

Artificial Intelligence as a Dual-use Technology

Éva AMBRUS¹

The aim of this article is to give an overview of the state of artificial intelligence regarding malware attacks, its uses in the military and views regarding if it should be classified as a dual-use technology. As an emerging technology, with a wide variety of use and capabilities, more could be done to overview its uses, and some form of control over it. While the classical exports control might be counterproductive, a more closed approach towards critical information dissemination might be advisable until the full range of capabilities of artificial intelligence will be known.

Keywords: *artificial intelligence, dual-use technology, military use, malware*

Introduction

The security paradigm is changing. Until a new definition comes forward, policy-makers, academia and users will debate its nature and possible effects. Asymmetrical warfare, hybrid warfare, ‘grey area’ warfare, (dis)information warfare, unpeace are just a few names used trying to pinpoint the development of (IT) technology on security. Warfare and security includes more and more cyberspace, including cyber weapons, cyber espionage and cybersecurity. One driver of this change is the advances made in the last decade regarding artificial intelligence (AI). In this article I will present the idea that AI should be classified as a dual-use technology, meaning that it can be used for both civilian and military applications. I will start with presenting where AI weapons are today, followed by the nature of the relationship between state and technology. I will then present a case for thinking about AI as a dual-use technology.

AI as a weapon

Writing an article about artificial intelligence and its uses can leave one with more questions than answers. And as the ‘grey area’ warfare, or this era of ‘unpeace’, even questions have a high complexity. With these in mind, my aim with this article is to shed light to some of the questions asked today regarding the malign use of artificial intelligence. There is little

¹ PhD student, University of Public Service, Faculty of Military Science and Officer Training, Doctoral Schools of Military Sciences and Military Engineering, e-mail: ambrus.eva.eszter@gmail.com; ORCID: <https://orcid.org/0000-0002-8354-1296>

question that AI will make cyber warfare more powerful, increasing its scale, speed and power. Merriam-Webster defines artificial intelligence as: (1) a branch of computer science dealing with the simulation of intelligent behaviour in computers; and (2) the capability of a machine to imitate intelligent human behaviour.² Although this encompasses the basic notion, it is still not a complete definition. This lack of certainty also comes from the notion of intelligence. What is intelligence? In case of AI, science has focused on different aspects, such as learning, reasoning, problem solving, perception, language and many others. Another approach is to focus on the goals and aims of AI. This is the notion put forward by Russel and Norvig, as well. There is no agreed definition of artificial intelligence. Russel and Norvig summarised the four main schools of thoughts as AI in the following way: (1) thinking humanly; (2) thinking rationally; (3) acting humanly; and (4) acting rationally.³ For this article, I will use the 4th approach, defining AI as a system that acts rationally, thus AI can be called a rational agent. ‘A rational agent is one that does the right thing’,⁴ the ‘right thing’ being the most successful outcome for the agent (in this case, AI).

A branch of AI is machine learning. Machine learning refers to the ability of a computer to learn using large sets of data (not just predefined rule sets).⁵ Machine learning can basically be supervised, reinforced or unsupervised. In supervised learning, the machine is trained to perform a specific task, such as recognising cats in pictures. For it to learn to distinguish this, it needs large amounts of tagged data, and this also includes checking the correct answers. Supervised training is used for tasks requiring information classification (for example filtering spam messages). Reinforced learning is giving direct feedback to the autonomous system about its output (for example did it classify correctly). In case of unsupervised learning, the program is not assigned any task and the data is unlabelled, so it is free to find its own correlations in the data. Learning from the data, the machine creates clusters in the given data and sets association rules that combine the various variables in the data. In cybersecurity, this can be the detection of malware.

An interesting part regarding machine learning and adaptability in the concepts of the human-in-the-loop (HITL). Basically, it combines human and artificial intelligence to create machine learning models with humans directly involved in training, tuning and testing the data. It is understood that HITL is important in cases when the cost of error is too high, when the ML algorithm cannot have any margin of error.⁶ This would be the case in any military application, but also for autonomous driving. Literature distinguishes between human-in-the-loop, human-on-the-loop and human-off-the-loop. In the first case (HITL), the human has the final say in the execution of the lethal force (for example drones). In the second case, the decision can be made without the human operator, but the operator can override it. In the third scenario, the human operator cannot override the weapon system’s triggering mechanism, so there is no human intervention possible.⁷

² ‘Artificial intelligence’, Merriam-Webster dictionary.

³ Stuart J Russell and Peter Norvig, *Artificial Intelligence. A Modern Approach* (New Jersey: Prentice Hall, 2010), 2.

⁴ Russell and Norvig, *Artificial Intelligence*, 4.

⁵ ‘Machine learning’, [Dictionary.com](https://www.dictionary.com).

⁶ Mothi Venkatesh, ‘What is Human-in-the-Loop for Machine Learning?’, [Hackernoon.com](https://hackernoon.com), July 17, 2018.

⁷ Seumas Miller, *Dual Use Science and Technology, Ethics and Weapons of Mass Destruction* (Springer International Publishing, 2018), 100.

Reviewing the literature, it seems generally believed that an antivirus system with artificial intelligence and machine learning is the solution to modern malware attacks.⁸ Malware can be defined as unwanted software which performs non-benign operations on any system. Avi Pfeifer et al. gave two insights into malwares. The first is that like biological viruses, they are rarely de novo; malwares are re-used to avoid detection and hide similarities. Secondly, regarding the functioning of malwares, it is harder to obfuscate what it does try to accomplish than how it wants to do it – it leaves a trace. It can be seen from these analogies, that malwares are changing fast, and the adaptability of AI and ML needs to be high.⁹

Adaptability faces at least two challenges: costs and adversarial attack. As we have seen, for ML to work well, it needs data (either labelled or unlabelled) to learn from. If the task is that it is trained for changes day-by-day, it means that it needs continuous training to be able to differentiate between benign or non-benign software. The second challenge, adversarial attack refers to the tactic of ‘poisoning the well’, or in this case, the training data sets. Recent research shows that deep learning is sensitive to contrasting, contradictory examples where the opponent can manipulate the input of the deep learning model in such a discreet way by adding minimal disruption to the input material to produce the desired result, that is, misclassification.

Its structure is due to the system of so-called neural networks. Neural networks are made up of elementary computing units – so-called neurons, which form interconnected layers. Each neuron applies an activation function to its input to create a specific output. Starting with model input, each network layer produces an output that the next layer uses as input. Networks with a single intermediate layer – hidden – are considered shallow neural networks, while models with multiple hidden layers are deep neural networks. They are sensitive to the manipulation of opposite examples of their inputs. ‘Adversarial examples are inputs to a classifier specifically crafted to deceive the model, causing misclassification’.¹⁰ Training models developed based on real and simulated data may be significantly more secure, but its development can come with a higher cost. One solution to this is counter-narrator training,¹¹ which is to improve the model’s generalisation ability, that is, the prediction of patterns outside the learner’s data set. Good generalisation also generally makes the classification less sensitive to minor disturbances and therefore, more resistant to conflicting examples.

Cybersecurity is often thought as passive, meaning the systems are waiting for the attack and all AI can do is help detect, categorise and respond to the attack.¹² AI-enhanced cyber weapons will have wider scope and greater speed than today’s adversarial AI. One of the threats of the future would be that these AI-enabled tools would enable, for example

⁸ Sherif Saad, William Briguglio and Haytham Elmiligi, ‘The Curious Case of Machine Learning in Malware Detection’, [Arxiv.org](https://arxiv.org/abs/2019.05.18), May 18, 2019.

⁹ Avi Pfeiffer, Brian E Ruttenberg, Lee Kellogg, Michael Howard, Catherine Call, Alison M O’Connor, Glenn Takata, Scott Neal Reilly, Terry Patten, Jason Taylor, Robert Hall, Arun Lakhotia, Craig Miles, Daniel Scofield and Jared Frank, ‘Artificial Intelligence Based Malware Analysis’, [Arxiv.org](https://arxiv.org/abs/2017.04.27), April 27, 2017.

¹⁰ Nuno Martins, José Magalhães Cruz, Tiago Cruz and Pedro Henriques Abreu, ‘Adversarial Machine Learning Applied to Intrusion and Malware Scenarios: A Systematic Review’, *IEEE Access* 8 (2020), 35417.

¹¹ Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens and Zbigniew Wojna, ‘Rethinking the Inception Architecture for Computer Vision’, *CVPR* (2016), 2818–2826.

¹² Anna L Buczak and Erhan Guven, ‘A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection’, *IEEE Communications Surveys Tutorials* 18, no 2 (2016), 1153–1176.

‘data-poisoning’. Data poisoning attacks are attacks when ‘malicious users inject false training data with the aim of corrupting the learned model’.¹³ As we have seen machine learning depends on datasets, and tampering with the input data to divert its results would mean it is at best underperforming, undetectably. The time for the human-in-the-loop to detect such functioning anomaly is essential, especially as more and more systems can (and will) become autonomous. Adaptability on the defence side faces more constraint than on the attackers’ side. Some of these factors are financial (costs), personnel (retraining), performance (should not degrade system performance), usability and manageability, operations (fitting into the security operation), design (built-in preferred), perception (of the usefulness of defensive technology).¹⁴

Dual-use technology

There are valid concerns about the increase of autonomy of weapons system, and ethical questions are raised. Non-governmental organisations like International Committee for Robot Arms Control would limit the research and development of AI to civilian use only. On the other hand, these technologies are already being developed and tested. Large scale deployment will not happen until the margin of error of these systems will be close to zero, thus the importance of the mentioned human-in-the-loop in the process. But it is imaginable that in the future these autonomous systems will be more accurate than humans, thus new ethical questions will be raised.

Artificial intelligence and its uses in attacks or defences are emerging technologies. Emerging technologies can be described as ‘technologies that have disruptive potential but have not yet been developed to their fullest potential’.¹⁵ AI throughout its history had several ‘growth’ periods, when technological advances made it possible to develop it further. Its full scope of practical uses still cannot be determined and it is used in both military and civilian industry. In theory, there is the possibility of misuse by different actors. One can argue, that most technology can be used for more than one purpose, but the term ‘dual-use’ is reserved for technology that has a significant government application (and thus pertains to national security) and a private sector application, as well. At one point PlayStation 2 was briefly considered a dual-use technology by Japan.¹⁶

Gregory Lewis et al. in their article present the case of information hazard in biotechnology. Their view is that both openness and secrecy of information may backfire. They suggest ‘that mitigation of these hazards can be improved if one can: (1) anticipate hazard potential before scientific work is performed; (2) consider how much the new

¹³ Jacob Steinhardt, Pang Wei Koh and Percy Liang, ‘Certified Defenses for Data Poisoning Attacks’, [Arxiv.org](#), November 24, 2017.

¹⁴ Sean M Price, ‘Adaptive threats and defences’, in *Information Security Management Handbook*, vol. 4, ed. by Harold F Tipton and Micki Krause (Auerbach Publications, 2019), 44–45.

¹⁵ Daniele Rotolo, Diana Hicks and Ben R Martin, ‘What Is an Emerging Technology?’, [Arxiv.org](#), January 4, 2016, 4.

¹⁶ Associated Press, ‘Sony’s High-Tech Playstation2 Will Require Military Export License’, *Los Angeles Times*, April 17, 2000.

information would likely help both good and bad actors; and (3) aim to disclose information in the manner that maximally disadvantages bad actors versus good ones'.¹⁷

As per the EU's definition 'dual-use items are goods, software and technology that can be used for both civilian and military applications'.¹⁸ We have seen the definition in the EU, but for its research and development programme, 'Horizon 2020, is more specific, requiring applicants for funding to ensure that "research and innovation activities carried out under Horizon 2020 shall have an exclusive focus on civil applications," and they are required to complete an ethics checklist to demonstrate that they comply with this requirement'.¹⁹

AI systems and their design knowledge can be used for both civilian and military applications, and more broadly for beneficial and harmful purposes. Artificial intelligence is dual-use in the same sense as human intelligence. Many of the tasks that would be useful to automate are themselves dual-use. For example, software vulnerability detection systems have both offensive and defensive applications, and there is little difference between the capabilities of an autonomous drone used to transport packages and an autonomous drone used to transport explosives. In addition, basic research aimed at understanding AI, enhancing its abilities and controlling it is inherently dual-use. Machine learning (and AI) is a fairly open field, where researchers share details about their models and codes on the internet, as well. Hagendorff introduces the notion of 'forbidden knowledge', which is akin to Bostrom's 'information hazard'. While the latter is 'a risk that arises from the dissemination or the potential dissemination of (true) information that may cause harm or enable some agents to cause harm',²⁰ while the former is defined as (scientific) knowledge that is too dangerous to be disseminated unrestrictedly, for example in the fields of IT security or synthetic biology.²¹

One of the trends is the increase of existing threats. The cost of attacks is decreasing with the spread of AI to perform tasks that generally require human work, intelligence and expertise. As a result, the range of actors capable of carrying out certain attacks, the speed at which the attacks are executed and the set of potential targets will increase. An emerging threat of using AI systems to perform tasks that are virtually unmanageable by humans is the fact that attackers can exploit vulnerabilities in AI systems. The typical nature of threats is changing. With the increasing use of AI, attacks will become more effective, highly targeted, difficult to associate with a perpetrator, and are likely to exploit vulnerabilities in AI systems. Cybersecurity is an area that takes early and enthusiastic advantage of AI. The adaptability of AI systems can also change the strategic environment of cybersecurity, the attack/defence balance. The systems currently in use are quite effective against typical human-made malware, and research has already shown that AI systems will soon be able to circumvent their protection.

¹⁷ Gregory Lewis, Piers Millett, Anders Sandberg, Andrew Snyder Beattie and Gigi Gronvall, 'Information Hazards in Biotechnology', *Risk Analysis* 39, no 5 (2019), 1.

¹⁸ European Commission, 'Dual-use trade controls'.

¹⁹ Tara Mahfoud, Christine Aicardi, Saheli Datta and Nikolas Rose, 'The Limits of Dual Use', *Issues in Science and Technology* 34, no 4 (2018).

²⁰ Nick Bostrom, 'Information Hazards: A Typology of Potential Harms from Knowledge', *Review of Contemporary Philosophy* 10 (2011), 45.

²¹ Thilo Hagendorff, 'Forbidden knowledge in machine learning. Reflections on the limits of research and publication', *Arxiv.org*, November 2019, 3.

Trends

Attitudes of society might change regarding AI as the coming generations are more and more confident in it. As with other technologies before, it may not be possible to clearly separate the civilian and military uses of it. Like the dual-faced God Janus, many of our everyday technology can be used for both purposes. We have to accept the fact that this technology will be used by actors for harm and prepare for it.

Carrick Flynn wrote a brief issue on the export control of artificial intelligence. His findings (regarding the *state of play* in the United States) are summarised in four main points:

- (1) New export control regulations on general purpose AI software, untrained algorithms, and datasets without military use are unlikely to succeed and should not be implemented.
- (2) Highly application-specific AI software, trained algorithms, and militarily sensitive data sets are useful targets for export control, but are already covered by the current export control regime.
- (3) Equipment for manufacturing AI chips is likely a highly effective point of export control.
- (4) The effectiveness of export controls on AI chips will depend on early implementation of export controls on chip manufacturing equipment. AI chips themselves are not yet a promising target for expanded regulation.

Per his findings, the (3) option could be a next step in export control regulation. In his view: ‘The computing power required for AI increasingly relies on specialized microprocessors (AI chips) optimized for AI applications. AI chips are produced using highly advanced semiconductor manufacturing equipment that is relatively easy to define, monitor, track, and control.’²²

New export control (option 1) would go against the fact that at this point innovation in AI relies on openness in the field, thus could harm research, as well as damaging relations between governments and the industry.

AI is becoming an important factor in maintaining the economic and national security of most countries, but as new technologies develop, so should new tools to address them. In the 20th century, one way of maintaining technological superiority was through export controls. Export controls are a web of regulations that prohibit the transfer of certain commodities or information, motivated by national security concerns or trade objectives, or both. Effectiveness of such measures in a globalised value-chain world remains dubious, as it needs to balance research and development needs (which thrive in an open environment) as well the interests of multinational technological corporations, allied countries and scientists.²³

Norms prevailing in the AI research community show a strong tendency towards openness. Most new research is published online, often sharing all the information from

²² Carrick Flynn, ‘Recommendations on Export Controls for Artificial Intelligence’, *Centre for Security and Emerging Technology*, February 2020.

²³ Jade Leung, Sophie-Charlotte Fischer and Allan Dafoe, ‘Export controls in the age of AI’, *War on the Rocks*, August 28, 2019.

the outline, through the algorithmic details to the source code. This level of openness has clear benefits in enabling researchers to rely on each other's work, fostering collaboration. Rather, it is a thought-provoking idea of what solutions may be needed when it comes to moving away from openness for security reasons. Should a risk assessment be conducted before publishing in detail AI attacks that can be used for attack? This is the norm, for example, in the field of biotechnology. Or would it be too early for this measure to await its widespread adoption, assessing which technical research is most important for safety? Should a community be established in which certain types of research results can be selectively shared among a predetermined set of criteria that meet certain criteria, such as effective information security and appropriate ethical standards? What can be learned from other models of dual-use technology sharing?

AI in the military

Regarding the uses of artificial intelligence in the military, one of the areas concerned is about decision-making.²⁴ The use of AI in analysis, classification is already in use (for example in the automotive industry) and will help decision-makers 'by providing easy-to-understand analysis and recommendations based on big data'²⁵. The question of control in the military is especially important, although most articles underline that humans will remain the final decision-makers. Scott D Sagan explores the connection between ethics, technology and war. In his view, new technologies could reduce collateral damage, but also lower the political cost of engagement, thus more conflicts would emerge.²⁶ This is a trend seen in the 'grey area' conflicts, like cyber weapons. Artificial intelligence is changing the nature and principles of warfare by making decision-making cycles faster, advantages provided by AI will be for those that can apply AI in the broadest sense, and thus a revision of concepts regarding the organisation, control and command of military forces will be needed.²⁷

Another distinction that can be made is that technology shapes warfare, the conduct of war. This distinction is important because more and more violence nowadays is not confined to a precise geographic area, and the opponents are also not clearly defined. Grasping a new definition of conflict, or extending its definition poses a problem. And with the rise of new technologies, an erosion of the state's monopoly over the use of force, combined with the proliferation of new technologies to non-state actors poses a new threat. On the other hand, wars are costly enterprises, even this new era of unpeace. As Sterling Pavelec notes, 'modern military technology is costly, funded by government resources and will require massive amounts of funding, brainpower, and a society that is willing and capable of technological evolution'.²⁸

²⁴ Gordon Cooke, 'Magic Bullets: The Future of Artificial Intelligence in Weapons Systems', [Army.mil](https://www.army.mil), June 11, 2019.

²⁵ Cooke, 'Magic Bullets'.

²⁶ Scott D Sagan, 'Ethics, Technology and War', *Daedalus* 145, no 4 (2016), 6–11.

²⁷ Imre Porkoláb and Imre Négyesi, 'A mesterséges intelligencia alkalmazási lehetőségeinek kutatása a haderőben', *Honvédségi Szemle* 147, no 5 (2019), 17.

²⁸ Sterling M Pavelec, *War and Warfare since 1945* (Routledge, 2017), 156.

‘War made the state and state made the war’ is a famous saying, and technology can help or hinder both. After the Second World War, theorists talk about postmodern warfare as it became less physical; states could reach their aims below the threshold of war. Technological arms race persists today: we face uncertainty, complexity and over-reliance on technology. Alex Roland argues in his essay, that ‘technology is like an “open door”, as it adds what most accounts of technological innovation lack: human agency. Humans must decide if they are going to, or can, take up a given military innovation. And they must adapt it to their circumstances. Technology is a possibility, not an imperative’.²⁹ As the emerging and developing technologies raise new questions about ethics, morals and legality, it is important to notice that although the door is opening, we have not yet passed through it. As Feldman et al. wrote, ‘context is critical: training exercises may look like war, but they are actually between allies; cold war may look like peace, but it isn’t exactly. Any intelligent system (human, human/machine, or machine) must be aware of these and other complicating concerns’.³⁰

Regarding the use of AI and ML in the military, in a recently published article James Johnson argues that ‘the fusion of AI machine learning and human judgment to gauge an adversary’s intentions (and predict escalation) for the purposes of planning and directing future wars for the pursuit of political objectives, is, therefore, a far less unlikely prospect in the near future than the use of AI to achieve tactical and operational ends (e.g. drone swarming and cyber defence)’.³¹ The connection between war and technology is undeniable, but their logic differ on a fundamental level – ‘technology perceives the universe as functioning rationally and predictably, while in war no success is possible which is not grounded in an ability to tolerate uncertainty, cope with it and make use of’.³²

Conclusion

In this article I have attempted to give an overview of the main points regarding artificial intelligences’ adversarial use, its place in the military and the questions regarding dual-use technology distinction. The development of new technologies raises the question of its uses, and in the case of AI, more so are questions of ethical nature. As yet it is an emerging technology, its full capabilities are hard to predict, but it is my view that erring on the side of caution would be preferable. Understanding that a technological race might unfold, a limited dissemination of information could be one of the possible solutions until we are aware of all possibilities that artificial intelligence may provide. Finally, at one point it would be advisable to have a distinction between military-level and civilian-level AI, either through capabilities, aims or by other criteria.

²⁹ Alex Roland, ‘War and Technology’, *Foreign Policy Research Institute*, February 27, 2009.

³⁰ Philip Feldman, Aaron Dant and Aaron Massey, ‘Integrating Artificial Intelligence into Weapon Systems’, [Arxiv.org](https://arxiv.org/abs/1905.08111), May 10, 2019.

³¹ James Johnson, ‘The AI-cyber nexus: implications for military escalation, deterrence and strategic stability’, *Journal of Cyber Policy* 4, no 3 (2019).

³² Martin van Creveld, *Technology and War: From 2000 B.C. to the Present* (New York: The Free Press, 1991), 316.

References

- ‘Artificial intelligence’, Merriam-Webster dictionary. Online: www.merriam-webster.com/dictionary/artificial%20intelligence
- Associated Press, ‘Sony’s High-Tech Playstation2 Will Require Military Export License’. *Los Angeles Times*, 17 April 2000. Online: www.latimes.com/archives/la-xpm-2000-apr-17-fi-20482-story.html
- Bostrom, Nick, ‘Information Hazards: A Typology of Potential Harms from Knowledge’. *Review of Contemporary Philosophy* 10 (2011), 44–79.
- Buczak, Anna L and Erhan Guven, ‘A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection’. *IEEE Communications Surveys Tutorials* 18, no 2 (2016), 1153–1176. DOI: <https://doi.org/10.1109/comst.2015.2494502>
- Cooke, Gordon, ‘Magic Bullets: The Future of Artificial Intelligence in Weapons Systems’. *Army.mil*, 11 June 2019. Online: www.army.mil/article/223026/magic_bullets_the_future_of_artificial_intelligence_in_weapons_systems
- European Commission, ‘Dual-use trade controls’. Online: <https://ec.europa.eu/trade/import-and-export-rules/export-from-eu/dual-use-controls/>
- Feldman, Philip, Aaron Dant and Aaron Massey, ‘Integrating Artificial Intelligence into Weapon Systems’, *Arxiv.org*, 10 May 2019. Online: <https://arxiv.org/pdf/1905.03899.pdf>
- Flynn, Carrick, ‘Recommendations on Export Controls for Artificial Intelligence’. *Centre for Security and Emerging Technology*, February 2020. Online: <https://cset.georgetown.edu/wp-content/uploads/Recommendations-on-Export-Controls-for-Artificial-Intelligence.pdf> DOI: <https://doi.org/10.51593/20190001>
- Hagendorff, Thilo, ‘Forbidden knowledge in machine learning. Reflections on the limits of research and publication’. *Arxiv.org*, November 2019. Online: <https://arxiv.org/pdf/1911.08603.pdf> DOI: <https://doi.org/10.1007/s00146-020-01045-4>
- Johnson, James, ‘The AI-cyber nexus: implications for military escalation, deterrence and strategic stability’. *Journal of Cyber Policy* 4, no 3 (2019), 442–460. DOI: <https://doi.org/10.1080/23738871.2019.1701693>
- Leung, Jade, Sophie-Charlotte Fischer and Allan Dafoe, ‘Export controls in the age of AI’. *War on the Rocks*, 28 August 2019. Online: <https://warontherocks.com/2019/08/export-controls-in-the-age-of-ai/>
- Lewis, Gregory, Piers Millett, Anders Sandberg, Andrew Snyder Beattie and Gigi Gronvall, ‘Information Hazards in Biotechnology’. *Risk Analysis* 39, no 5 (2019), 975–981. DOI: <https://doi.org/10.1111/risa.13235>
- ‘Machine learning’, *Dictionary.com*. Online: www.dictionary.com/browse/machine-learning
- Mahfoud, Tara, Christine Aicardi, Saheli Datta and Nikolas Rose, ‘The Limits of Dual Use’. *Issues in Science and Technology* 34, no 4 (2018). Online: <https://issues.org/the-limits-of-dual-use/>
- Martins, Nuno, José Magalhães Cruz, Tiago Cruz and Pedro Henriques Abreu, ‘Adversarial Machine Learning Applied to Intrusion and Malware Scenarios: A Systematic Review’. *IEEE Access* 8 (2020), 35403–35419. DOI: <https://doi.org/10.1109/ACCESS.2020.2974752>
- Miller, Seumas, *Dual Use Science and Technology, Ethics and Weapons of Mass Destruction*. Springer International Publishing, 2018. DOI: <https://doi.org/10.1007/978-3-319-92606-3>

- Pavelec, Sterling M, *War and Warfare since 1945*. Routledge, 2017. DOI: <https://doi.org/10.4324/9781315175478>
- Pfeffer, Avi, Brian E Ruttenberg, Lee Kellogg, Michael Howard, Catherine Call, Alison M O'Connor, Glenn Takata, Scott Neal Reilly, Terry Patten, Jason Taylor, Robert Hall, Arun Lakhotia, Craig Miles, Dan Scofield and Jared Frank, 'Artificial Intelligence Based Malware Analysis'. *Arxiv.org*, 27 April 2017. Online: <https://arxiv.org/pdf/1704.08716.pdf>
- Porkoláb Imre and Négyesi Imre, 'A mesterséges intelligencia alkalmazási lehetőségeinek kutatása a haderőben'. *Honvédségi Szemle* 147, no 5 (2019), 3–20.
- Price, Sean M, 'Adaptive threats and defences', in *Information Security Management Handbook*, vol. 4, ed. by Harold F Tipton and Micki Krause. Auerbach Publications, 2019, 42–65.
- Roland, Alex, 'War and Technology'. *Foreign Policy Research Institute*, 27 February 2009. www.fpri.org/article/2009/02/war-and-technology/
- Rotolo, Daniele, Diana Hicks and Ben R Martin, 'What Is an Emerging Technology?' *Arxiv.org*, 4 January 2016. Online: <https://arxiv.org/abs/1503.00673>
- Russell, Stuart J and Peter Norvig, *Artificial Intelligence. A Modern Approach*. New Jersey: Prentice Hall, 2010.
- Saad, Sherif, William Briguglio and Haytham Elmiligi, 'The Curious Case of Machine Learning in Malware Detection'. *Arxiv.org*, 18 May 2019. Online: <https://arxiv.org/pdf/1905.07573.pdf> DOI: <https://doi.org/10.5220/0007470705280535>
- Sagan, Scott D, 'Ethics, Technology and War'. *Daedalus* 145, no 4 (2016), 6–11. DOI: https://doi.org/10.1162/daed_e_00407
- Steinhardt, Jacob, Pang Wei Koh and Percy Liang, 'Certified Defenses for Data Poisoning Attacks'. *Arxiv.org*, 24 November 2017. Online: <https://arxiv.org/abs/1706.03691>
- Szegedy, Christian, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens and Zbigniew Wojna, 'Rethinking the Inception Architecture for Computer Vision', *CVPR* (2016), 2818–2826. DOI: <https://doi.org/10.1109/cvpr.2016.308>
- Van Creveld, Martin, *Technology and War: From 2000 B.C. to the Present*. New York: The Free Press, 1991.
- Venkatesh, Mothi, 'What is Human-in-the-Loop for Machine Learning?' *Hackernoon.com*, 17 July 2018. Online: <https://hackernoon.com/what-is-human-in-the-loop-for-machine-learning-2c2152b6dfbb>