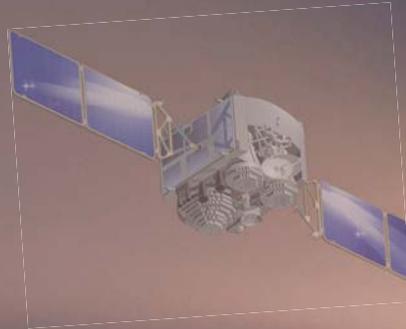




Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Federal Department of Defence,
Civil Protection and Sport DDPS
armasuisse
Science and Technology S+T



Defence Future Technologies

What we see on the horizon



Under the supervision of Dr. Quentin Ladetto

November 2017

The Impact of Autonomous Weapons Systems on International Security and Strategic Stability

Autonomous weapons systems (AWS) are for most people akin to science fiction. Although fully autonomous weapons systems do not yet exist, recent rapid progresses in artificial intelligence compels us to think about the potential impact of this weapon on the conduct of war, international security and international stability. This paper starts by reviewing what artificial intelligence is all about, its recent progress and then looks at the likely impact of AWS on strategic stability. It argues that AWS will likely be very destabilizing for international stability because they favour the offensive. It follows that strategies of pre-emption are likely to emerge to thwart the use of AWS.

Keywords: Artificial intelligence, proliferation, strategic stability, autonomous weapons, swarming, deterrence, pre-emption

Author: Dr. Jean-Marc Rickli, Geneva Centre for Security Policy (GCSP)

What is artificial intelligence?

Artificial intelligence is not new and has been developed since the 1950s when computer scientists rallied around the term at the Dartmouth Conferences in 1956. In the words of the founders of the discipline of artificial intelligence, AI represents the ability of “making a machine behave in ways that would be called intelligent if a human were so behaving”. [1] In other words it is the capability of a computer systems to perform task that normally require human intelligence.

For a very long time, AI was not considered seriously because scientific advances were almost inexistent. In the words of Sergey Brin, Google’s cofounder, “I didn’t pay attention to it (AI) at all. Having been trained as a computer scientist in the 1990s, everybody knew that AI didn’t work. People tried it, they tried neural nets and none of it worked.” Fast-forward a few years and Brin says that AI now “touches every single one of Google’s main projects, ranging from search to photos to ads...everything we do. The revolution in deep nets has been very profound, it definitely surprised me, even though I was sitting right there.” [2] If the development of AI is a surprise for specialist, imagine how disruptive it is for policy makers.

Two technological developments brought AI to the fore. Firstly, advances in the miniaturization of transistors, the building blocks of modern computer hardware, have allowed the doubling of the computing power in about every 18 months since the 1960s. This is known as the Moore’s law. The size of the current smallest transistors is 14nm with 10nm semiconductors expected this year or in 2018. [3] 5nm seems to be a physical barrier but scientists are already working on bringing this limitation down to 1nm. As a way of comparison, the diameter of human hairs varies from 0.017 to 0.18 millimeters. This represents more than 1000x the size of current transistors.

Secondly, the number of data generated on a daily basis has exploded notably with the emergence of mobile and connected devices. It is estimated that 2.5 exabytes are produced every day. This is the equivalent of 250’000 libraries of Congress or 530’000’000 songs. [4] With the rise of the Internet of Things (IoT) it is estimated that 8.4 billion connected things will be in use worldwide by the end of 2017 and 20.4 billion by 2020. [5] These will generate 44 zetabytes of data which is the equivalent of 5’200 gigabytes for every individuals on earth. [6] By way of comparison, 1 gygabyte can hold the content of about 10 meters of books on a shelf, so 5’200 GB is about 52km of books.

The combination of increasing computing power and data has allowed making some breakthroughs in artificial intelligence notably by the application of machine learning [7] techniques such as artificial neural networks and especially deep learning. Deep learning represents a class of machine learning algorithms that are inspired by our understanding of the biology of our brains. Deep learning processes data through different layers of the neural network where at each step information is extracted.

As this learning process is very data intensive, it is only recently with the explosion of big data that it became possible to run massive amounts of data through the system to train it. Thus, major milestone in AI were met recently.

In 2015, Baidu, Microsoft and Google managed to sort million of images with an error rate inferior to 5% which is the typical human error rate. [8] This is all the more remarkable as the best algorithm had an error rate of 28.2 percent in 2010 and is now down to 2.7 percent in 2017. [9] In 2016, Google Deepmind, created the AlphaGo algorithm and beat the second best player of the game of Go in four out of five games. Go is considered as the most complex board game as there are more positions (10^{p170}) than atoms in the universe (10^{p80}). [10] In January 2017, Libratus, an algorithm developed by Carnegie Mellon University, played more than 150’000 hands in no-limit Texas Hold ‘Em Poker against four of the best world’s human players. In the end, Libratus won \$1’776’250. The player who lost the less lost \$85’649 and the biggest loser lost \$880’087.

The innovation with current algorithms is that unlike Deep Blue that was trained to beat Kasparov at chess through brute computational force to evaluate millions of positions, current algorithms are learning and are general purpose frameworks that can be applied to different issues. For instance, AlphaGo was then used by Google to manage power usage. It reduced by 40 percent the amount of electricity needed for cooling and that translated into a 15 percent reduction in overall power saving. [11] “By using reinforcement learning techniques and enough simulations, these algorithms are able to increasingly learn new tasks on their own. [12] Thus, while AlphaGo learned the game by analyzing 30 million Go moves from human players, Libratus learned from scratch. Libratus started by “playing at random, and eventually, after several months of training and trillions of hands of poker, it reached a level where it could not just challenge the best humans but play in ways they couldn’t – playing a much wider range of bets and randomizing these bets, so that rivals have more trouble guessing what cards it holds.” [13] Both AlphaGo and Libratus ended up playing in a very different way that humans play the game.

Unlike chess board games, poker is a game with incomplete information. Because of the possibility of bluffing at poker, “an AI player has to randomize its actions so as to make opponents uncertain when it is bluffing.” [14] There is here no single optimal move. Because of this, Libratus is a milestone in AI as a machine can now out-bluff a human. Moreover, algorithms such as Libratus will be able to play a role with everything from trading to cybersecurity, political negotiations and warfare.

The Militarization of AI

The militarization of artificial intelligence is translated in increasing autonomy in weapon systems. According to the US Department of Defense, “to be autonomous, a system must

have the capability to independently compose and select among different courses of action to accomplish goals based on its knowledge and understanding of the world, itself, and the situation.” [15] Autonomous systems can be of two categories: autonomy at rest and in motion: Whereas the former operate virtually such as software, the latter operate in the physical world such as robots or autonomous vehicles. [16] To be fully autonomous, autonomous systems in motion should fulfill three functions. They should move independently through their environment to arbitrary locations; select and fire upon targets in their environment and create and or modify their goals, incorporating observation of their environment and communication with other agents. [17] A recent study by Heather Roff and Richard Moyes of the Global Security Initiative at Arizona State University mapped 256 systems that contain some features of autonomous weapons systems (Figure 1). [18] The study shows that features of mobility (homing and navigation) are those that proliferated the most since they are also the oldest one. Target acquisition and identification technologies are the next technologies as they are a key component of early offensive systems especially air-to-air missiles. Self-engagement technologies such as target image discrimination and loitering are the most recent emerging technologies. Target image discrimination has been boosted by recent improvement in image recognition as mentioned previously. Though this technology is currently on a low number of deployed systems, this feature is prioritized in the development of new weapon systems such as LRASM and TARES. [19]

Loitering technologies are being incorporated in new unmanned aerial vehicles (UAVs) such as nEUROn and Taranis to make them persistent. [20] The combination of target discrimination combined with loitering represents as rightly stated by Roff and Moyes, “a new frontier of autonomy, where the weapon does not have a specific target but a set of potential targets and it waits in the engagement zone until an appropriate target is detected.” [21] This represents a step towards offensive autonomous weapons.

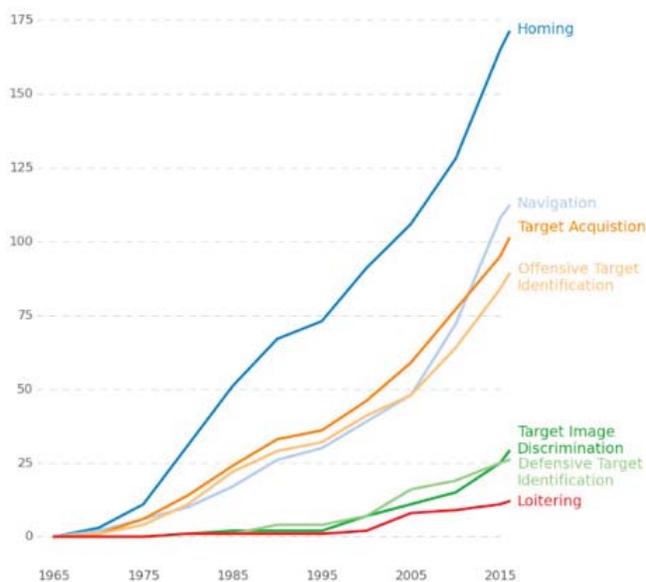


Figure 1: Development of various autonomous technologies over time (Roff and Moyes) [22]

These developments demonstrate that although fully autonomous weapons systems do not yet exist, they are probably within less than a generation reach. Indeed, artificial intelligence and the associated technologies of the so-called Fourth Industrial Revolution are growing exponentially because they “can build on and diffuse over the digital networks” developed by the Third Industrial Revolution. [23] It is therefore crucial to think about their potential impact on international security and warfare. The next part looks at the impact of AWS on strategic stability and proliferation.

The impact of AWS on strategic stability

One crucial aspect of the development of new weapons for international security is their impact on strategic stability. Strategic stability refers to “the condition that exists when two potential adversaries recognize that neither would gain an advantage if it were to begin a conflict with the other.” [24] Strategic stability is neither simple nor static, but should be viewed broadly as a result of effective deterrence.

Throughout history, the major military modes of achieving deterrence has been to build a military force large enough to establish the credibility of threatened punishment if vital interests are impinged upon.

During the Cold War, deterrence was thought of as a mechanism to prevent the opponent’s use of nuclear weapons through the threat of retaliation. Maintaining a second strike capability was the cornerstone of the doctrine of Mutual Assured Destruction (MAD) and also the key enabler of strategic stability. Strategic stability was thus achieved by “making sure that each side has enough offensive forces to retaliate after a first strike and it assumes that neither side has the defensive capability to impede the other sides’ ability to deliver its devastating retaliatory strike.” [25]

Moving beyond the nuclear dimension, deterrence can be more generally defined as “the maintenance of such a posture that the opponent is not tempted to take any action which significantly impinges on his adversary’s vital interests”. [26] Deterrence therefore relies on maintaining an offense-defense balance, which is in favor of the defense.

It follows that “conflict and war will be more likely when offense has the advantage, while peace and cooperation are more probable when defense has the advantage.” [27] An important consideration when it comes to AWS is therefore the likely impact they will have on the offense-defense balance.

With the development of AWS, the threshold for the use of force will likely be lowered and thus favor the offense for two reasons. Firstly, even if the international community manages to impose limitations on targeting - which is debatable if we look at dynamics of contemporary conflict that make civilians, targets of choice - which could be more discriminate indeed, the key problem remains that states will likely exercise less restraint in the use of these weapons because the social cost incurred is lowered by the fact that no human life will be put at risk on the attacker’s side. One constraint in the use of AWS is economic. One could argue that the price of these technologies is a significant restraining factor on the propensity of states using and developing them. While at the early stages of development of any technology, the initial cost may be high, however, as has been demonstrated with the development of the computer or digital technologies, the high price of early adopters is greatly reduced over time. The same is likely to be true for AWS, particularly because of the dual-use nature of AI technology, which is primarily driven by the private sector.

Secondly, the offensive nature of these weapons is also strengthened by their likely tactical use, which relies on swarming tactics. The latter relies on overwhelming and saturating the adversary’s defense system by coordinating and synchronizing a series of simultaneous and concentrated attacks. Such tactics are aimed at negating the advantage of any defensive posture. In October 2016, the U.S. DoD conducted an experiment where 103 Perdrix micro drones were launched from FA/18 combat aircrafts and were assigned four objectives. In the words of William Roper, director of the Strategic Capabilities Office at the U.S. Department of Defence, the drones shared “one distributed brain for decision-making and adapting to each other like swarms in nature.” [28] The drones collectively decided that a mission was accomplished, flew on to the next mission and carried out that one. [29] On 9 February 2017, two teams from Georgia Tech Research Institute and the U.S. Naval Postgraduate School pitted two swarms of autonomous drones against one another. This represented “the first example of a live engagement between swarms of unmanned air vehicles (UAVs). [30] Three days later, on 12 February 2017, China set a

world record when “a formation of 1’000 drones performed at an air show in Guangzhou.” [31]

With the use of swarms of AWS it is very likely that the offense-defense balance will shift towards the former. In that case, deterrence will no longer be the most effective way to guarantee territorial integrity. In an international environment that favors the offensive, the best strategy to counter the offensive use of force is one that relies on striking first. It follows that strategies of pre-emption are very likely to become the norm if AWS are becoming the weapons of choice in the future. Striking first before being attacked will provide a strategic advantage. The concept of pre-emption, however, is a clear violation of the current international regime on the use of force, which relies on self-defense and authorization granted by the UN Security Council and falling under the Chapter VII of the UN charter.

Another consequence of favoring offense is the greater likelihood of international arms races.

As mentioned above, the ability to strike first represents a strategic advantage. In order to deny the adversary’s ability to do the same, states are very likely to invest and improve current AWS technology. This, in turn, is likely to initiate an arms race. We can already observe this dynamic at play. Last year, for instance, the Defense Science Board of the U.S. Department of Defense released its first study on autonomy. The report stated that “autonomous capabilities are increasingly ubiquitous and are readily available to allies and adversaries alike. The study therefore concluded that DoD must take immediate action to accelerate its exploitation of autonomy while also preparing to counter autonomy employed by adversaries. (...) Rapid transition of autonomy into warfighting capabilities is vital if the U.S. is to sustain military advantage.” [32]

AI and autonomy are seen pivotal technologies of the U.S. Third Offset strategy to secure military advantage over China and Russia. These two powers as well lesser one have also increasingly invested in autonomous weapons technologies. In August 2016, the state-run China daily reported that the country had embarked on “the development of cruise missile systems with a “high level” of artificial intelligence to counter the US Navy LRASM. [33] At the 2017 Russian Army Expo, Yermak, a fully automated, AI-powered anti-aircraft defence system was unveiled. [34] Russian famous AK-47 company, Kalashnikov, also recently announced that they are building “a range of products based on neural networks, including a fully automated combat module that can identify and shoot at its target.” [35] This is in line with President Putin recent remarks that “the one who becomes the leader in this sphere (i.e. AI) will be the ruler of the world.” [36]

When developing AWS, states and the international community should very carefully think about the consequences of this technology also falling into the hands of radical terrorist groups such as the Islamic State (ISIS). These groups massively rely on suicide bombings for tactical (breaching a front line, for instance) and strategic purposes (shocking the international community). The acquisition of AWS by these groups would act as a massive force multiplier as they could use the same suicide bombing tactics but with a greater concentration of mass (since they would rely on more machine acting as “suicide bombers”).

The vertical proliferation of autonomous weapons systems to non-state actors will potentially democratise the access to the use of military force. As mentioned in a study by the Harvard Kennedy School Belfer Center [37], the commercial AI industry concentrates most of AI advances. Beyond the issues of double-use technologies, commercial applications become increasingly widely available with time passing. In addition, the cost of replicating algorithms, which are lines of codes, is almost non-existent. Thus, scenarios where violent non-state actors such as ISIS, terrorist groups or criminal organisations could acquire AWS on black markets, or by building their own cannot be excluded. [38] Current developments in the weaponization of drones are an indicator supporting this possibility. [39]

Conclusions

Although autonomous weapons systems do not yet exist, the intended impact of these weapons on international security has the potential to be very destabilizing for the international system be it because it can upset the strategic balance and favor an offensive defense posture favoring pre-emptive strategies or because these technologies could be used beyond their intended limitations shall they fall into the hands of non-state actors or terrorist organizations. The international community – not only limited to the United Nations but also the AI industry players and scientific community – shall therefore be very vigilant about not granting too much power and autonomy to weaponised robots and algorithms as the consequences might well be dystopian.



Dr. Jean-Mark Rickli

He is the head of global risk and resilience at the Geneva Centre for Security Policy (GCSP). He is also a research fellow at King’s College London, where he was an assistant professor at the Department of Defence Studies, and a non-resident fellow in modern warfare and security at TRENDS Research and Advisory in Abu Dhabi. He is a senior advisor for the AI (Artificial Intelligence) Initiative at the Future Society at Harvard Kennedy School and an expert on autonomous weapons systems for the United Nations and for the United Nations Institute for Disarmament and Research (UNIDIR). Dr. Rickli received his PhD and MPhil in International Relations from Oxford University, UK, where he was also a Berrow scholar at Lincoln College.

Literature

- [1] McCarthy, John et al. (1955). “A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence” 31 August, <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>
- [2] Chainey, Ross (2017). “Google Co-Founder Sergey Brin: I Didn’t See AI Coming,” World Economic Forum, 19 January, <https://www.weforum.org/agenda/2017/01/google-sergey-brin-i-didn-t-see-ai-coming/>
- [3] Galeon, Dom (2016). “Nevermind Moore’s Law: Transistors Just Got a Whole Smaller,” Futurism, 8 October, <https://futurism.com/nevermind-moores-law-transistors-just-got-a-whole-lot-smaller/>
- [4] Khoso, Mikal (2016). “How Much Data is Produced Every Day,” Northeastern University, 13 May, <http://www.northeastern.edu/levelblog/2016/05/13/how-much-data-produced-every-day/>
- [5] EGHAM (2017), “Gartner says 8.4 Billion Connected “Things” Will Be in Use in 2017, Up 31 Percent From 2016”, Gartner, 7 February 2017: <http://www.gartner.com/newsroom/id/3598917>
- [6] 1 zettabyte is approximately 1’000 exabytes. 1 exabyte is 1’000 petabytes and 1 petabytes is 1’000’000 gigabytes. 1 gygabyte can hold the content of about 10 meters of book a shelf, see <http://whatsabyte.com>, See also, Mearian, Lucas (2012). “By 2020 there will be 5’200 GB of data for every person on Earth”, Computer World, 11 December, <http://www.computerworld.com/article/2493701/data-center/by-2020--there-will-be-5-200-gb-of-data-for-every-person-on-earth.html>
- [7] Machine learning is “the practice of using algorithms to parse data, learn from it and then make a determination or prediction about something in the world.” Copeland, Michael (2016). “What is the Difference Between Artificial Intelligence, Machine Learning and Deep Learning?” NVidia Blog, 29 July, <https://blogs.nvidia.com/blog/2016/07/29/>

whats-difference-artificial-intelligence-machine-learning-deep-learning-ai/

- [8] Hern, Alex (2015). "Computers Now Better than Human at Recognizing and Sorting Images," *The Guardian*, 13 May, <https://www.theguardian.com/global/2015/may/13/baiduminwa-supercomputer-better-than-humans-recognising-images>
- [9] Gershgorn, Dave (2017). "The Data That Transformed AI Research – and Possibly the World," *Quartz*, 26 July, <https://qz.com/1034972/the-data-that-changed-the-direction-of-ai-research-and-possibly-the-world/>
- [10] Lei, Leon (no date). "Go and Mathematics", *The American Go Foundation* <http://agfgo.org/downloads/Go%20and%20Mathematics.pdf>
- [11] Vincent, James (2016). "Google Uses Deepmind AI to Cut Data Center Energy Bills", *The Verge*, 21 July, <https://www.theverge.com/2016/7/21/12246258/google-deepmind-ai-data-center-cooling>
- [12] Metz, Cade (2017). "Google's Alphago Levels Up From Board Games to Power Grids", *Wired*, 24 May, <https://www.wired.com/2017/05/googles-alphago-levels-board-games-power-grids/>
- [13] Metz, Cade (2017). "Inside Libratus, the Poker AI that Out-Bluffed the Best Humans", *Wired*, 2 January, <https://www.wired.com/2017/02/libratus/>
- [14] Metz, Cade (2017). "Inside Libratus, the Poker AI that Out-Bluffed the Best Humans", *Wired*, 2 January, <https://www.wired.com/2017/02/libratus/>
- [15] Defense Science Board (2016). Summer Study on Autonomy. Washington, Department of Defense, Office of the Under Secretary of Defense for Acquisition, Technology and Logistics June, p. 4.
- [16] Defense Science Board (2016). Summer Study on Autonomy. Washington, Department of Defense, Office of the Under Secretary of Defense for Acquisition, Technology and Logistics June, p. 5.
- [17] Roff, Heather and Moyes, Richard (2016). "Autonomy, Robotics and Collective Systems", *Global Security Initiative*, Arizona State University, <https://globalsecurity.asu.edu/robotics-autonomy>
- [18] Roff, Heather and Moyes, Richard (2016). "Autonomy, Robotics and Collective Systems", *Global Security Initiative*, Arizona State University, <https://globalsecurity.asu.edu/robotics-autonomy>
- [19] LRASM is a stealthy long range anti-ship cruise missile developed in the United States while TARES (Tactical Advanced Recce Strike) is an unmanned combat air vehicle (UCAV) for standoff engagement developed by the German firm Rheinmetall.
- [20] The nEUROn is an experimental stealthy, autonomous UCAV developed by an international cooperation led by French Dassault. Taranis is a British demonstrator programme for stealthy UCAV technology developed by BAE which first flew in 2013.
- [21] Roff, Heather and Moyes, Richard (2016). "Autonomy, Robotics and Collective Systems", *Global Security Initiative*, Arizona State University, <https://globalsecurity.asu.edu/robotics-autonomy>
- [22] Roff, Heather and Moyes, Richard (2016). "Autonomy, Robotics and Collective Systems", *Global Security Initiative*, Arizona State University, <https://globalsecurity.asu.edu/robotics-autonomy>
- [23] Schwab, Klaus (2017). *Shaping the Fourth Industrial Revolution: Handbook for Citizens, Decision-Makers, Business Leaders and Social Influencers*, Geneva, World Economic Forum, p. 15 (forthcoming).
- [24] Hildreth, Steven and Woolf, Amy. F (2010). *Ballistic Missile Defence and Offensive Arms Reductions : a Review of the Historical Record*. Washington : Congressional Research Service, p. 4.
- [25] Hildreth, Steven and Woolf, Amy. F (2010). *Ballistic Missile Defence and Offensive Arms Reductions: a Review of the Historical Record*. Washington: Congressional Research Service, p. 7. When security is achieved through deterrence by punishment "the ability to retaliate supports a defensive strategy, while the ability to deny retaliatory capabilities, that is, to limit damage, could support and offensive strategy". MAD therefore is defense-dominant as retaliatory capabilities dominate damage limitation capabilities. Glaser, Charles L. (2014). *Analyzing Strategic Nuclear Policy*. Princeton: Princeton University Press, pp. 106-107, see also, Jervis, Robert (1978). "Cooperation under the Security Dilemma," *World Politics*, vol. 30, no. 2, January, pp. 206-210.
- [26] Amoretta M Hoeber (1968): "Strategic Stability," *Air University Review*, July-August, <http://www.airpower.maxwell.af.mil/airchronicles/aureview/1968/jul-aug/hoerber.html>
- [27] Sean M. Lynn-Jones (1995). "Offense-Defense Theory and Its Critics", *Security Studies*, vol. 4, no. 4, p. 691.
- [28] U.S. DoD (2017). "Department of Defense Announces Successful Micro-Drone Demonstration" U.S. Department of Defense, 9 January, <https://www.defense.gov/News/News-Releases/News-Release-View/Article/1044811/department-of-defense-announces-successful-micro-drone-demonstration/>
- [29] Mizokami, Kyle (2017). "The Pentagon's Autonomous Swarming Drones are the Most Unsettling Thing You'll See Today." *Popular Mechanics*, 9 January, <http://www.popularmechanics.com/military/aviation/a24675/pentagon-autonomous-swarming-drones/>
- [30] Toon, John (2017). "Swarms of Autonomous Aerial Vehicles test Dogfighting Skills" *Georgia Tech News Centre*, 21 April, <http://www.news.gatech.edu/2017/04/21/swarms-autonomous-aerial-vehicles-test-new-dogfighting-skills>
- [31] Le Miere, Jason (2017). "Russia Developing Autonomous "Swarm" of Drones in New Arms Race with U.S., China." *Newsweek*, 15 May, <http://www.newsweek.com/drones-swarm-autonomous-russia-robots-609399>
- [32] Defense Science Board (2016). Summer Study on Autonomy. Washington, Department of Defense, Office of the Under Secretary of Defense for Acquisition, Technology and Logistics June, pp. iii and 3.
- [33] Markoff, John and Rosenberg, Matthew (2017). "China's Intelligent Weaponry Gets Smarter," *New York Times*, 3 February, <https://www.nytimes.com/2017/02/03/technology/artificial-intelligence-china-united-states.html?mcubz=0>
- [34] Jotham, Immanuel (2017). "Russian Army Expo 2017: Autonomous Anti-Aircraft Missile System Unveiled," *International Business Times*, 28 August, <http://www.ibtimes.co.uk/russian-army-expo-2017-weapons-maker-tikhomirov-unveils-new-autonomous-anti-aircraft-missile-system-1636920>
- [35] Tucker, Patrick (2017). "Russian Weapons Maker to Build AI-Directed Guns," *DefenceOne*, 14 July, <http://www.defenseone.com/technology/2017/07/russian-weapons-maker-build-ai-guns/139452/>
- [36] Associated Press (2017), "Putin: Leader in artificial intelligence will rule world", *The Washington Post*, 1 September: https://www.washingtonpost.com/business/technology/putin-leader-in-artificial-intelligence-will-rule-world/2017/09/01/969b64ce-8f1d-11e7-9c53-6a169beb0953_story.html?utm_term=.1aed353f12ac
- [37] Allen, Greg and Chan Taniel, (2017). *Artificial Intelligence and National Security*, Harvard Kennedy School Belfer Center Study for Science and International Affairs, July: <http://www.belfercenter.org/sites/default/files/files/publication/AI%20NatSec%20-%20final.pdf>
- [38] Horowitz C. Michael (2016). "Who'll want artificially intelligent weapons? ISIS, democracies, or autocracies?" *Bulletin of the Atomic Scientists*, 29 July: <http://thebulletin.org/who%E2%80%99ll-want-artificially-intelligent-weapons-isis-democracies-or-autocracies9692>
- [39] Warrick Joby (2017) "Use of weaponized drones by ISIS spurs terrorism fears" *The Washington Post*, 21 February: https://www.washingtonpost.com/world/national-security/use-of-weaponized-drones-by-isis-spurs-terrorism-fears/2017/02/21/9d83d51e-f382-11e6-8d72-263470bf0401_story.html?utm_term=.a64b5f10e771