autonomy in Weapon Systems* (Washington DC: Department of Defense, November 21, 2012), 2. <http://www.dtic.mil/whs/directives/corres/pdf/300009p.pdf>.

³³ Unmanned Carrier-Launched Airborne Surveillance and Strike.

³⁴ Marina Malenic, "Surveillance or Strike?," *IHS Jane's Defence Weekly* (June 3, 2015): 31.

³⁵ *ibid.*

³⁶ Noel Sharkey, "Saying 'No!' to Lethal Autonomous Targeting," *Journal of Military Ethics* 9:4 (2010): 370.

This article has been republished online with Open Access.

Ministry of Defence © Crown Copyright 2023. The full printed text of this article is licensed under the Open Government Licence v3.0. To view this licence, visit <https://www.nationalarchives.gov.uk/doc/open-government-licence/>. Where we have identified any third-party copyright information or otherwise reserved rights, you will need to obtain permission from the copyright holders concerned. For all other imagery and graphics in this article, or for any other enquires regarding this publication, please contact: Director of Defence Studies (RAF), Cormorant Building (Room 119), Shrivenham, Swindon, Wiltshire SN6 8LA.

 **ROYAL
AIR FORCE**
**Centre for Air and
Space Power Studies**

OGL