JOURNAL OF INTERNET LAW

VOLUME 22

NUMBER 5

NOVEMBER 2018

EDITED BY DLA PIPER

Artificial Intelligence and Export Controls: Conceivable, But Counterproductive?

By Roszel C. Thomsen II

INTRODUCTION

Artificial intelligence (AI)[1] increasingly is viewed as an important technology for protecting the national and economic security interests of the United States. The Trump Administration's National Security Strategy summarized it succinctly:

> *To maintain our competitive advantage, the United States will prioritize emerging technologies critical to economic growth and security, such as data science, encryption, autonomous technologies, gene editing, new materials, nanotechnology, advanced computing technologies, and artificial intelligence. From self-driving cars to autonomous weapons, the field of artificial intelligence, in particular, is progressing rapidly.*[2]

Where security interests lead, export controls frequently follow. Ambassador Philip Griffiths, Head of Secretariat Wassenaar Arrangement (WA) on Export Controls for Conventional Arms and Dual-Use Goods and Technologies,[3] likewise identified AI as a priority:

> *Looking ahead, the WA Lists review process can be expected to continue to address new*

Continued on page 15

Roszel C. Thomsen II is a senior partner in the law firm of Thomsen and Burke LLP and founder of the Alliance for Network Security, a trade association composed of leading Information Technology companies. His practice focuses on international trade and investment issues, including export controls, economic sanctions, and foreign investment reviews.

Wolters Kluwer

Artificial Intelligence and Export
from page 1

> *technologies of security concern, including additive manufacturing or 3-D printing (where the WA is pursuing technical dialogues with the NSG and the MTCR respectively), thermal batteries, terrestrial equipment and components for satellites, as well as artificial intelligence and the integration of advanced sensors and navigation equipment to increase autonomy of weapons systems and robotisation of the battlefield.*[4]

In addition to export controls under the WA, the President also could implement export controls on AI as an "emerging and foundational technology." The new National Defense Authorization Act of 2019 (Pub. L. No. 115-232),[5] Title XVII – Review of Foreign Investment and Export Controls, Subtitle B – Export Control Reform Section 1758, authorizes the President to identify and control the export of "emerging and foundational technologies" (as yet to be defined).[6] [7]

At its annual BIS Update conference on May 14–15, 2018, the U.S. Department of Commerce's Bureau of Industry and Security (BIS) announced that it intends to publish a proposed rule that will define how specific "emerging and foundational technologies" will be identified and evaluated under the Export Administration Regulations. Most recently, on May 29, 2018, the White House promised new export controls, saying: "To protect our national security, the United States will implement … enhanced export controls for Chinese persons and entities related to the acquisition of industrially significant technology. The … export controls will be announced by June 30, 2018, and they will be implemented shortly thereafter."[8]

To date, no such export controls have been implemented.[9] However, on August 8, 2018, the BIS announced a recruitment for new candidates to serve on the Emerging Technology Technical Advisory Committee to advise the Department of Commerce on emerging technologies.[10]

This paper summarizes the current export controls on AI hardware, technology and software. It also explores whether any new export controls on AI, whether pursuant to the Wassenaar Arrangement Dual Use List or the new rubric of "emerging and foundational technologies," are likely to be effective, or possibly could prove to be counterproductive. It offers cautionary notes for policymakers who may be considering new export controls on AI. In addition, it makes recommendations for industry and academia, because AI likely will prove indispensable to future success across a wide range of applications.

AI ARCHITECTURE

In the absence of regulatory guidance on the topic, the presentation by Professor Artur Dubrawski of Carnegie Mellon University at the recent BIS Update Conference provides the best available summary of how the BIS may be thinking about AI. According to Professor Dubrawski, AI can be thought of as consisting of four distinct layers, each with its own challenges and related technologies: (1) the perception layer, (2) the learning layer, (3) the decision layer, and (4) the action layer.

The perception layer includes the sensors (hardware and software) that collect data. For example, it can include sensors for detecting the proximity of objects or optics used to capture facial features for biometric uses. According to Professor Dubrawski, the technologies at the perception layer are considered mature technologies, developed largely independent of AI. The learning layer includes data science and machine learning algorithms, which he also considers to be mature technologies. Access to large data sets also is required in order to have AI systems that learn quickly. The United States, Europe, and China are currently the main players in developing technology for the learning layer.

By comparison, the decision layer and action layer are still in the early stages of development. The decision layer is where planning occurs and where the software chooses a path to take based on the data available. The action layer is where the AI system realizes autonomy and takes an action based on its decision. Important research in this layer also includes human–machine interfaces.

For example, consider the application of AI to autonomous vehicles. The perception layer would be represented by the sensors on the vehicle. The learning layer would be represented by the software which analyzes the surrounding environment using the data from the sensors. The decision layer would be represented by the analysis of the learned information to conclude that a pedestrian was about to step in front of the vehicle. The action layer would be the transmission of the information to the vehicle's braking system, so that the pedestrian would not be harmed.

Not only are specific verticals like autonomous vehicles organizing around AI, Amazon reportedly is organizing its entire business around AI.[11] Clearly, the potential for AI to transform industries and companies is vast, but what does Dr. Dubrawski's presentation tell us about "chokepoint" technologies that may be susceptible of effective export controls?

One possible implication from Dr. Dubrawski's presentation is that export controls should not be focused on the relatively mature technologies at the perception and learning layers. However, it is precisely these two layers that are subject to export controls, today. Export controls on sensors (e.g., under Category 6 on the Commerce Control List of the EAR) are beyond the scope of this paper. Rather, we will focus on the electronics, computers, software, and technology for AI.

EXISTING EXPORT CONTROLS ON AI HARDWARE

More than two decades ago, the WA introduced multilateral export controls on AI hardware in its Dual Use List (WADUL),[12] and the United States adopted conforming amendments to the Commerce Control List (CCL) of the Export Administration Regulations.[13] (We also will consider important AI technologies that are *not* controlled under either the WADL or CCL.)

By way of background, in the early 1990s, the U.S. Department of Defense identified neural networks as "militarily critical" technologies:

> *Neural net technology, wherein a computational device is designed or modified to emulate in a simplistic manner the computational processes of the brain by utilizing a multiplicity of simple computational devices (artificial neurons) arrayed in large networks that can be trained – hence "neural networks."*
>
> *Neural networks, interconnected optically, and combined with Artificial Intelligence Expert Systems are projected to be developed into a hybrid, self-improving pattern recognition system.*[14]

About the same time, export controls on "neural network integrated circuits" were implemented under WADUL entry 3A1a9 and corresponding CCL entry 3A001.a.9. (The term "neural network integrated circuit" remains undefined.) "Neural computers" also were controlled under WADUL entry 4A4b and corresponding CCL entry 4A004.b. The term "neural computer" is defined as

> *A computational device designed or modified to mimic the behavior of a neuron or a collection of neurons (i.e., a computational device that is distinguished by its hardware capability to modulate the weights and numbers of the interconnections of a multiplicity of computational components based on previous data)*.[15]

Despite (or, perhaps because of?) these controls on neural network integrated circuits and neural computers, most of the hardware used for AI applications today consists of general purpose central processing units (CPUs) and graphics processing units (GPUs) combined into commodity cluster computer systems.[16] [17] Computer systems combining the CPUs with GPUs have demonstrated significantly superior performance to computers based on CPUs alone. These CPUs,[18] GPUs,[19] and computer systems[20] built therefrom generally are not controlled under the WADUL and are subject only to minimal controls under the CCL.

CPUs are classified on the CCL under 3A991 and controlled only for Anti-Terrorism reasons. GPUs are designated EAR99 subject to lowest level of controls. Computer systems built of CPUs and GPUs are controlled only for Anti-Terrorism reasons based on their Adjusted Peak Performance, because the contribution of the GPUs is not included in the calculation of the APP, and the APP of multiple CPUs is not aggregated pursuant to Note 5.[21] Indeed, anecdotal evidence suggests that neural network integrated circuits[22] and neural computers[23] remain research projects, with commercialization still perhaps decades hence. This experience certainly should give regulators pause, before they assume that they can identify and effectively subject any "emerging" or "foundational" AI hardware to export controls, today!

Why were these controls ineffective in the past, and what might be controlled effectively today? Perhaps because general purpose CPUs and GPUs are so inexpensive and ubiquitous, they rendered Neural Network integrated circuits largely irrelevant (up to this point). The computers that are performing most AI compute tasks today – general purpose computer clusters – have been decontrolled from the WADUL for many years, and they are subject to minimal controls on the CCL.[24] Like the CPUs and GPUs, they are inexpensive, ubiquitous, and so have rendered Neural Computers largely irrelevant (up to this point). Trying to control AI based on general purpose processors and computing clusters is not likely to be effective, given their widespread foreign availability,[25] worldwide.

The next question is whether it might be possible to control the technology for AI. We will address general purpose AI technology below. However, if we limit our consideration for the moment to controls on technology for the next generation of specialized AI processors, the current controls would not seem likely to be effective. There has been a veritable tsunami of investment in special purpose AI hardware, both inside and outside the United States.[26,27,28,29]

As PC sales have declined, Intel, for example, has invested in AI hardware because it has been depending on its sales to data centers which provide services to mobile and web-based apps.[30] Those apps rely on AI for features like photo and speech recognition. Adjusting to this new business strategy, Intel has modified its CPUs to become more than 200 times better at AI training, which resulted in $1 billion in sales of its Xeon processors in 2017 for this market.[31]

In the United States, special purpose accelerators are under development by established companies, including Alphabet, Apple, Intel, and Microsoft among others.[32] Venture capital funding is abundant for startup companies, including Cerebras, Graphcore, Groq, SimpleMachines, Wave Computing,[33] and others. While each approach is somewhat different, these special purpose accelerators really reflect known mathematical operations optimized for various AI algorithms.

Outside the United States, the landscape is similar. For example, established Chinese companies,[34] like Baidu and Huawei, and startups, like Cambricon,[35] are reportedly developing AI chips that are competitive with non-Chinese offerings. In 2017, China accounted for 48% of worldwide funding for AI startups, versus 38% for the United States,[36] (although some suggest that the wave of money washing over the Chinese AI startups may be about to recede).[37]

The most likely tweaks to the WADUL and CCL that would result in new controls on the integrated circuits and underlying technologies are amendments to WADUL items 3A1, 3E2 and CCL 3A001 and 3E002. These controls currently apply only to microprocessors and related technology. However, simple changes could extend their coverage to include not only CPUs, but also GPUs, and various items sometimes described as "neural processing units," "neuromorphic processors," "tensor processing units," and "vision processing units."[38]

The most likely tweaks to the WADUL and CCL that would result in new controls on the computer systems are amendments to the Notes to the Calculation of Adjusted Peak Performance in Category 4 on the WADUL and CCL. In particular, Notes 3, 4, 5, and 6 specify the conditions under which the performance of computing elements must be aggregated.[39] Requirements to include GPUs and other accelerators for AI applications could result in a significant increase in the controls on computer systems used for AI application.

However, controls on these types of items would only be addressing the learning layer of AI, which is already a mature technology. The controls would not be capturing the technologies that are still in the early stages of development, which are those dealing with the decision and action layers of AI. Therefore, such controls are likely to have little impact on foreign advancement of AI.

PUTATIVE CONTROLS ON AI TECHNOLOGY

It is difficult to discern the precise reasons why the WA and its members have not attempted to implement export controls on AI technology, because the minutes of the WA's discussion are not a matter of public record. However, one might speculate that the reasons include how old and how widely published AI's key concepts really are. Academic literature demonstrates that the academic research on AI began over sixty years ago. An early example is the 1955 Dartmouth Summer Research Project Proposal on Artificial Intelligence authored by John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon.[40]

"Modern" AI can be dated to Geoffrey Hinton's breakthrough paper, with colleagues David Rumelhart and Ronald Williams, which was published in 1986.[41] The paper elaborated on a technique called backpropagation, or "backprop." Backprop, in the words of Jon Cohen, a computational psychologist at Princeton, is "what all of deep learning is based on – literally everything."[42] Commentators have offered complicated taxonomies of AI.[43] There are clever algorithms with modern AI, such as merging of Bayesian statistics with deep learning. This offers more probabilistic inferences about the world, allowing AI systems to handle uncertainty better. Nevertheless, distilled to its essence, AI today is deep learning, and deep learning is backprop. How can you control backprop technology that was published in the open literature more than 30 years ago? Under the current EAR, such technology is exempt from control, as "fundamental research."[44]

Indeed, even if we assume for the sake of the argument that AI is more than backprop, the next question is whether it is possible to control "emerging" and "foundational" AI technology, today. Again, the answer is probably negative. Many of the leading minds split time between academia and companies. Examples include not only the aforementioned Geoffrey Hinton[45] (University of Toronto and Google) but also Yann LeCun [46] (NYU and Facebook) and many others. Controlling their work would require a radical re-consideration of the "fundamental research" exemption. In addition, these leading researchers do not always remain in the United States, or go back and forth. Qi Lu,[47] an important AI researcher, who left an Executive Vice President role at Microsoft to become the Chief Operating Officer at Baidu, recently announced that he would return to the United States. Another interesting fact about these three leading minds of AI: None of them was born in the United States.

The list goes on. Andrew Chi-Chih Yao,[48] a Turing Award winner who renounced US citizenship, is now researching AI theory development in China. Tim Byrnes,[49] an Australian physicist, is researching quantum computing at NYU's campus in Shanghai. Zhang Liang-jie,[50] formerly employed by IBM in its Watson AI team, is now chief scientist of the enterprise software group at Kingdee in Shenzhen. Zenglin Xu,[51] formerly of Purdue University, now leads AI research at the (notorious) University of Electronic Science and Technology in China.[52]

If we exclude individual academic researchers, and consider controlling corporate research on AI, we see a similar cross-pollination. Google has set up an AI research facility in China. Google AI China Center is intended to help the company conduct AI research in the country and hire employees with backgrounds in machine learning.[53] Vice versa, Chinese internet company Baidu has set up an AI research facility in America. Baidu revealed plans to invest in autonomous car technology research and development (R&D) from its base in Silicon Valley. It would be extremely difficult, or perhaps impossible, to control corporate AI research within national boundaries.

According to a report in the New York Times, the Trump Administration is considering new controls on the release of technology to foreign researchers[54] at American colleges and universities. Although the article does not mention AI, specifically, it is not difficult to imagine that any new rules could cover "emerging" and "foundational" technologies perhaps including some related AI.

PUTATIVE CONTROLS ON AI SOFTWARE

As in the case of AI technology, neither the WADUL nor the CCL implement explicit controls on AI software. Could companies' proprietary AI software be controlled effectively? Some companies in both the United States and China retain their AI software as proprietary, and hence conceivably susceptible of effective control. Amazon's AI software for shopping recommendations and making its physical operations more efficient, Netflix's recommendation engine, or Tesla and Baidu's autonomous vehicle software, might fall into this category.

On the other hand, there is plenty of high-quality open source software for AI, including Google's TensorFlow,[55] Facebook's Caffe2[56] and Pytorch,[57] and the OpenAI consortium[58] (which has no corporate affiliation). Imposing export controls on proprietary AI software would do nothing more than favor one business model (open source) over another business model (proprietary code), unless the United States attempts to prohibit publication of open source software, with the attendant First Amendment issues looming large.[59]

PUTATIVE CONTROLS ON AI DATA

The next question is whether the data used in AI training can be controlled effectively? After all, AI requires quite a lot of data, today.[60] Moreover, data are probably the one area where China already may have an advantage. There simply are more Chinese people, more Chinese cars, more Chinese almost everything connected to the internet, generating more data, than exist in the United States. Based on simple demographics, that is not likely to change, anytime soon. Export controls that focus on access to data likely would perpetuate China's advantage.

There are two problems with trying to control the data used in AI. The first is that the United States only controls "technology" under the EAR, and most of the data used to train AI are not "technology" so-defined.[61] The second is that AI is being optimized to operate on smaller and smaller data sets. Any control on data today could prove to be ephemeral.[62]

AI CONTROLS BASED ON END-USER AND END-USE

If list-based export controls on AI hardware, software and technology are not likely to be effective, then the next question is whether there is a role for export controls to play in AI, based on end-user or end-use? There are at least two possible approaches.

The first approach would be to control AI having exclusively military applications as a military item under the WA Munitions List (WAML)[63] and the corresponding United States Munitions List (USML)[64] of the International Traffic in Arms Regulations (ITAR). For example, training a conventional AI system using technical data controlled under the ITAR likely would result in a model that similarly is subject to the ITAR. If indeed this is the result, then it may have the perverse result of retarding the development of AI in the private sector, because not many companies will want to have their models subject to the ITAR. (See "Export Controls on AI Could be *Counterproductive*" below.)

The second approach would be to identify and control only specific end-users and end-uses that raise security concerns. For example, exporters could be required to apply for an export license, if they know or have reason to know that a particular end-user was engaged in military activities.[65] The government could impose an export license requirement for end-users it knows to be engaged in military activities.[66] The export controls on AI would resemble the current controls on proliferation of weapons of mass destruction.[67]

Nevertheless, twenty-five years of experience with end-user and end-use controls suggests that we should be cautious about expecting either of these mechanisms to be successful, and they could have unforeseen consequences. A good example is the sanctions on the National University of Defense Technology and the three high-performance computing centers that were imposed several years ago, ostensibly due to nuclear proliferation concerns.[68] While those sanctions may have retarded the development of the Sunway Tianhe 2 computer, which relied on Intel CPUs and NVIDIA GPUs, its successor, the more powerful Sunway TaihuLight, now uses only indigenous Chinese processors.[69] The net result is that the end-user and end-use controls may simply have accelerated the Chinese development of indigenous processors, while at the same time reducing the available market in China for Intel and NVIDIA. Indeed, it may be one of the reasons why the Chinese government reportedly has increased its investments in key semiconductor technologies, thereby further reducing the long-term opportunities for U.S. companies in the Chinese market.

AI AND CHINA

Although neither the WA nor its members explicitly target potential adversaries, the "elephant in the room" that cannot be ignored is China. Are the United States and the members of the WA doomed to be surpassed by China in AI?

There is no clear answer, but there are some intriguing facts. Chinese researchers are publishing more academic papers, but American research papers are more frequently cited, suggesting that they are more influential. The United States still leads the world in AI patents, but the number of Chinese patents is growing faster. For five years, China had the world's fastest computer, a symbolic achievement for a country trying to show that it is a technology powerhouse. Earlier this year, the United States retook the lead thanks to a machine, called Summit, built for the Oak Ridge National Laboratory in Tennessee.[70] However, China has the most supercomputers. From a list of the 500 fastest machines, Chinese companies and government produced 206 of them, while the United States made 124.[71] It is interesting to note the correlation to the United States imposed sanctions in 2015. Would sanctions on exports of AI technology also perversely accelerate Chinese investment to the detriment of U.S. industry?

A recent Oxford University study[72] posited that the United States remains preeminent in almost every important area of AI, except perhaps access to data. Other, similar Chinese initiatives in the fields of semiconductor technology and biotechnology, for example, suggest that extravagant claims in the Chinese State Council's "New Generation AI Development Plan" should be taken with a grain (perhaps, even a boulder) of salt.

EXPORT CONTROLS ON AI COULD BE *COUNTERPRODUCTIVE*

The importance of AI to the Department of Defense is well documented.[73] However, if the objective of export controls is to maintain a military advantage, William Carter, Deputy Director and Fellow, Technology Policy Program, Center for Strategic and International Studies before the House Armed Services Committee Subcommittee on Emerging Threats and Capabilities, January 9, 2018, testified that export controls on AI could be *counterproductive*:

> *Trying to control our own commercial technologies as "dual use" only deters private companies from working with DoD to protect their freedom to market their products internationally, and paying defense contractors to re-invent the wheel by building bespoke versions of commercial technologies for a DoD client has proven ineffective and wasteful, draining our resources and causing the military to fall dramatically behind the private sector in even simple day-to-day technologies. Perhaps it is time for a new way of thinking about maintaining a technological edge for the military.*[74]

Some might suggest that these concerns are overstated, and that private companies will join the national security establishment in facing an existential challenge to American leadership in AI. To its credit, the Pentagon has been looking for new ways to engage industry and academia in response to the reported Chinese "military-civil fusion" of AI research and development.[75] Such optimism should be tempered by evidence that American private companies may not be inclined to support the Defense Department or Homeland Security, at least not without reservations.[76]

For example, Google's employees have circulated an open letter opposing the company's work on autonomous vehicles and AI for the Department of Defense,[77] and Facebook's controlling shareholder and CEO has testified before the Congress that new legislation would be required for his company to support facial recognition technology for the Homeland Security.[78] In the long run, the commercial AI market likely is much larger than the defense market for AI. Google's overall commitment to AI surely dwarfs its revenue from Project Maven. Project Maven was allocated $16 million in last year's NDAA, but funding increased to $93.1 million in this year's bill.[79] Ultimately, Google decided to terminate Project Maven. Google's new corporate policy declared that it will not support certain kinds of future Department of Defense (DoD) programs.[80] Now, the DoD will be seeking a new private sector partner to replace Google, since Project Maven is explicitly a public–private partnership – "a commercial technology initiative that inserts commercial AI into existing programs of records."[81]

Employees of Amazon, Microsoft, Salesforce, and other large information technology companies have expressed similar concerns. Amazon employees protested the sale of facial recognition software to law enforcement.[82] Microsoft employees protested the sale of AI technology to Immigration and Customs Enforcement.[83] Salesforce employees pushed back against contracts with military and other government agencies.[84] Elon Musk, Deep Mind founders, and others pledge not to develop lethal AI.[85] Perhaps, recognizing that it may not be able to rely on commercial technologies, the Pentagon recently signed an $885 million contract with one of its traditional suppliers (Booz Allen) to develop AI for its requirements, ranging from national security to health care.[86] At the same time, companies like NVIDIA continue to develop systems that support the development of AI, for non-military applications.[87,88] Perhaps, as Mr. Carter suggests, it is time to consider a new approach, beyond export controls, to retain the U.S. military's advantage.

The new NDAA is evidence that the Congress agrees with Mr. Carter that new approaches may be required. Section 238(e) of the NDAA mandates a complete study of past and current advances of AI, and the future of AI, including the methods and means necessary to advance the state of the art.[89] Requirements for the study include: (1) reviewing advances in AI and associated technologies relevant to the need of the DoD and the armed forces, (2) the competitiveness of the DoD in AI, (3) recommendations on how the DoD can maintain a technical advantage, (4) recommendations on establishing data standards and incentives for sharing open training data, (5) recommendations for engagement with other relevant agencies involved with AI, and (6) recommendations for legislative action relating to AI, including recommendations to more effectively fund and organize the DoD.[90]

Investment in AI for the U.S military also is included in the new NDAA. The Act includes $15 million to "enhance and accelerate artificial intelligence research" in the service branches, $10 million for the Air Force, and $5 million for the Army.[91] The Pentagon wants to establish new agencies for AI and requested $70 million for a new Joint Artificial Intelligence Center (JAIC).[92] The bill also authorized $10 million for a new National Security Commission on AI.[93]

CONCLUSION

AI has been cited as "militarily critical technology" for over 25 years, and it has been subject to dual use export controls for over 20 years. Nevertheless, export controls have had a relatively minor impact on AI development.

Perhaps, the reason export controls have not been effective is that there is no "chokepoint" AI technology that is susceptible of effective control. Whether we consider hardware, technology, or software, the opportunities for effective export controls are elusive. It may be possible to implement export controls on a particular military application and related data set. For example, pattern recognition hardware, software, and technology used to identify a cat, and to identify an enemy aircraft, are quite similar, except for the data sets used to train the system. However, broad based export controls on dual use technology are problematic, at best.

Despite cautionary voices, like Dov Zakheim, former Under Secretary (Comptroller) and Chief Financial Officer for the Department of Defense from 2001 to 2004 and a Deputy Under Secretary of Defense from 1985 to 1987, who recently opined that "AI often is erroneously identified as an emerging technology,"[94] policymakers may be tempted to implement new export controls on AI hardware, software, and technology in the name of national and economic security. These new export controls most likely would be under the WADUL, under the CCL, or under the rubric of "emerging and foundational" dual use technologies. Well-intentioned as such attempts might be, any such export controls must be carefully crafted. There is a significant possibility that they will be ineffective, at best, or possibly even counterproductive.

Industry and academia have a crucial role to play, helping to ensure that any new export controls are narrowly tailored to protect legitimate national and economic security interests without unnecessarily hindering the development of AI.[95] Industry and academic participation in the regulatory process will be important, both through the existing Technical Advisory Committees, as well as its recently chartered Emerging Technologies Technical Advisory Committee. It will be important to file public comments on any proposed rules or notices of inquiry that the relevant agencies may publish. The difficult task of designing and implementing effective export controls on AI and eschewing controls that could be counterproductive has only just begun.

NOTES

[1] John S. McCain National Defense Authorization Act for Fiscal Year 2019, Pub. L. No. 115-232, § 238(g), 132 Stat. 1636 (2018). Artificial Intelligence includes (1) Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight or that can learn from experience and improve performance when exposed to data sets.

(2) An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action. (3) An artificial system designed to think or act like a human, including cognitive architectures and neural networks. (4) A set of techniques, including machine learning, that is designed to approximate a cognitive task. (5) An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.

[2] National Security Strategy of the United States, December 18, 2017, https://www.whitehouse.gov/wp-content/uploads/2017/12/NSS-Final-12-18-2017-0905-2.pdf.

[3] www.wassenaar.org (follow "About Us" hyperlink).

[4] Ambassador Philip Griffiths, Remarks at the 2017 Export Control Forum on WMD/Military Proliferation Trends and Emerging Technologies of Concern (transcript available at http://trade.ec.europa.eu/doclib/docs/2017/december/tradoc_156485.pdf).

[5] Press Release, The White House, Statement by President Donald J. Trump on H.R. 5515 (Aug. 13, 2018) (available at: https://www.whitehouse.gov/briefings-statements/statement-president-donald-j-trump-h-r-5515/).

[6] John S. McCain National Defense Authorization Act for Fiscal Year 2019, Pub. L. No. 115-232, § 1703(a)(6), 132 Stat. 1636 (2018). Title XVII, Subtitle A – Committee on Foreign Investment in the United States authorizes restrictions on investments in "critical technologies," including "emerging and foundational technologies controlled pursuant to Section 1758 of the Export control Reform Act of 2018."

[7] Although this paper is not concerned primarily with investment restrictions, it may be of interest that Xilinx of the United States recently announced its acquisition of Dee Phi, an AI startup in Beijing. Most Chinese investments in the United States have been smaller investments in venture capital and private equity funds that, in turn, invest in the United States.

[8] Press Release, The White House, Statement on Steps to Protect Domestic Technology and Intellectual Property from China's Discriminatory and Burdensome Trade Practices (May 29, 2018) (available at: https://www.whitehouse.gov/briefings-statements/statement-steps-protect-domestic-technology-intellectual-property-chinas-discriminatory-burdensome-trade-practices/).

[9] 83 Fed. Reg. 58201 (November 19, 2018).

[10] Emerging Technology Technical Advisory Committee (ETTAC), 83 Fed. Reg. 39,054 (August 8, 2018).

[11] Blake Morgan, *How Amazon Has Reorganized Around Artificial Intelligence And Machine Learning*, Forbes, July 16, 2018, https://www.forbes.com/sites/blakemorgan/2018/07/16/how-amazon-has-re-organized-around-artificial-intelligence-and-machine-learning/#549c9ddd7361.

[12] Wassenaar Dual Use List.

[13] Commerce Control List of EAR.

[14] Department of Defense, Militarily Critical Technologies List, 1992.

[15] EAR Part 772.

[16] Stephen Nells, *Intel sold $1 billion of artificial intelligence chips in 2017,* Business Insider, Aug. 8, 2018, https://www.businessinsider.com/r-intel-sold-1-billion-of-artificial-intelligence-chips-in-2017-2018-8.

[17] Seth Archer, *Nvidia gave away its newest AI chips for free – and that's part of the reason why it's dominating the competition (NVDA)*, Business Insider, Jul. 25, 2017, https://markets.businessinsider.com/news/stocks/nvidia-stock-price-gave-its-first-volta-v100-for-free-and-is-dominating-its-competition-2017-7-1002201278?utm_content=bufferdd755&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer.

[18] *See generally* Intel Processors Page, https://www.intel.com/content/www/us/en/products/processors.html.

[19] *See generally* NVIDIA Store, https://www.nvidia.com/en-us/shop/.

[20] *See generally* Amazon Web Services Home Page, https://aws.amazon.com/, Microsoft Azure Home Page, https://azure.microsoft.com/en-us/, Google Cloud Home Page https://cloud.google.com/.

[21] *But see,* Note 3, if the manufacturer makes certain performance claims.

[22] *See* IBM Research Home Page, http://www.research.ibm.com/articles/brain-chip.shtml.

[23] Differentiable Neural Computers, https://deepmind.com/blog/differentiable-neural-computers/ (Oct. 12, 2016).

[24] Most commercial computer clusters are classified under ECCN 4A994.b or 5A992.c, subject only to AT controls.

[25] 15 CFR §768 (2001).

[26] Cade Metz, *Big* Bets on A.I. Open a New Frontier for Chip Start-Ups, Too, *The New York Times*, Jan. 14, 2018, https://www.nytimes.com/2018/01/14/technology/artificial-intelligence-chip-start-ups.html.

[27] Yair Snir, AI Is Making Hardware Sexy Again, Forbes, Jun. 21, 2018, https://www.forbes.com/sites/startupnationcentral/2018/06/21/ai-is-making-hardware-sexy-again/#a9034d5277e2.

[28] Sean Gallagher, China producing x86 chips nearly identical to AMD server processors, arsTECHNICA, July 9, 2018, https://arstechnica.com/information-technology/2018/07/china-producing-x86-chips-nearly-identical-to-amd-server-processors/.

[29] Andy Patrizio, *The AI revolution has spawned a new chips arms race,* arsTECHNICA, July 9, 2018, https://arstechnica.com/gadgets/2018/07/the-ai-revolution-has-spawned-a-new-chips-arms-race/.

[30] Stephen Nellis, *Intel sold $1 billion of artificial intelligence chips in 2017,* Reuters, Aug. 8, 2018, https://www.reuters.com/article/us-intel-tech/intel-sold-1-billion-of-artificial-intelligence-chips-in-2017-idUSKBN1KT2GK?utm_source=applenews.

[31] *Id.*

[32] Karl Freund, *Will ASIC Chips Become The Next Big Thing in AI?,* Forbes, Aug. 4, 2017, https://www.forbes.com/sites/moorinsights/2017/08/04/will-asic-chips-become-the-next-big-thing-in-ai/#7062cb0e11d9.

[33] *See generally* Cerebras Home Page https://www.cerebras.net/, Graphcore Home Page, https://www.graphcore.ai/, Groq Home Page, https://groq.com/, Simple Machines Home Page, http://www.simplemachinesinc.com/wp/, Wave Computing Home Page, https://wavecomp.ai/.

[34] Yiting Sun, *China wants to make the chips that will add AI to any gadget*, MIT Technology Review, Jan. 24, 2018, https://www.technologyreview.com/s/609954/china-wants-to-make-the-chips-that-will-add-ai-to-any-gadget/.

[35] *See generally,* Cambricon Home Page, http://www.cambricon.com/.

[36] Peter H. Diamandis, *China Is Quickly Becoming an AI Superpower,* SingularityHub, Aug. 29, 2018, https://singularityhub.com/2018/08/29/china-ai-superpower/#sm.00015zktky12hfj1pzk2c0t33tdzu.

[37] Sarah Dai, *Investor warns of day of reckoning for 90 per cent of Chinese AI start-ups as funding dries up,* South China Morning Post, Aug. 27, 2018, https://www.scmp.com/tech/article/2161387/investor-warns-day-reckoning-90-pc-chinese-ai-start-ups-funding-dries.

[38] Jennifer Chu, *Engineers design artificial synapse for "brain-n-a-chip" hardware,* MIT News, Jan. 22, 2018, http://news.mit.edu/2018/engineers-design-artificial-synapse-brain-on-a-chip-hardware-0122?utm_campaign=Revue%20

newsletter&utm_medium=Newsletter&utm_source=The%20Wild%20Week%20in%20AI.

39. Bureau of Industry and Security Commerce Control List (CCL), https://www.bis.doc.gov/index.php/licensing/commerce-control-list-classification/commerce-control-list-ccl/17-regulations/139-commerce-control-list-ccl (last visited Aug. 15, 2018).
40. *See* https://aaai.org/ojs/index.php/aimagazine/article/view/1904/1802.
41. *See* http://www.cs.toronto.edu/~hinton/absps/naturebp.pdf.
42. James Somers, *Is AI Riding a One-Trick Pony,* MIT Technology Review, Sept. 29, 2017, https://www.technologyreview.com/s/608911/is-ai-riding-a-one-trick-pony/.
43. Francesco Corea, *AI Knowledge Map: How to Classify AI Technologies,* Forbes, Aug. 22, 2018, https://www.forbes.com/sites/cognitiveworld/2018/08/22/ai-knowledge-map-how-to-classify-ai-technologies/#5b8cff437773.
44. Section 734.8 of the EAR
45. *See* http://www.cs.toronto.edu/~hinton/bio.html.
46. *See* http://yann.lecun.com/ex/bio.html.
47. *See* http://ir.baidu.com/phoenix.zhtml?c=188488&p=irol-govBio&ID=255244.
48. *See* http://iiis.tsinghua.edu.cn/yao/.
49. *See* https://shanghai.nyu.edu/academics/faculty/directory/tim-byrnes.
50. *See* https://www.igi-global.com/affiliate/liang-jie-zhang/79289.
51. *See* http://www.bigdata-research.org/people/faculty/5.html.
52. 77 FR 58006 (no. 182). Note that the University of Electronic Science and Technology is listed on the Entity List of the EAR.
53. *See generally* Jacques Bughin et al., June 2017, *Artificial Intelligence The Next Digital Frontier?* (https://www.mckinsey.com/~/media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx) (explaining that companies are expanding their search for talent abroad by opening R&D offices in other countries).
54. Ana Swanson and Keith Bradsher, *White House Considers Restricting Chinese Researchers Over Espionage Fears,* The New York Times, Apr. 30, 2018, https://www.nytimes.com/2018/04/30/us/politics/trump-china-researchers-espionage.html.
55. *See generally* TensorFlow Home Page https://www.tensorflow.org/.
56. *See generally* Cafe2 Home Page https://caffe2.ai/.
57. *See generally* PyTorch Home Page http://pytorch.org/.
58. *See generally* Open AI Home Page https://openai.com/.
59. *Bernstein v. United States Dept. of Justice*, 192 F.3d 1308 (9th Cir. 1999).
60. *See generally* Jacques Bughin et al., June 2017, *Artificial Intelligence The Next Digital Frontier?* (https://www.mckinsey.com/~/media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx) (explaining that "the world is generating vast quantities of the fuel that powers AI – data. Billions of gigabytes of it every day.").
61. *See* Part 772 of the EAR Technology definition. https://www.bis.doc.gov/index.php/documents/regulation-docs/434-part-772-definitions-of-terms/file (last visited Aug. 15, 2018).
62. Yiting Sun, *More efficient machine learning could upend the AI paradigm,* MIT Technology Review, Feb. 2, 2018, https://www.technologyreview.com/s/610095/more-efficient-machine-learning-could-upend-the-ai-paradigm/.
63. Wassenaar Arrangement, *List of dual-use goods and technologies and munitions list.* (Dec., 2017) https://www.wassenaar.org/app/uploads/2018/01/WA-DOC-17-PUB-006-Public-Docs-Vol.II-2017-List-of-DU-Goods-and-Technologies-and-Munitions-List.pdf.
64. ITAR Part 121.1.
65. EAR Part 744.21.
66. Bureau of Industry and Security Entity List, https://www.bis.doc.gov/index.php/policy-guidance/lists-of-parties-of-concern/entity-list (last visited Aug. 15, 2018).
67. EAR Sections 744.2, 744.3, 744.4.
68. 80 FR 8524.
69. James Vincent, *Chinese super computer is the world's fastest – and without using US chips,* The Verge, June 20, 2016, https://www.theverge.com/2016/6/20/11975356/chinese-supercomputer-worlds-fastes-taihulight.
70. Steve Lohr, *Move Over, China: U.S. Is Again Home to World's Speediest Supercomputer*, *The New York Times*, June 8, 2018, https://www.nytimes.com/2018/06/08/technology/supercomputer-china-us.html.
71. Steve Lohr, *China Extends Lead as Most Prolific Supercomputer Maker, The New York Times*, June 25, 2018, https://www.nytimes.com/2018/06/25/technology/china-supercomputers.html.
72. Jeffrey Ding, Deciphering China's AI Dream, The context, components, capabilities, and consequences of China's strategy to lead the world in AI (2018), https://www.fhi.ox.ac.uk/wp-content/uploads/Deciphering_Chinas_AI-Dream.pdf.
73. Richard Potember, January 2017, *Perspectives on Research in Artificial Intelligence and Artificial General Intelligence Relevant to DoD*, https://fas.org/irp/agency/dod/jason/ai-dod.pdf.
74. William Carter, Statement Before the House Armed Services Committee Subcommittee on Emerging Threats and Capabilities "Chinese Advances in Emerging Technologies and their Implications for U.S. National Security", (Jan. 9, 2018) (transcript available at https://csis-prod.s3.NVIDIAaws.com/s3fs-public/congressional_testimony/ts180109_Carter_Testimony.pdf?zmxasiIZi6jHZPgAAsMYcSiSMwdw6LgJ).
75. Cade Metz, *Artificial Intelligence Is Now a Pentagon Priority. Will Silicon Valley Help? The New York Times*, Aug. 26, 2018, https://www.nytimes.com/2018/08/26/technology/pentagon-artificial-intelligence.html?rref=collection%2Fsectioncollection%2Ftechnology&action=click&contentCollection=technology&region=stream&module=stream_unit&version=latest&contentPlacement=2&pgtype=sectionfront.
76. Scott Malcomson, *Why Silicon Valley Shouldn't Work With the Pentagon, The New York Times*, Apr. 19, 2018, https://www.nytimes.com/2018/04/19/opinion/silicon-valley-military-contract.html?rref=collection%2Fsectioncollection%2Fopinion&action=click&contentCollection=opinion&region=rank&module=package&version=highlights&contentPlacement=6&pgtype=sectionfront.
77. *See* Scott Shane & Daisuke Wakabayashi, "*The Business of War": Google Employees Protest Work for the Pentagon*, The New York Times, Apr. 4, 2018, https://www.nytimes.com/2018/04/04/technology/google-letter-ceo-pentagon-project.html.
78. *See* Rachel Kraus, *When ICE asks Facebook for Surveillance help, Zuck Pledges to Say "No,"* Mashable, Apr. 10, 2018, https://mashable.com/2018/04/10/facebook-wont-comply-with-ice-extreme-vetting-immigrants/.
79. Jay Cassano, *Pentagon's artificial intelligence programs get huge boost in defense budget,* Fast Company, Aug. 15, 2018, https://www.fastcompany.com/90219751/pentagons-artificial-intelligence-programs-get-huge-boost-in-defense-budget?partner=rss&utm_campaign=rss+fastcompany&utm_content=rss&utm_medium=feed&utm_source=rss.
80. *See generally* Artificial Intelligence at Google Page https://ai.google/principles.
81. *See supra* note 79.
82. Ali Breland, *Amazon Employees Protest Sale of Facial Recognition Tech to Law Enforcement*, The Hill, June 21, 2018, http://thehill.

com/business-a-lobbying/393583-amazon-employees-protest-sale-of-facial-recognition-tech-to-law.

83. Sheera Frankel, *Microsoft Employees Protest Work with ICE as Tech Industry Mobilizes over Immigration*, The New York Times, June 19, 2018, https://www.nytimes.com/2018/06/19/technology/tech-companies-immigration-border.html.

84. Caroline O'Donovan, *Employees of Another Major Tech Company Are Petitioning Government Contracts*, BuzzFeed, June 26, 2018, https://www.buzzfeed.com/carolineodonovan/salesforce-employees-push-back-against-company-contract?utm_term=.ttRLjnElg#.ogmm12VLM.

85. James Vincent, *Elon Musk, DeepMind Founders, and Others Sign Pledge to Not Develop Lethal AI Weapon Systems*, The Verge, July 18, 2018, https://www.theverge.com/2018/7/18/17582570/ai-weapons-pledge-elon-musk-deepmind-founders-future-of-life-institute.

86. Pentagon Signs $885 Million Artificial Intelligence Contract with Booz Allen, https://blogs.wsj.com/cio/2018/07/30/pentagon-signs-885-million-artificial-intelligence-contract-with-booz-allen/ (July 30, 2018, 5:06 PM ET).

87. Kevin Cook, *NVIDIA Gaming Drives the Deep Learning-AI Revolution*, Zacks, Aug. 21, 2018, https://www.zacks.com/commentary/177771/nvidia-gaming-drives-the-deep-learning-ai-revolution?cid=CS-APPLE-FT-177771.

88. Hector Marinez, *NVIDIA to Collaborate with DARPA to Develop Systems for Post-Moore's Law Era*, NVIDIA, July 24, 2018, https://blogs.nvidia.com/blog/2018/07/24/darpa-research-post-moores-law/.

89. John S. McCain National Defense Authorization Act for Fiscal Year 2019, Pub. L. No. 115-232, § 238(e)(1), 132 Stat. 1636 (2018).

90. *Id.*

91. Jay Cassano, *Pentagon's Artificial Intelligence Programs Get Huge Boost in Defense Budget*, Fast Company, Aug. 15, 2018, https://www.fastcompany.com/90219751/pentagons-artificial-intelligence-programs-get-huge-boost-in-defense-budget?partner=rss&utm_campaign=rss+fastcompany&utm_content=rss&utm_medium=feed&utm_source=rss.

92. *Id*.

93. *Id*.

94. Dov S. Zakheim, *What America Must Do to Remain the World's High-tech Leader*, The Hill, Aug. 9, 2018, http://thehill.com/opinion/technology/401034-what-america-must-do-to-remain-the-worlds-high-tech-leader.

95. *See generally* Jacques Bughin et al., *Notes from the Frontier: Modeling the Impact of AI on the World Economy* (McKinsey & Company, Sept. 2018), https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-frontier-modeling-the-impact-of-ai-on-the-world-economy (explaining that there is large potential for AI to contribute to global economic activity and there are several barriers which might hinder AI adoption and absorption).