Reith, Russell, and the Robots: AI, Warfare, and Shaping the Debate

Keith Scott
De Montfort University, UK
jklscott@dmu.ac.uk

Abstract: On December 8th, 2021, Professor Stuart Russell delivered the second of that year's Reith Lectures, presented under the banner title 'Living With Artificial Intelligence'. This specific talk dealt with 'The Future Role of AI in Warfare', and in this paper I propose a reading of Russell's address which both summarises and critiques his argument and stance, to determine what, if anything, can be taken from his position as effectively a public philosopher and applied in the realm of modern warfare, where ethical questions are taken from the seminar room and enacted in battlespace. The Reith lectures occupy a unique place in public discourse; given each year by a leading figure in the field under discussion, they help to shape opinion and debate. In considering the role of AI, and in particular its deployment in combat, there is undoubtedly a need for multi- and transdisciplinary thought, but the choice of Russell as the lecturer is not unproblematic. He is undoubtedly an expert in the field of AI, but he has no direct experience of working with the military, and is clearly not a neutral witness. He has been a leading figure in the campaign to ban research into autonomous weapon systems, and was closely involved in the production of *Slaughterbots*, a short film which presents a nightmare vision of swarming drones as agents of political repression. There are deep and serious questions to be asked about the role of AI in warfare, but Russell's position that we must stop all research in the field is arguably naïve. Our adversaries will surely not be as punctilious. At the heart of the debate lie complex issues concerning human agency and control (and 'control' lies at the etymological root of 'cyber'); this paper will use Russell's lecture as a starting point for the consideration of how we might develop an ethical doctrine for the use of AI, resting on the idea of human-machine teaming. It will, in short, argue for a cybernetic solution to the problems of cyber warfare.

Keywords: AI, conflict, cyber politics, discourse analysis

1. Introduction

1.1 'Cyber', systems, and society

> *I have spoken of machines, but not only of machines having brains of brass and thews of iron. When human atoms are knit into an organization in which they are used, not in their full right as responsible human beings, but as cogs and levers and rods, it matters little that their raw material is flesh and blood. What is used as an element in a machine, is in fact an element in the machine. (Wiener 1989, p.185)*

The essential problem at the heart of this paper was described by C.P. Snow in his 1959 lecture 'The Two Cultures', where he argues that there is a fundamental division between the two dominant modes of thought in Western society:

> *Literary intellectuals at one pole - at the other scientists, and as the most representative, the physical scientists. Between the two a gulf of mutual incomprehension - sometimes (particularly among the young) hostility and dislike, but of all lack of understanding. (Snow 1961, p.4)*

More recently, the idea of a binary opposition between Humanities and Sciences has been replaced by a more complex and richer vision of societies composed of intersecting, overlapping divisions (Arts/Sciences, male/female, white/people of colour, heterosexual/LGBQT+…), but the essential point remains the same: put crudely, we do not all live in the same world, or view and experience it in the same way. What has this to do with the business of cyber security? My answer is: everything. We work in a discipline which forms a system of knowledge, but that system is constrained by and inscribed within a network of other interacting and often conflicting systems which we must acknowledge and respect if we are to achieve what we wish to achieve. If our work is commercially funded, we must show that it satisfies the commercial imperatives of our employers; if we are conducting government funded research, we must demonstrate that what we do aligns with the political imperatives of those in power. 'Cyber' is never purely about technology; it is about the application of that technology within the world, and the interactions between technology and its users. Cybernetics is, as Wiener put it, about 'control and communication in the animal and the machine'; our work operates within contexts which are social, cultural, and in the widest sense political. It is not enough to innovate; we must *communicate* what our work offers, if it is to be accepted and permitted. This is, for many, a source of frustration, seen as a distraction from our core business, but if we cannot present our work in a way which is socially approved of and accepted, it will be rejected. How our work is presented in the public sphere is of vital importance, not least

because we are living through a period where mistrust of science and of expert knowledge in general is rife (Eyal 2019; Pew Research Center 2019). Here we look at an example of discourse on cyber warfare delivered in the public realm, and consider how it presents a topic of huge concern to both ourselves and the wider world, but does so in a way which arguably fails to engage in a measured consideration of the theme. I contend that we cannot allow discussion of our work in the public sphere to be dominated by any one point of view; debate is essential to democracy and science alike. This will require us to be adept not only at research and policy construction, but at communicating our ideas to stakeholders, the citizenry, and opinion formers. We need to be able to bridge the gap between the Snow's two cultures, and convey scientific and technical thought in engaging writing/speech, as was done by Turing (1950 – as much a philosophical and theological treatise as a founding text of AI studies) or Haldane (1924; 1925). The text I will examine here is just such a work; but its conclusions, I argue, are antithetical to many of our opinions as researchers and practitioners in cyber warfare.

1.2 The Reith Lectures and the partial expert

In 1948, Bertrand Russell delivered the first of the BBC's annual Reith Lectures, instituted in memory of the Corporation's first Director General, who declared the BBC's purpose as being to 'inform, educate, and entertain'. Every year since then[1], a series of intellectuals, politicians, and the great and the good have addressed the listening public 'to advance public understanding and debate about significant issues of contemporary interest' (https://www.bbc.co.uk/radio4/features/the-reith-lectures/about/). The lecturer for 2021 was Stuart Russell, Professor of Computer Science at the University of California, Berkeley. Over four lectures (each with a brief question and answer session), delivered in four different locations in the UK, Russell spoke on the topic of 'Living With Artificial Intelligence'. This paper will consider his second lecture, delivered at the Whitworth Hall in Manchester on November 3rd, on the subject of 'The Future role of AI in warfare'. [https://www.bbc.co.uk/programmes/m00127t9 TRANSCRIPT: https://downloads.bbc.co.uk/radio4/reith2021/BBC_Reith_Lectures_2021_2.pdf]. Russell's talk is interesting, informative, and well-expressed; it is however neither neutral nor impartial, and this aspect of the talk is problematic, given the role of the BBC as a shaper of opinion in the English-speaking world.

The BBC occupies a unique position in British society; viewed as the voice of the Establishment, it is the State broadcaster, but independent of the government, funded neither by a levy nor taxation, but through the licence fee paid by all consumers of its product. It has a worldwide reputation for neutrality and balance, which is arguably accurate, given that it has been attacked throughout its history from both the Left and Right. However, there are two issues that must be considered when looking at supposed impartiality. The first is what Robin Lustig (2017) terms ''bothsidesism', because it subscribes to the principle that both sides of an argument always require a hearing'; this leads to a gross imbalance, when, for example, a sceptic of androgenic climate change is given equal airtime to a representative of the view held by the overwhelming majority of scientists (equality does not equal equity). Secondly, the BBC's own Editorial Guidelines make the following key points:

> *The BBC is committed to achieving **due** impartiality in all its output. This commitment is fundamental to our reputation, our values and the trust of audiences. The term 'due' means that the impartiality must be adequate and appropriate to the output, taking account of the subject and nature of the content, the likely audience expectation and any signposting that may influence that expectation.*
>
> *Due impartiality […] does not require absolute neutrality on every issue*
>
> *[…]*
>
> *Audiences expect artists, writers and entertainers to have freedom to explore subjects from one perspective and to create content that reflects their own distinctive voice. **It must be clear to audiences where personal views are being expressed.***
>
> *(https://www.bbc.co.uk/editorialguidelines/guidelines/impartiality. My emphasis)*

- in other words, *due* impartiality ≠ *true* impartiality. When we turn to examine Stuart Russell's credentials as a Subject Matter Expert, it is clear that he occupies a very specific position with regard to the use of AI in warfare, and in particular the development and possible use of autonomous weapons systems. He gained his PhD in 1986, and since then has gone on to become a major figure in AI research, cowriting one of the key textbooks in the

[1] With the exception of 2015, when Stephen Hawking's lectures were postponed through illness, and 1992, when "the BBC simply couldn't find anyone to do them".

field (Norwig and Russell 1986) and producing a huge output of academic research (people.eecs.berkeley.edu/~russell/publications.html). Of direct concern to us, however, is the extent to which he has become one of the major figures warning of the risks of autonomous weapons systems. Looking at the list of his articles in non-academic texts dealing with Lethal Autonomous Weapons Systems (LAWS) (people.eecs.berkeley.edu/~russell/research/LAWS.html) his stance on the topic is self-evident. While not a member of the Campaign to Stop Killer Robots (www.stopkillerrobots.org/) he has worked and continues to work with them, and has been described (albeit by a non-neutral source) as an 'influential fellow-traveller' (Freedberg 2019). He has also narrated and helped co-create the 2017 short film *Slaughterbots*, which presents a vision of autonomous killer microdrones, and helped promote its 2021 sequel, *Slaughterbots: If Human: kill()*. He is, in short, not a neutral voice. This is not to say, however, that he is of necessity wrong in his opinions, nor that he should not be allowed as wide a platform as possible to air his views. He is an intelligent, incredibly-well-informed expert in the field, even if the depiction of AI weaponry depicted in *Slaughterbots* takes a deliberately extreme position, owing more to the dystopian visions of a series like *Black Mirror* than a reasoned analysis of the future of autonomous ordnance. However, by giving him the chance to deliver the Reith lecture on this topic, he becomes the de facto expert voice on the subject, and the idea that there might be opposing positions is pushed to the background. Furthermore, by choosing a speaker who does not come from the military domain, or a related research area such as the UK's DSTL[2] or America's DARPA, the myth is perpetuated that there is no consideration of these issues within the armed forces; such is clearly note the case, as evidenced by works such as *The Tallinn Manual* and *Tallinn 2.0* or the *UK Joint Concept Note 1/18:*

Human-Machine Teaming. While anecdote is not evidence, on the basis of conversations with members of the UK military working in the cyber domain, the author has seen no evidence whatsoever of a desire to introduce LAWS in the battlespace.

For those who are concerned that the public debate over AI in warfare should be grounded in presenting as full a picture of the facts as possible, then, the choice of Russell as effectively the unchallenged voice of authority is a source of some concern. When we turn to look at the content of his lecture, a further range of questions appears.

2. Analysis

"Yeah, well, you know, that's just, like, your opinion*, man" (The Big Lebowski)*

A lecture such as this can in effect be viewed as an influence operation; the speaker is seeking to persuade the audience (both in the lecture hall and beyond). He or she will draw on the fundamental persuasive tools of oratory and rhetoric, making use of ethos, pathos and logos (the appeals to personal character, emotion, and reason) to present their message in the way that best convinces their audience. As stated above, the BBC operates a principle of due impartiality, but the lecture as recorded strives for balance (or is presented as so doing). The entire recording runs for just under an hour, and while (as might be expected), Russell speaks for the majority of the time, his is not the only voice heard. The lecture is introduced by the broadcaster Anita Anand, who also questions the speaker, and chairs the subsequent discussion and question and answer session. There are contributions (questions and points of information) from a number of individuals, of whom a number are identified by affiliation/area of expertise, as shown below:

Table 1: Responses to Russell; identification and airtime

	Affiliation/Expertise	Airtime
Chi Chi Ekweozer	'tech entrepreneur and also a start-up founder'	0'16'' (0.46%)
Air Commodore David Rowland	Head Of Strategy, Information Systems and Services , Joint Forces Command at UK Ministry of Defence	1'20'' (2.30%)
Phil Horn	Unknown	0'30'' (0.86%)
Kathy New	Unknown	0'25'' (0.72%)
Dr Steven Meers	Head of AI Lab, DSTL	1'42'' (2.94%)
Ruben Boss	Unknown	0'23'' (0.66%)
Professor David Balman	Unknown	0'19'' (0.55%)
Dr Keith Dear	Director of AI Innovation at Fujitsu Defence and National Security.	1'41'' (2.91%)
Jo Hooker	Unknown	0'20'' (0.58%)
Virginia Watson	Unknown	0'15'' (0.43%)

[2] Defence Science and Technology Laboratories; a representative from DSTL does participate in the programme.

The range of contributors is clearly designed to present a diversity of opinions, and to address concerns about possible bias. However, there are a number of issues which must be considered when considering just how full a picture of the subject is presented:

- 1. we are presented with a recording of the event; what content might have been edited out before broadcast?
- 2. How were the questioners/participants selected, and by whom?
- 3. How was the audience selected? Is it a random group of participants, or invited as a representative sample of the population?
- 4. What proportion of the programme is given over to participants other than Russell? He speaks for **42'41''** (**73.72%** of the airtime); as the table above shows, other participants are given much less time to advance their views.

The programme, of course, is titled a 'lecture', not a 'debate', but the chosen format gives an impression of a free exchange of views and opinions; this impression is not matched by the facts. This is not in any way meant to imply a deliberate attempt to deceive the listener, but we need to remember the difference between *due* and *true* impartiality.

Any individual's knowledge is invariably partial, in both senses of the word; it is incomplete, because no one person can know all there is to know on a subject, and it is prejudiced, because all human beings embody a unique combination of personal and sociocultural beliefs – we are all biased, however much we may strive not to be. The problem with this lecture is that masks the inescapable prejudice of the single point of view by giving the appearance of being a level playing field for discussion and debate; the talk time analysis shows this is not the case. At a structural level, the programme is fundamentally skewed towards Russell's point of view.

If we examine the actual *content* of the lecture, the same problem arises, from the very title of the programme. We are promised a discussion of 'The Future role of AI in warfare'; what we actually get, is in effect 'Why Lethal Autonomous Weapons Systems are a Bad Thing'. Now, there is absolutely nothing wrong with a talk on this subject, but it should be badged as such and the speaker's stance declared openly from the onset. To do otherwise is misleading in the extreme.

It is arguably asking too much to ask a single speaker to deliver an informed overview of the topic at hand in an hour (why, then, not have a panel discussion or debate?); however, even within the limitations of a single speaker, we can see just how much Russell fails to discuss, or even mention. The listener is left with the impression that LAWS represent the only issue concerning the use of AI in combat, and this is a complete misrepresentation.

Even the most cursory examination of a selection of the material dealing with the implementation of AI in a military context shows that LAWS do not represent the area where it is most likely to be deployed, for both ethico-legal and practical reasons (Freedberg 2018**b**; Maxwell 2020). Much of the implementation of AI will be in areas which are much less obviously dramatic than LAWS, but arguably much more important in terms of delivering an improved combat capability.

The 2018 US DOD Summary of AI Strategy (p.11) identifies a number of key areas where it foresees 'the fielding of AI systems that augment the capabilities of our personnel by offloading tedious cognitive or physical tasks and introducing new ways of working'. These can be grouped into two principal categories:

2.1 Logistics

Central to the effectiveness of any combat force is the so-called 'tooth-to-tail ratio' (3TR), i.e. how many troops are required to support that element of the force directly involved in combat. As McGrath (2007) shows, throughout history, the number of troops making up the 'tail' has far outnumbered the 'teeth' (see **Table 2** below).

Table 2: Overall US Combat Forces Proportions (Source: McGrath 2007, p. 105)

	World War 1 1918	World War 2 (ETO)	Korea 1953	Vietnam 1968	Cold War 1974 Germany	KTO 1991	Iraq 2005
Total Army Strength	3. 514, 137	4, 308, 114*	1, 533, 615	1, 570, 343	783, 330	572, 423	468, 578
Combat	984, 116	828, 402	115, 581	113, 820	51, 170	99, 781	52, 768
Percent	28	19	7.5	7.2	6.5	17	11

* 72% of the Army's 1945 strength (the proportion of the Army in the ETO)

Jacobs (2020) discusses the benefits of AI-enabled systems for improving logistics in the commercial sector, and writers discussing a specifically military context (*inter alia* Freedberg 2018**a** and **b**, Abadicio 2019, and Scharre 2019) make similar observations. Abadicio identifies 5 key areas where AI offers real potential for improving the 3TR (preventive maintenance, cloud services, supply chain management, medical aid and driverless resupply). This last area (driverless resupply) is arguably much more likely to occur in the near to mid-term future than Russell's vision of autonomous swarming 'killerbots'. As Freedberg (2018**b**) observes:

> *Current battery technology may only be able to power ground robots and exoskeletons for less than a full day — but at a fuel dump or ammo depot, unlike a forward patrol, you can just plug them into a diesel generator. Current AI may not be able to navigate around potholes or landmines without a human guide — but at a forward base, unlike the front line, you can just bulldoze the ground flat and mark obstacles with reflective tape or radio-frequency ID tags. And, of course, as with any new technology, robots and exoskeletons might break down unexpectedly — but in a maintenance unit, unlike an infantry squad, your mechanics are right there to fix it and the enemy isn't right there to kill you before it's fixed.*

Technological innovation (particularly within large, hierarchical organizations) tends to be iterative, gradual, and cautious. In the case of the introduction of AI, this is surely a welcome factor, as it allows a proper process of evaluation and review. Russell presents a nightmare scenario where the military rushes into the introduction of LAWS, taking no account of existing legal and ethical frameworks, let alone technological limitations. As Freedberg (2018**b**) puts it:

> *artificial intelligence, for all its rapid progress, is still pretty terrible at coping with uncertainty, ambiguity, and limited information — the very essence of frontline combat. The problem is doubly difficult for ground troops: Drones flying through empty air face a much simpler problem than robots rolling along on land, where mud, dust, and ditches bedevil AI long before an actual enemy shows up. And the machine-learning techniques that let computers beat humans at chess or go may not translate to the chaos of combat, where the playing pieces aren't color-coded black and white, the terrain isn't a nice flat grid, and your opponent doesn't wait for you to take your turn.*

It is of course entirely probable that adversaries (whether states or VNSA) will be less punctilious (as Russell observes, a Kargu-2 drone has already been used in Libya - see Wadhwa and Salkever 2021 – and the risk of LAWS use by terrorist groups is a real threat[3]). However, for Western forces, who are bound by existing legal frameworks relating to armed conflict, and complex and protracted procurement processes, it seems more likely that the application of AI will be more concerned, for the foreseeable future, with supply chains than 'slaughterbots'.

2.2 Decision-Making and cognitive shift

> *The most profound applications of AI are likely to be in information processing and command and control. Just as industrialization changed the physical aspects of warfare, artificial intelligence will principally change the cognitive aspects of warfare. Militaries augmented with AI will be able to operate faster and with more numerous systems, and conduct more complex and distributed operations.*
>
> *(Scharre (2019), p.3)*

We are entering an era when we will be able to capture data from a huge range of sensors in battlespace, from missile telemetry and satellite imagery to fuel levels of planes and tanks, to video feeds from individual soldiers'

[3] See Ware 2019 and Pledger 2021 for discussion of current activity and future possibilities in this area.

bodycams; the 'Internet of Battlefield Things' offers huge opportunities for harnessing AI to process this (Zhu, Majumdar, and Ekenna, C. 2021; Kott, Anathram, and West 2016). No single human (or arguably group of humans) would be able to sift through this data rapidly enough to identify what is significant and act upon it in a battlefield environment; information processing is what AI is designed to do, and there is a growing body of research (see Blaha 2018; Rasch, Kott, and Forbus 2003; Smallegange, Bastiaansen, Venema, and Bronkhorst 2018; Szabadföldi 2021) examining the potential use of AI as a tool to help in decision-making in combat. Dewees, Umphres, and Tung (2021) suggest that such a system could support (but *not* supplant) human decisions by taking the following steps:

- 1. *list alternatives* available to decision-maker
- 2. *define the relevant possible results* for each alternative
- 3. *make a conditional probability estimate* for each alternative
- 4. *present a value trade-off*, a cost-benefit analysis for each possible choice.

The ability to process large volumes of data and generate possible outcomes also raises the possibility of more accurate battle planning, where various scenarios can be run and rerun virtually before entering combat, and/or to analyse previous actions to generate more effective future strategy and tactics (Kerbusch, Keisjer, and Smit, S. 2018). At no point does Russell even raise this as a possibility; he is so fixated on the idea of AI controlling assets in battlespace that he does not consider the idea of using it as a planning tool, or a means of modelling combat to *reduce* death tolls, but bringing a conflict to a more rapid conclusion. Moreover, he neglects the repeated idea that runs through writing on AI in warfare, that argues that it is best used to aid, rather than replace human control. As Scharre (2019) says, the 'most effective militaries are likely to be those that optimally combine AI with human cognition in so-called centaur approaches' (p.4).

3. Conclusion

I have attempted to show that the view of AI in combat presented by Russell is highly partial and fundamentally flawed; rather than seeking to give a balanced and comprehensive view of the terrain, his vision is monocular. It aims to close off, rather than promote, debate. Given the role of the Reith Lectures in setting out the 'approved version' of the topic under discussion, this is a truly missed opportunity. I reiterate the point made above; it is not that Russell is *wrong* when he discusses the risks of LAWS; while he may be discussing the risks of a currently non-existent threat, he is laying out an ethical position which is worthy of respect, and which seems likely to be of real value in coming years (although his argument that we should not research LAWS denies the possibility of researching how to combat them). I would not also expect him to cover all possible aspects of the ostensible subject in the space of a single lecture; my main concern is that the BBC has not taken the opportunity to launch a true debate on the subject, and promote a plurality of opinions. As it stands, any response to Russell's lectures will inevitably be seen as a response to his argument, and this will merely lead to an adversarial approach, rather than a dialogue. In approaching the use (or non-use) of AI in combat, we need a pooling of opinions, from public perception to expert knowledge; we need, in short, a truly 'cybernetic' approach (it would be interesting to see what would emerge from a discussion of the subject following the 'Syntegrity' model proposed by Stafford Beer (1994)). The origins of 'cyber' lie not just with Norbert Wiener, or Ross Ashby, or Gordon Pask, or any of the pioneers of a discipline named 'cybernetics', they lie in the discussions of the heterogeneous group of individuals – including Wiener, Von Neumann, Shannon, Margaret Mead and Gregory Bateson – who participated in the Macy Conferences (Heims 1991). This series of meetings presents a model for bringing together a range of worldviews to create a 'hybrid field of knowledge existing between and within disciplines' (Kline 2015, p.3). AI-enabled combat, it is argued, will be based on the interaction of human and machine; how we decide to conduct it must be based on the bringing together of different spheres of human thought.

References

Abadicio, M. (2019). 'Artificial Intelligence for Military Logistics – Current Applications'. *emerj.com*. 30 April. [online]. https://emerj.com/ai-sector-overviews/artificial-intelligence-military-logistics/.

Beer, S. (1994). *Beyond Dispute: The Invention of Team Syntegrity*. Chichester: John Wiley.

Blaha, L.M. (2018). 'Interactive OODA Processes for Operational Joint Human-Machine Intelligence'. The Specialist Meeting Big Data & Artificial Intelligence for Military Decision Making, NATO, Bordeaux 30-1-21.

Dewees, B., Umphres, C., and Tung, M. (2021). 'Machine Learning and Life-and-Death Decisions on the Battlefield' *Warontherocks.com*, 11 January. [online]. https://warontherocks.com/2021/01/machine-learning-and-life-and-death-decisions-on-the-battlefield/.

Eyal, G. (2019). *The Crisis of Expertise*. New York: Wiley.

Freedberg, S.J. (2018**a**). 'AI Logistics Let Combat Units Move Faster: Uptake's DIUX Contract'. . *Breakingdefense*, 27 June. [online]. https://breakingdefense.com/2018/06/ai-logistics-can-speed-up-army-tanks-uptakes-diux-contract/.

Freedberg, S.J. (2018**b**). 'Forget the Terminator: Robotics for Logistics 1st, Combat 2nd'. *Breakingdefense*, 4 December. [online]. https://breakingdefense.com/2018/12/forget-the-terminator-ii-robotics-for-logistics-1st-combat-2nd/.

Freeman D. et al (2022).'Coronavirus conspiracy beliefs, mistrust, and compliance with government guidelines in England'. *Psychological Medicine* 52, 251–263.

Haldane, J.B.S. (1925). *Callinicus: a Defense of Chemical Warfare*, Kegan Paul: London.

Haldane, J.B.S. (1924). *Dedalus or science and the future* Kegan Paul: London.

Heims, S.J. (1991). *The Cybernetics* Group. Cambridge, MA: MIT Press.

Jacobs, T. (2020). 'Artificial Intelligence (AI) in Supply Chain & Logistics Supply– commercial sector'. *Throughput, inc.* 23 January. [online]. https://throughput.world/blog/topic/ai-in-supply-chain-and-logistics/.

Kerbusch, P., Keisjer, B., and Smit, S. (2018). 'Roles of AI and Simulation for Military Training', The Specialist Meeting Big Data & Artificial Intelligence for Military Decision Making, NATO, Bordeaux 30-1-21.

Kline, R. R. (2015). *The Cybernetics Moment: Or Why we Call our Age the Information Age*. Baltimore: John Hopkins Univ. Press.

Kott, A., Ananthram, S., and Bruce J. West, B. J. (2016). 'The Internet of Battle Things'. *Computer* 49.12 (2016): pp. 70-75.

Lustig, R. (2017). ' Robin Lustig: The BBC has confused balance with impartiality, it should never be impartial between good and evil'. *Press Gazette*, 20 February. [online]. pressgazette.co.uk/robin-lustig-the-bbc-has-confused-balance-with-impartiality-it-should-never-be-impartial-between-good-and-evil/.

Maxwell, P. (2020). 'Artificial Intelligence is the Future of Warfare (Just Not in the Way You Think)', *Modern Warfare Institute*, April 20. [online]. https://mwi.usma.edu/artificial-intelligence-future-warfare-just-not-way-think/.

McGrath, J.J. (2007). *The Other End of the Spear: The Tooth-to-Tail Ratio (T3R) in Modern Military Operations.* Combat Studies Institute Press: Fort Leavenworth.

Ministry of Defence. (2018). *UK Joint Concept Note 1/18: Human-Machine Teaming*. Shrivenham: Development, Concepts and Doctrine Centre.

Pew Research Center, August 2019, "Trust and Mistrust in Americans' Views of Scientific Experts".

Pledger, T.G. (2021), 'The Role of Drones in Future Terrorist Attacks'. Land Warfare Paper No. 137, Association of the United States Army.

Rasch, R., Kott, A., and Forbus, K. (2003). Incorporating AI into Military Decision Making: An Experiment. *IEEE Intelligent Systems* 18 (4), pp.18-26.

Russell, S. (2021). *BBC Reith Lectures 2021 – Living With Artificial Intelligence*, Lecture 2, 'The Future Role of AI in Warfare', BBC Radio 4, 8 December. https://www.bbc.co.uk/programmes/m00127t9. TRANSCRIPT: https://downloads.bbc.co.uk/radio4/reith2021/BBC_Reith_Lectures_2021_2.pdf.

Scharre, P. (2019). 'Military Applications of Artificial Intelligence: Potential Risks to International Peace and Security'. Stanley Center for Peace and Security, Muscatine, Iowa.

Schmitt, M.N., ed. (2013). *Tallinn Manual on the International Law Applicable to Cyber Warfare*. Cambridge: Cambridge University Press.

Schmitt, M.N., ed. (2017). *Tallinn Manual 2.0 on the International Law Applicable to Cyber Operations*. Cambridge: Cambridge University Press.

Smallegange, J.A.P., Bastiaansen, H.J.M., Venema, A.P., and Bronkhorst, A.W. (2018). 'Big Data and Artificial Intelligence for Decision Making: Dutch Position Paper'. Paper, The Specialist Meeting Big Data & Artificial Intelligence for Military Decision Making, NATO, Bordeaux 30-1-21.

Snow, C.P. (1961). *The Two Cultures and The Scientific Revolution*. Cambridge: Cambridge University Press.

Szabadföldi, I. (2021). 'Artificial Intelligence In Military Application – Opportunities And Challenges'. *Land Forces Academy Review* Vol. XXVI, No. 2(102), pp.157 –65.

Turing, A.M. (1950). 'Computing Machinery and Intelligence'. *Mind*, Volume LIX, Issue 236, October 1950, Pages 433–460

US Department of Defense (2018). SUMMARY OF THE 2018 DEPARTMENT OF DEFENSE ARTIFICIAL INTELLIGENCE STRATEGY: Harnessing AI to Advance Our Security and Prosperity. Washington: Department of Defense. https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMARY-OF-DOD-AI-STRATEGY.PDF.

Wadhwa, V, and Salkever A. (2021). 'Killer Flying Robots Are Here. What Do We Do Now?'. *Foreign Policy*, 5 July. [online]. https://foreignpolicy.com/2021/07/05/killer-flying-robots-drones-autonomous-ai-artificial-intelligence-facial-recognition-targets-turkey-libya/.

Ware, J. (2019). 'Terrorist Groups, Artificial Intelligence, and Killer Drones'. *warontherocks.com*, 24 September. [online]. https://warontherocks.com/2019/09/terrorist-groups-artificial-intelligence-and-killer-drones/

Wiener, N. (1989). *The Human Use of Human Beings*. London: Free Use Books. First published 1950.

Zhu, L., Majumdar, S., and Ekenna, C. (2021). 'An invisible warfare with the internet of battlefield things: A literature review'. *Human Behavior and Emerging Technologies*, Volume 3, Issue 2 (April), pp. 255-60.