



Homeland Defense & Security  
Information Analysis Center



# Artificial Intelligence

*State of the Art Report*

DISTRIBUTION A. APPROVED FOR PUBLIC RELEASE: DISTRIBUTION UNLIMITED

**STATE OF THE ART REPORT:  
Artificial Intelligence and Machine Learning for Defense Applications**

Gregory Nichols, Homeland Defense & Information Analysis Center  
Sue Ellen Haupt, National Center for Atmospheric Research  
David John Gagne, National Center for Atmospheric Research  
Anthony A. Rucci, GRIDSMART Technologies, Inc.  
Joel Hewett, Homeland Defense & Security Information Analysis Center  
Gopikrishna Deshpande, Auburn University  
Pradyumna Lanka, Auburn University  
Susan A. Youngblood, Auburn University

2017

Contract Number: FA8075-13-D-0001

**Sponsorship Statement:**

HDIAC is a DoD sponsored Information Analysis Center (IAC) with policy oversight provided by the Under Secretary of Defense for Research and Engineering (USD (R&E)), and administratively managed by the Defense Technical Information Center (DTIC). HDIAC is operated by Information International Associates (IIa) in Oak Ridge, Tenn. Reference herein to any specific commercial products, processes or services by trade name, trademark, manufacturer or otherwise, does not necessarily constitute or imply its endorsement, recommendation or favoring by the United States government or HDIAC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or HDIAC and shall not be used for advertising or product endorsement purposes.

**About This Publication:**

Copyright 2018 by IIa. This State of the Art Report was developed by IIa under HDIAC contract FA8075-13-D-0001. The government has unlimited free use of and access to this publication and its contents in both print and electronic versions. Subject to the rights of the government, this document (print and electronic versions) and the contents contained within it are protected by U.S. copyright law and may not be copied, automated, resold or redistributed to multiple users without the written permission of HDIAC.

**Distribution A:**

Approved for Public Release; Distribution Unlimited

This page intentionally left blank.

## Contents

<b>AUTHORS</b> .....	<b>4</b>
<b>ACRONYMS</b> .....	<b>6</b>
<b>EXECUTIVE SUMMARY</b> .....	<b>7</b>
<b>1 INTRODUCTION</b> .....	<b>10</b>
<b>2 USE OF ARTIFICIAL INTELLIGENCE IN ENABLING RENEWABLE ENERGY</b> .....	<b>23</b>
Introduction, 23	
Forecasting Renewable Energy and AI, 25	
A Look to the Future, 34	
Summary and Conclusions, 36	
<b>3 BIOMETRICS, CRITICAL INFRASTRUCTURE PROTECTION, AND HOMELAND DEFENSE AND SECURITY</b> .....	<b>45</b>
<b>4 CHEMICAL, BIOLOGICAL, RADIOLOGICAL, AND NUCLEAR DEFENSE/WEAPONS OF MASS DESTRUCTION</b> .....	<b>52</b>
<b>5 ARTIFICIAL INTELLIGENCE APPLICATIONS TO MILITARY MEDICINE</b> .....	<b>59</b>
Introduction, 59	
Applications of AI in Mental Health Disorders and Brain Injuries, 60	
Data Sharing and Big Data Neuroimaging Initiatives, 63	
Conclusion, 64	
<b>6 CONCLUSION</b> .....	<b>71</b>

## Executive Summary

The Homeland Defense & Security Information Analysis Center (HDIAC) develops State of the Art Reports on relevant, key topics for the Department of Defense (DoD). One of the most significant topics within the Defense community is artificial intelligence (AI). AI and some of its branches, primarily machine learning, were prominent in the 2014 DoD Third Offset Strategy, which guided DoD's approach to combat operations and defensive posture [1]. AI research has also become pervasive in academia and nearly all industrial sectors. Some of this research produces outcomes being used by DoD and also highlights new risks and challenges that could have implications for the DoD regarding how it conducts operations in the future.

AI research in general, as well as its use in military applications, has its roots in the 1950s. Alan Turing, a British mathematician and computer scientist, asked the question, "Can machines think [2]?" This introduced a realistic tenor to the field of AI, a subject once isolated to the realm of science fiction. Almost immediately after the publication of Turing's article, the Office of Naval Research (ONR) funded early work into artificial neural networks, which form the basis of machine learning and are required for deep learning [3]. Subsequently, AI efforts, both in defense and civilian research, were led by the Advanced Research Projects Agency (ARPA, which was later renamed Defense Advanced Research Projects Agency [DARPA]), which continues to this day [4]. However, additional research is now being conducted in the private sector, which poses its own set of challenges. Just over half a century later, the concept of machine thought is rapidly approaching reality. Recent advances in the field demonstrate new AI capabilities, such as sight, speech, and even moral/ethical reasoning [5-7]. These advances all show a common theme that machines are able to learn, adapting or amending their predetermined programming based on new input. The ability for machines to learn can be incorporated into weaponry and other applications used within DoD to support DoD strategy. AI is redefining the way wars are fought and won, as has been the case with many emerging technologies throughout history.

Since the 1970s, the bulk of DoD efforts regarding AI have focused on applied research [4]. Accordingly, this report focuses on current and potential applications of AI—not basic research. In many ways, this narrower focus makes it easier to describe the proposed capabilities and current capacity of AI. Additionally, new research is published almost daily, making it challenging to stay up-to-date on essential algorithms, coding principles, chip processing speeds, and other pertinent information needed to develop AI.

This report focuses on applications of AI that are relevant to DoD and other governmental agencies that share similar goals, such as the U.S. Department of Homeland Security (DHS) and the Intelligence Community. HDIAC collaborated with subject matter experts to obtain interpretations on relevant applications of AI to DoD and other agencies that could similarly incorporate AI into their operations. Discussions narrowed on relevant applications that have been in use for the past three years as well as advances that could become commercialized within the next 18–24 months. Given the vast amount of information, this report is not all-inclusive and is only a survey of prominent developments.

Although AI is often thought of as a part of computer science and typically viewed from a cyber perspective, applications of AI are relevant to all eight HDIAC focus areas. This report covers applications to Alternative Energy; Biometrics; Chemical, Biological, Radiological, and Nuclear (CBRN) Defense; Critical Infrastructure Protection (CIP); Cultural Studies; Homeland Defense and Security (HDS); Medicine; and Weapons of Mass Destruction (WMD). In some cases, the discussions group together areas that either overlap well or where information on AI applications is thin, and it made more sense to integrate relevant areas in order to provide a more robust perspective.

## References

1. Lange, K. (2016, March 30). 3rd Offset Strategy 101: What it is, what the tech focuses are. DoDLive. Retrieved from <http://www.dodlive.mil/2016/03/30/3rd-offset-strategy-101-what-it-is-what-the-tech-focuses-are/>
2. Turing, A. M. (1950). I.—Computing machinery and intelligence. *Mind, LIX*(236), 433-460. doi:10.1093/mind/lix.236.433
3. Miller, W., Mckenna, T., & Lau, C. (1992). Office of Naval Research contributions to neural networks and signal processing in oceanic engineering. *IEEE Journal of Oceanic Engineering, 17*(4), 299-307. doi:10.1109/48.180298
4. National Research Council. (1999). Funding a revolution: Government support for computing research. Washington, DC: The National Academies Press. doi:10.17226/6323
5. Metz, C. (2017, June 02). The unsettling performance that showed the world through AI's eyes. Retrieved from <https://www.wired.com/2017/04/unsettling-performance-showed-world-ais-eyes/>
6. DepMind. (n.d.). WaveNet: A generative model for raw audio. Retrieved from <https://deepmind.com/blog/wavenet-generative-model-raw-audio/>
7. Vicky Charisi, V., Dennis, L., Fisher, M., Lieck, R., Matthias, A., Slavkovik, M., Loh (Sombetzki), J., Winfield, F.T., & Yampolsky, R. (2017). Towards moral autonomous systems. Retrieved from <https://arxiv.org/abs/1703.04741>

## 4

## Chemical, Biological, Radiological, and Nuclear Defense/Weapons of Mass Destruction

*Joel Hewett and Gregory Nichols, HDIAC*

U.S. warfighters have already deployed AI and machine learning technologies on the battlefield—or, rather, approximately 10,000 feet above it. In December 2017, DoD’s newly-commissioned Algorithmic Warfare Cross-Functional Team successfully deployed an AI-enabled tool in an operational theater for the first time in American military history. The effort, known as Project Maven, applied an object recognition algorithm to the live video feed of unmanned aircraft systems engaged in surveillance during Operation Inherent Resolve to analyze the drone’s reams of footage to generate actionable intelligence [1].

AI promises to drastically reduce the amount of time and effort needed to perform “tedious work,” such as “low-level counting activity” and data entry, which will allow for quicker analysis of still-image and video stream datasets to identify major threats [1]. Moreover, prominent defense analysts have highlighted imagery analysis as one of AI’s most “remarkable capabilities” suited for near-term military use [2].

AI-enabled imagery analysis is likely to be used for the detection and monitoring of CBRN and WMD threats and assets. Former Deputy Defense Secretary Robert Work, who established Project Maven, said in November 2015 these programs could soon be used to efficiently monitor worldwide WMD sites, triggering an alert when a “facility is moving from a benign posture to a threatening posture [3].” Such monitoring may also be further tailored to search for specific assets types, such as “a [ballistic missile] transporter that’s 15 meters long, 4.7 meters wide,” Work offered as an example [3].

While at the time of writing Project Maven is not explicitly focused on CBRN or WMD threats, its status as DoD’s first field deployment of an AI-enabled warfighting tool—and the unique and critical role that imagery analysis plays in the CBRN Defense and Countering WMD missions [4,5]—make it an exemplar of the state-of-the-art in this subfield. The parameters of Project Maven closely reflect how AI and machine learning algorithms are most likely to be applied by DoD to the CBRN/WMD defense mission over the next 5 to 10 years. This general point is supported by two conclusions drawn from HDIAC’s technical analysis of the project, both of which are elaborated on in the next section:

1. AI-enabled imagery analysis software for CBRN/WMD applications is years—if not a decade or more—away from functioning in the “plug-and-play” mode desired by civilian as well as uniformed DoD leaders.
2. Future deployments of AI-enabled algorithms for imagery analysis are all but certain to grow increasingly processor-intensive, suggesting that machine learning engineers must codesign AI software and supporting hardware *simultaneously* in order to deliver workable applications useful to DoD activities.

## Project Maven

Ultimately, technologically-minded strategic thinkers and military leaders hope to apply the immense computing power of AI-based systems to sensing systems in a modular, self-contained configuration. Indeed, when speaking to *Seapower* magazine in May 2016, CPT Jeffrey J. Czerewko, now Commander, Naval Air Force, U.S. Pacific Fleet, described his vision for Naval AI with exactly that tenor of language. He described ideal field-deployable machine learning systems as “platform and domain agnostic” in their use, capable of being hooked up to any existing sensor system as “little black boxes” whose internal workings will require no specific maintenance or continual monitoring [6].

That vision remains powerful, and—it should be noted—fully viable. However, through the end of 2017, Project Maven has demonstrated the infeasibility of delivering this advanced type of “plug-and-play” capability for American warfighting use in the medium-term.

First, the construction and coding of Project Maven’s foundational object identification algorithm was labor-intensive and slow. Program staff manually identified and labeled more than 150,000 images simply to build the starting dataset from which the AI algorithm would teach itself to analyze drone footage [1]. Continued use of the imagery analysis system beyond its initial deployment date will require the manual cataloging of as many as 750,000 *additional* images into the dataset. This will be performed to:

- optimize the algorithm’s performance to match the distinct attributes of its specific desert environment [1]
- decrease the likelihood of the system committing misidentifications due to noisy or substandard image quality [7]
- defend it from disruptive or adversarial attempts to trick or spoof its computer vision [8]

The same requirements for operational maintenance will apply to AI-enabled systems implemented in the near-term, until more general-purpose technological imagery pre-processing programs can be developed to shorten the period needed to train an AI-enabled machine program to begin teaching itself. Indeed, CNA recently noted that “little attention” has been paid to the “amount and kind of data necessary to train military AI systems” [9].

Second, despite the immense amount of computer processing power available to DoD entities (to tap or directly procure), Project Maven’s AI-enabled analytical code once compiled required the assembly of custom-built IT hardware to support its operation [1]. Even when the AI was training itself on only the base dataset of 150,000 images, its operations-per-second load was so processor- and graphics-intensive that it was deemed unable to run off pre-existing DoD networks, including those that were fully cloud-based [1]. The requirement of bespoke coding and engineering labor in operating an AI-enabled tool *for a single video feed* is not a workable proposition beyond the immediate near-term.

This challenge is a novel one even for Project Maven’s state-of-the-art system architecture. Previous instances of AI-enabled imagery analysis systems, typically

ensconced in the laboratories of R1 academic research institutions and highly-capitalized private enterprises, have simply drawn on brute force processing resources to operate. When Google designed an AI-equipped neural network in 2012 to sort through millions of YouTube videos and images to recognize cat pictures, the search company's team simply assembled an array of 16,000 standard-run, commercially-available central processing units (CPUs) to perform the feat [10].

However, future iterations of the Google AI system (renamed Google Brain shortly thereafter) butted heads against fundamental limitations posed by current computing hardware. During 2016 and 2017, Google engineers turned to an emerging, little-known type of processing chip to boost their network's nameplate processing capacity. Tensor processing units (TPU) promise to enable AI-system-wide performance improvements as much as 30 times greater than previous chips [11]. In mid-2017, Google installed an even more powerful version of the initial TPU model with positive results [12].

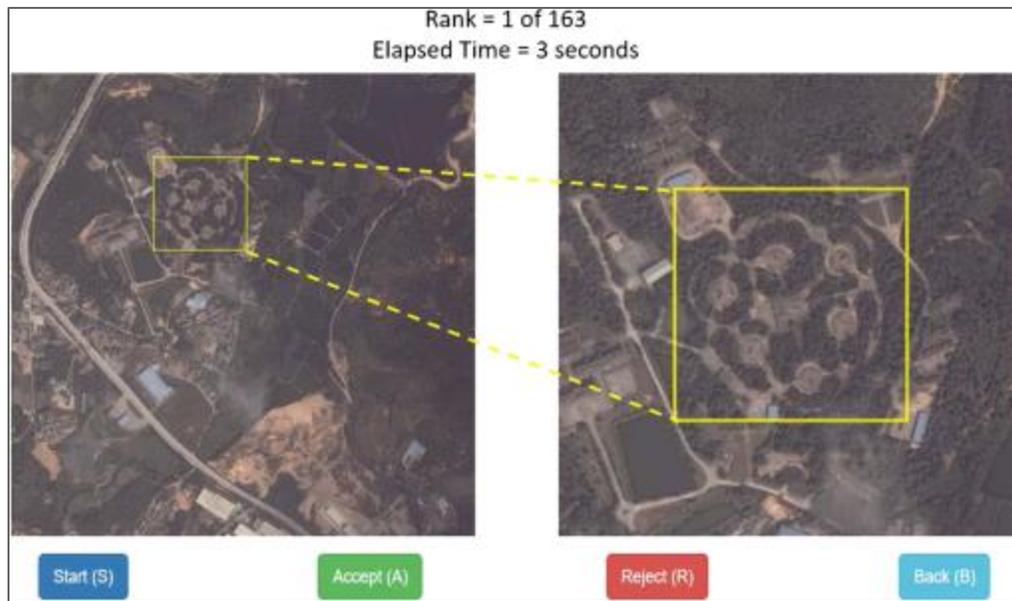
Notably, even given the relative simplicity of Google's unplug-and-replace operation (removing CPU and graphics processing units to make room for the tensor units), the effort was not a simple upgrade job. The TPUs required many months of debugging and software-hardware integration work before they could even function [13].

### **CBRN/WMD Imagery Analysis**

As the 2017 *National Security Strategy of the United States of America* stresses, the detection of WMD assets remains a key priority action for the nation's defense community. The strategy states that the United States will "better integrate intelligence...operations to ensure that frontline defenders have the right information and capabilities to respond to WMD threats" [14]. Elsewhere, the Army has called upon the scientific and technical enterprise of the nation to help provide next-generation "environmental monitoring and threat warning services" for DoD's global ballistic missile defense mission [15]. Therefore, delivering capabilities to achieve this mission has been of high interest to AI researchers and engineers.

In July 2017, researchers at the University of Missouri reported the results of a landmark effort to apply deep convolutional neural networks to the analysis of Chinese surface-to-air missile (SAM) launch sites. (These sites display a distinctive configuration, typically characterized by the presence of four to six small launch pads, set in a broad symmetrical and circular shape [16].)

Drawing upon a training dataset of around 1.6 million open-source images of known SAM sites, the researchers restricted the AI program's access to no more than 100 such positive or affirmative image examples from which it could learn. They then directed the program to conduct a broad-area survey of satellite imagery collected over 55,078 square miles of terrain in southeastern China. The study trained multiple AI-enabled algorithms to auto-select potential candidate sites and prioritize them for imagery personnel to analyze (see Figure 1).



**Figure 1.** Web user interface used for human review of AI-identified and ranked candidate SAM sites. The image on the left presents a contextual view of a candidate site, while the image on the right offers a zoomed-in view of the site, providing the viewer with additional detail for inspection. This candidate was correctly identified as a SAM site [16]. (Released)

By doing so, the system reduced the search area requiring human review to only 0.15 percent of the 55,078-square-mile search area. Moreover, this process, as the authors explain, allowed “four novice imagery analysts with no prior imagery analysis experience” to complete the site search nearly 81 times “faster than a traditional visual search [conducted] over an equivalent land area...while achieving nearly identical statistical accuracy” [16].

As the University of Missouri research effort, Project Maven, and Google Brain make clear, an AI-enabled image analysis algorithm can train itself to detect and identify nearly any type of target, provided that a sufficiently large training dataset exists. In other words: replace the exceptionally small training set of SAM sites (<100 images) with images of chemical weapons production facilities, biological weapons testing sites, or radioactive material extraction fields—and AI-enabled programs will significantly advance the CBRN/WMD defense missions of the United States.

### Emerging Hardware Considerations

As the experience of Project Maven indicates, hardware and processing power access considerations are likely to place limitations on black box-like applications of AI-enabled algorithms for imagery analysis in the medium-term. This issue is apparent in the SAM identification study performed at the University of Missouri: when using the reduced final training set of <100 images, the *total* processing time for the neural network program was roughly 23 hours, even when performed on a high-powered commercial-grade desktop computer station with a high-rate connection to network attached storage [16].

As Google's experience with TPUs makes clear, the path forward to achieving AI-enabled analysis faster and more generalized lies, in part, with the development of hardware specifically designed and optimized for AI use. Multiple leading firms, including Nvidia, Intel, Qualcomm, ARM, and Microsoft are indeed seeking to do just that [17,18]. In the far-term, the usefulness of AI in the analysis of massive datasets may depend on successful yet-to-be-determined developments in quantum computing and/or neuromorphic spiking neural nets [19]. As the Defense Science Board noted in 2016, "technology advances in high-throughput, embedded processing, and machine learning offer the promise of onboard processing of high-resolution, multi-source airborne ISR sensor data" [20].

In order to place AI-enabled imagery analysis at the edge of deployed platforms—or as close to the field as possible—future systems will need to be designed with both hardware and software components in mind. DoD has already taken some steps in this direction [21].

R&D removed from the operational side of AI-enabled systems has recently underscored this future path of development. In December 2017, after conducting a comprehensive review of AI-enabled uses of deep-learning neural networks, a team of engineers and entrepreneurs from MIT and Nvidia cautioned that the continued design of AI systems untethered to specific hardware considerations would reduce their accuracy and efficiency. Instead, systems engineers should codesign algorithm models and hardware contemporaneously, to both "maximize accuracy and throughput, while minimizing energy and cost" [22].

## Conclusion

DoD leaders have stated that the most promising yet untapped datasets collected by their forces are image- and/or video-based. As technology and national security analyst Gregory Allen explains:

Every day, US spy planes and satellites collect more raw data than the Defense Department could analyze even if its whole workforce spent their entire lives on it. [...] Unfortunately, most of the imagery analysis involves tedious work—people look at screens to count cars, individuals, or activities, and then type their counts into a PowerPoint presentation or Excel spreadsheet [1].

"Worse," Allen concludes, "most of the sensor data just disappears—it's never looked at—even though the department has been hiring analysts as fast as it can for years [1]."

Plug-and-play applications of AI in the CBRN/WMD detection space remain years away from operational use. For DoD to tap into the wealth of imagery sensor data it collects every day, future efforts in this subfield must: (a) investigate ways of streamlining algorithm coding and self-training procedures and (b) work in tandem with hardware manufacturers to codesign machine learning hardware-software systems that optimize the function of both.

## References

1. Allen, G. C. (2017, December 21). Project Maven brings AI to the fight against ISIS. *Bulletin of the Atomic Scientists*. Retrieved from <https://thebulletin.org/project-maven-brings-ai-fight-against-isis11374>
2. Allen, G., Chan, T. (2017, July). *Artificial Intelligence and National Security* (Rep.). Cambridge, MA: Harvard University, Harvard Kennedy School, the Belfer Center for Science and International Affairs. Retrieved from <https://www.belfercenter.org/sites/default/files/files/publication/AI%20NatSec%20-%20final.pdf>
3. Weisgerber, M. (2016, April 28). The increasingly automated hunt for mobile missile launchers. DefenseOne. Retrieved from <https://www.defenseone.com/technology/2016/04/increasingly-automated-hunt-mobile-missile-launchers/127864>
4. Gwertzman, B. (1979, October 26). U.S. monitors signs of atom explosion near South Africa. *New York Times*. Retrieved from <https://www.nytimes.com/1979/10/26/archives/us-monitors-signs-of-atom-explosion-near-south-africa-seeks-more.html>
5. Park, J., & Helmboldt, J. (2017). Detection of large scale man-made explosions by the Global Navigation Satellite System. *Journal of the Homeland Defense & Security Information Analysis Center*, 4(2), 24-29. Retrieved from [https://www.hdiac.org/wp-content/uploads/files/journals/Journal\\_Volume\\_4\\_Issue\\_2\\_Final/HDIAC%20Journal\\_Volume%204\\_Issue%202\\_FINAL%20WEB.html](https://www.hdiac.org/wp-content/uploads/files/journals/Journal_Volume_4_Issue_2_Final/HDIAC%20Journal_Volume%204_Issue%202_FINAL%20WEB.html)
6. Matthews, W. (2016, May 18). Navy's Minotaur system is a step toward automated data analysis. *Seapower*. Navy League of the United States. Retrieved from <http://seapowermagazine.org/stories/20160518-data.html>
7. Hosseini, H., Ziao, B., & Poovendran, R. (2017, July 20). *Google's cloud vision API is not robust to noise*. Paper presented at the 16th IEEE International Conference on Machine Learning and Applications, Cancun, Mexico. Retrieved from <https://arxiv.org/pdf/1704.05051.pdf>
8. Kurakin, A., Goodfellow, I. J., & Bengio, S. (2016, July 8). *Adversarial examples in the physical world* (Tech. Rep.). Google, Inc. Retrieved from <https://arxiv.org/pdf/1607.02533v1.pdf>
9. Iliachinski, A. (2017, October.) *Artificial intelligence & autonomy: Opportunities and challenges* (Rep.). Arlington, VA: CNA. Retrieved from [https://www.cna.org/CNA\\_files/PDF/DIS-2017-U-016388-Final.pdf](https://www.cna.org/CNA_files/PDF/DIS-2017-U-016388-Final.pdf)
10. Markoff, J. (2012, June 25). How many computers to identify a cat? 16,000. *New York Times*. Retrieved from <https://www.nytimes.com/2012/06/26/technology/in-a-big-network-of-computers-evidence-of-machine-learning.html>
11. Sato, K., Young, C., & Patterson, D. (2017, May 12). An in-depth look at Google's first tensor processing unit (TPU) [Web log post]. Google Cloud Big Data and Machine Learning Blog. Retrieved from <https://cloud.google.com/blog/big-data/2017/05/an-in-depth-look-at-googles-first-tensor-processing-unit-tpu>
12. Simonite, T. (2017, July 25). The rise of AI is forcing Google and Microsoft to become chipmakers. *Wired*. Retrieved from <https://www.wired.com/story/the-rise-of-ai-is-forcing-google-and-microsoft-to-become-chipmakers/>

13. Lewis-Kraus, G. (2016, December 14). The great A.I. awakening: How Google used artificial intelligence to transform Google Translate, one of its more popular services—and how machine learning is poised to reinvent computing itself. *New York Times Magazine*. Retrieved from <https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html>
14. The White House. (2017, December). *National Security Strategy of the United States of America*. Retrieved from <https://www.whitehouse.gov/wp-content/uploads/2017/12/NSS-Final-12-18-2017-0905.pdf>
15. U.S. Army. (2017, September 22). *Warfighters' Science and Technology Needs Bulletin*. Fort Eustis, VA: U.S. Army Training and Doctrine Command. Retrieved from [http://www.arcic.army.mil/APP\\_Documents/Army-Warfighters-ST-Needs-Bulletin.pdf](http://www.arcic.army.mil/APP_Documents/Army-Warfighters-ST-Needs-Bulletin.pdf)
16. Marcum, R. A., Davis, C. H., Scott, G. J., & Nivin, T. W. (2017, October 6). Rapid broad area search and detection of Chinese surface-to-air missile sites using deep convolutional neural networks. *Journal of Applied Remote Sensing*, 11(4). doi:10.1117/1.JRS.11.042614
17. Melendez, S. (2016, December 7). Chipmakers are racing to build hardware for artificial intelligence. Fast Company. Retrieved from <https://www.fastcompany.com/3065843/artificial-intelligence-chips-race-vidia-intel>
18. Versace, M. (2017, October 20). Does artificial intelligence require specialized processors? The New Stack. Retrieved from <https://thenewstack.io/ai-hardware-software-dilemma/>
19. Vorhies, W. (2017, November 14). The three-way race to the future of AI. Quantum vs. neuromorphic vs. high performance computing [Web log post]. Data Science Central. Retrieved from <https://www.datasciencecentral.com/profiles/blogs/the-three-way-race-to-the-future-of-ai-quantum-vs-neuromorphic-vs>
20. Defense Science Board. (2016, July). *Report of the Defense Science Board Summer Study on Autonomy*. Department of Defense: Office of the Under Secretary of Defense for Acquisition, Technology and Logistics. Retrieved from <https://www.hsdl.org/?view&did=794641>
21. Weisberger, M. (2017, December 21). The Pentagon's new artificial intelligence is already hunting terrorists. DefenseOne. Retrieved from <https://www.defenseone.com/technology/2017/12/pentagons-new-artificial-intelligence-already-hunting-terrorists/144742/>
22. Sze, V., Chen, Y-H., Yang, T-J., & Emer, J. S. (2017, December). Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12), 2295-2329. doi:10.1109/JPROC.2017.2761740