

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.						
1. REPORT DATE (DD-MM-YYYY) 05/15/2020		2. REPORT TYPE Master's Thesis			3. DATES COVERED (From - To) Aug 2019 - May 2020	
4. TITLE AND SUBTITLE DOD's ARTIFICIAL INTELLIGENCE: DoD's Struggle to Adopt and Integrate Machine Learning and Autonomy Technologies				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Caspers, Matthew, S.				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Joint Forces Staff College Joint Advanced Warfighting School 7800 Hampton Blvd. Norfolk, VA 23511-1702				8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution is unlimited						
13. SUPPLEMENTARY NOTES Not for Commercial Use without the express written permission of the author						
14. ABSTRACT With Great Power Competition now coloring nearly all U.S. national security policymaking, Lee Sedol's defeat by AlphaGo, a machine learning algorithm, heightened the U.S.'s urgency to adopt artificial intelligence (AI) and autonomy technologies for defense. However, the DoD's sluggish adoption of these disruptive technologies cannot be solved by policy pressure and funding alone. Instead, the DoD must accelerate efforts to encourage widespread AI literacy, challenge legacy acquisition practices ill-suited to support software development, and pursue safety policies emphasizing a system-based approach to accident prevention and hazard mitigation.						
15. SUBJECT TERMS artificial intelligence, AI, autonomy, autonomous systems, systems safety						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			Matthew S. Caspers	
UNCLASS	UNCLASS	UNCLASS	Unclassified Unlimited	88	19b. TELEPHONE NUMBER (Include area code) 757-443-6400	

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATE COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33315-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report. e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

**NATIONAL DEFENSE UNIVERSITY
JOINT FORCES STAFF COLLEGE
JOINT ADVANCED WARFIGHTING SCHOOL**



DOD'S ARTIFICIAL INTELLIGENCE:

**DoD's Struggle to Adopt and Integrate Machine Learning and Autonomy
Technologies**

by

Matthew S. Caspers

Lieutenant Colonel, United States Air Force

This work cannot be used for commercial purposes without the express written consent of the author.

Page intentionally left blank.

DOD'S ARTIFICIAL INTELLIGENCE

DoD's Struggle to Adopt and Integrate Machine Learning and Autonomy Technologies

by

Matthew S. Caspers

Lieutenant Colonel, United States Air Force

A paper submitted to the Faculty of the Joint Advanced Warfighting School in partial satisfaction of the requirements of a Master of Science Degree in Joint Campaign Planning and Strategy. The contents of this paper reflect my own personal views and are not necessarily endorsed by the Joint Forces Staff College or the Department of Defense.

This paper is entirely my own work except as documented in footnotes.

Signature: _____



15 May 2020

Thesis Advisor:


Signature: _____



**Bryon Greenwald, Ph.D., Professor
Colonel (Ret), U.S. Army**


Approved by:

Signature: _____




**Kristian Smith, Colonel, U.S. Army
Committee Member**

Signature: _____



**Jody Owens, Colonel, U.S. Air Force
Committee Member**

Signature: _____



**Miguel L. Peko, Captain, U.S. Navy
Director, Joint Advanced Warfighting
School**

Page Intentionally Left Blank

Abstract

With the long wars in Iraq and Afghanistan distracting the United States for the last two decades, China began chipping away at the U.S.'s conventional overmatch by rapidly developing and fielding a multitude of anti-access/area denial weapons in the Pacific Theater. Meanwhile, Russia exploited the U.S.'s preoccupation to hone its methods for pursuing its global interests while avoiding direct U.S. confrontation. Concurrently in the private sector, the massive proliferation of user-generated data from smart phones and the Internet of Things contributed to a treasure trove of data from which to feed learning algorithm development. With Great Power Competition now coloring nearly all U.S. national security policymaking, Lee Sedol's defeat by AlphaGo, a machine learning algorithm, heightened the U.S.'s urgency to adopt artificial intelligence (AI) and autonomy technologies for defense. However, the DoD's sluggish adoption of these disruptive technologies prompted lawmakers to increase Congressional pressure and plow additional funding into research and development. But funding and policy alone will not solve the DoD's AI adoption problem. Instead, the DoD must accelerate efforts to encourage widespread workforce literacy regarding AI and autonomy, challenge legacy acquisition practices ill-suited to support software development, and pursue safety policies emphasizing a systems-based approach to accident prevention and hazard mitigation.

Page Intentionally Left Blank

Dedication

For Admiral Adama and the Final Five

Page Intentionally Left Blank

Acknowledgments

I am grateful to Colonel Douglas Wickert, Lt Col David Hoffman, and Lt Col Daniel Javorsek for their advice, support, and willingness to listen and offer ideas. I owe a special debt of gratitude to Mr. Jeff Turner. Without his help, this paper would be unintelligible and unreadable. And finally, I am especially grateful to my wife and children, who patiently waited for me to finish this research and suffered through my thinking out loud.

Page Intentionally Left Blank

Table of Contents

Chapter 1: Introduction	1
Chapter 2: Why AI?	7
Strategy and Policy	7
Disappointment and Achievement.....	10
Chapter 3: Developing the Foundation of AI Literacy	14
Why Literacy?.....	14
Lexicon	15
Autonomy and AI	17
Narrow AI Versus General AI	21
Modern AI Family Tree.....	22
Chapter 4: Enabling Acquisition Reforms.....	28
Never Finished - Agile & DevSecOps.....	30
Enterprise Data	36
Sustainment.....	39
Intelligence.....	40
Chapter 5: Safety-Critical AI Development	43
Humans Versus Automated Systems.....	44
Misplaced Trust	46
A Systems-Based Approach	48
Reliability Versus Safety	50
Software System Safety	54
Chapter 6: Conclusion.....	60
Bibliography	67

Chapter 1: Introduction

At the start of the 21st century, the United States continued to enjoy overwhelming conventional military supremacy, first demonstrated in the 1991 Gulf War and again during Operation Iraqi Freedom in 2003. Since the September 11, 2001 terrorist attacks and for most of the 21st century, the United States' national security strategy focused almost exclusively on defeating transnational terrorism and winning the wars in Iraq and Afghanistan. But as the U.S. remained focused on its long wars in the Middle East, China's military modernization began to chip away at U.S. military advantages by fielding anti-access and area-denial weapons systems capable of denying U.S. operational reach.¹ Similarly, Russia began revitalizing its armed forces and renewing its methods and tactics for competing below the threshold of armed conflict.² In response, then Secretary of Defense Chuck Hagel unveiled the Third Offset Strategy in 2014, a framework emphasizing affordable, distributed, and resilient approaches derived from commercial sources that leverage autonomy and artificial intelligence (AI) to deliver decisive effects to offset U.S. adversaries' advantages.³

Over the last decade, artificial intelligence research and development directed toward national security focused on small projects and investments intended to seed future defense innovations by adapting commercial approaches to spur advancement in

¹ U.S. Department of Defense, *National Defense Strategy* (Washington DC: Government Printing Office, June 2008), 3, <https://archive.defense.gov/pubs/2008NationalDefenseStrategy.pdf> (accessed February 2, 2020).

² U.S. Department of Defense, *National Defense Strategy*, 4, 10.

³ U.S. Department of Defense, *Secretary of Defense Speech: Reagan National Defense Forum Keynote* (Washington DC: Government Printing Office, 15 November 2014), <https://www.defense.gov/Newsroom/Speeches/Speech/Article/606635/> (accessed February 2, 2020).

military-specific domains.⁴ Autonomy and AI portfolios featured prominently in the newly formed Defense Innovation Unit, an organization charged with fostering new DoD partnerships with Silicon Valley to inject cutting-edge information technology into the DoD's shrinking technology base. Following China's 2017 declaration to lead global AI development by 2030, U.S. government pressure to accelerate AI research intensified.⁵ The National Defense Authorization Act (NDAA) of 2019 directed specific AI research and reporting for national security applications.⁶ President Trump followed by signing Executive Order 13859 that directed government support and coordination for the American AI Initiative, a whole of government approach to promote and protect academic and commercial investment in mathematics, science, and engineering disciplines that enable AI advancement.⁷ Despite extraordinary progress in commercial AI products, the flood of ambitious policy initiatives seeking to rapidly apply AI to national security solutions risks another wave of hype, inflated expectations, and the renewed potential for disappointment.⁸

⁴ Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence* (Santa Monica: RAND Corporation, 2019), 106, 129; Interview with a Defense Innovation Unit employee, December 28, 2019.

⁵ Tate Nurkin and Stephen Rodriguez, *A Candle in the Dark: U.S. National Security Strategy for Artificial Intelligence* (Washington DC: Atlantic Council, 10 December 2019) 6; Kelley M. Sayler, *Artificial Intelligence and National Security*, CRS Report No. R45178 Version 7 (Washington DC: Congressional Research Service, 2019), ii, 5-8.

⁶ *John S. McCain National Defense Authorization Act for Fiscal Year 2019*, Public Law 115-232, 115th Congress, 2nd Session (August 13, 2018), § 238, 61-64, <https://www.congress.gov/115/plaws/publ232/PLAW-115publ232.pdf> (accessed January 23, 2020).

⁷ Executive Order 13859, "Maintaining American Leadership in Artificial Intelligence," *Code of Federal Regulations*, title 3 (February 11, 2019), 1-6, <https://www.govinfo.gov/content/pkg/FR-2019-02-14/pdf/2019-02544.pdf> (accessed February 1, 2020).

⁸ Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence*, 22, 32, 127, 130; Stoney Trent and Scott Lathrop, "A Primer on Artificial Intelligence for Military Leaders," *Small Wars Journal*, <https://smallwarsjournal.com/jrnl/art/primer-artificial-intelligence-military-leaders> (accessed February 25, 2020).

Some characterize the intensity of government-directed AI research and development as a new arms race between the U.S. and China.⁹ Despite the increased interest, however, U.S. adoption of AI-enabled technology currently lags expectations. The RAND Corporation published a study in November 2019, sponsored by the Joint Artificial Intelligence Center (JAIC) as directed by the 2019 NDAA, assessing the DoD's posture and progress toward AI implementation.¹⁰ The RAND study highlighted several systemic challenges and barriers slowing DoD advancement and assessed the DoD as largely unprepared and poorly postured to take advantage of the commercial advances in AI and autonomy.¹¹ Google's former CEO and chairman of the National Security Commission on AI, Eric Schmidt, commented that the DoD "does not have an innovation problem; it has an innovation adoption problem."¹²

Part of the adoption problem stems from a lack of AI literacy across a broad cross-section of the DoD. Outside of science and engineering career fields dedicated to DoD's AI and autonomy research, the Department lacks proficiency in depth with AI concepts, which often translates into cynical skepticism that undercuts trust in

⁹ Franz-Stefan Gady, "Elsa B. Kania on Artificial Intelligence and Great Power Competition: On AI's Potential, Military Uses, and the Fallacy of an AI Arms Race," *The Diplomat*, December 31, 2019, <https://thediplomat.com/2020/01/elsa-b-kania-on-artificial-intelligence-and-great-power-competition/> (accessed January 1, 2020); Heather M. Roff, "The Frame Problem: The AI 'arms race' Isn't One," *Bulletin of the Atomic Scientists* 75, no. 3 (2019): 95-98; Paul Scharre, "Killer Apps: The Real Dangers of an AI Arms Race," *Foreign Affairs* 98, no. 3 (May/June 2019): 135-138.

¹⁰ Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence*, xi-xiii, 2, 7, 9.

¹¹ *Ibid.*, xii-xiii.

¹² Michael C. Horowitz and Lauren Kahn, "The AI Literacy Gap Hobbles American Officialdom," *War on the Rocks*, January 14, 2020, <https://warontherocks.com/2020/01/the-ai-literacy-gap-hobbling-american-officialdom/> (accessed January 22, 2020); House Armed Services Committee, "Statement of Dr. Eric Schmidt," *Promoting DoD's Culture of Innovation*, April 17, 2018, <https://docs.house.gov/meetings/AS/AS00/20180417/108132/HHRG-115-AS00-Wstate-SchmidtE-20180417.pdf> (accessed May 2, 2020).

autonomous systems.¹³ The currently widening gap between oversimplified popular articles and technical journals detailing an increasingly sophisticated technology exacerbates intellectual inaccessibility and presents obstacles to adoption. Agreeing on a widely accepted definition for AI remains a principal barrier to widespread literacy.¹⁴ The lack of general knowledge about AI constrains the DoD's ability to meaningfully explain its future concepts for employing AI and autonomous systems, a practice that could help dispel myths while demonstrating realistic capabilities and limitations of these emerging technologies.

AI also aggravates existing shortcomings and introduces new complications to legacy DoD acquisition policies, processes, and procedures. Transforming AI concepts from theory into practical tools requires effective software development, a process declared broken by the Defense Innovation Board's 2019 Software Acquisition and Practices study.¹⁵ Much of the existing Defense Acquisition System emphasizes a linear, waterfall acquisition approach to create a thorough set of weapon system requirements to mature technologies, reduce design risk, and improve production and deployment affordability.¹⁶ While well proven systems engineering principles, these processes

¹³ Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence*, 52-54; National Science and Technology Council, Select Committee on Artificial Intelligence, *The National Artificial Intelligence Research and Development Strategic Plan: 2019 Update*, (Washington DC: Government Printing Office, June 2019), 23-28, <https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf> (accessed February 5, 2020); National Security Commission on Artificial Intelligence, *Interim Report* (Washington DC: Government Printing Office, November 2019), 22, <https://www.nscail.gov/reports> (accessed February 7, 2020).

¹⁴ Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence*, 21-22, 147-153.

¹⁵ U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage* (Washington DC: Government Printing Office, May 3, 2019), i, https://media.defense.gov/2019/Apr/30/2002124828/-1/-1/0/SOFTWAREISNEVERDONE_REFACTORINGTHEACQUISITIONCODEFORCOMPETITIVEADVANTAGE_FINAL.SWAP.REPORT.PDF (accessed March 26, 2020).

¹⁶ U.S. Joint Chiefs of Staff, *Charter of the Joint Requirements Oversight Council (JROC) and Implementation of the Joint Capabilities Integration and Development System (JCIDS)*, CJCSI 5123.01H

remain geared toward minimizing development and procurement risks for hardware-intensive systems characterized by large capital investments.¹⁷ Even absent AI complications, the iterative nature of software development conforms poorly to the DoD's many linear processes.¹⁸

As the pace of AI research continues to accelerate, bringing increasingly capable AI into daily usage, the imperative for establishing norms and principles that guide the safe application of AI and autonomy increases.¹⁹ Companies seeking a competitive advantage may rush emerging technologies into the marketplace without adequate protections for consumer safety.²⁰ Similarly, states fearing a strategic disadvantage and tempted by the allure of game-changing technologies might develop and field immature weapons systems without sufficient safety controls, endangering both service members and the public.²¹ With AI already crowned as the next revolution of military affairs, AI safety standards remain controversial and nascent.²² Furthermore, the growing specter of an AI arms race between the U.S. and China exacerbates pressure to take shortcuts with lax safety protocols.

(Washington DC: Joint Chiefs of Staff, August 31, 2018), D-6, D-14, <https://www.jcs.mil/Portals/36/Documents/Library/Instructions/CJCSI%205123.01H.pdf?ver=2018-10-26-163922-137> (accessed March 1, 2020).

¹⁷ Susanna V. Blume and Molly Parrish, *Make Good Choices, DoD: Optimizing Core Decisionmaking Processes for Great-Power Competition* (Washington DC: Center for a New American Security, November 2019), 5.

¹⁸ Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, viii, xi, 11, 14; National Security Commission on Artificial Intelligence, *Interim Report*, 30-34.

¹⁹ Paul Scharre, "Killer Apps: The Real Dangers of an AI Arms Race," 143-144.

²⁰ *NOVA*, "Look Who's Driving," season 46, episode 19 (originally aired October 23, 2019).

²¹ Paul Scharre, "Killer Apps: The Real Dangers of an AI Arms Race," 140-144; Susanna V. Blume and Molly Parrish, *Make Good Choices, DoD: Optimizing Core Decisionmaking Processes for Great-Power Competition*, 17.

²² Christian Brose, "The New Revolution in Military Affairs: War's Sci-Fi Future," *Foreign Affairs* 98, no. 3 (May/June 2019): 122-124; U.S. Department of Defense, Defense Innovation Board, *AI Principles: Recommendation on the Ethical Use of Artificial Intelligence by the Department of Defense* (Washington DC: Government Printing Office, October 31, 2019), 2-4, https://media.defense.gov/2019/Oct/31/2002204458/-1/-1/0/DIB_AI_PRINCIPLES_PRIMARY_DOCUMENT.PDF (accessed March 13, 2020).

To combat these problems and accelerate AI adoption, the DoD must improve workforce literacy with AI and autonomy technologies to promote buy-in, widen accessibility to concept and tactics development, and democratize vulnerability evaluation. The DoD must also continue to challenge legacy acquisition methodologies and mindsets by embracing recent reforms, empowering program managers, and creatively tailoring acquisition strategies within the bounds of existing statutes and regulations. And finally, the DoD should continue to lead and expand initiatives for the safe application of autonomous and AI systems.

Chapter 2: Why AI?

Strategy and Policy

Attempting to arrest the DoD's technological decay and stagnant investments, then Secretary of Defense Chuck Hagel announced rebalanced priorities for the joint force in the 2014 Quadrennial Defense Review aimed at bolstering the U.S.'s ability to compete in Asia as China's military modernization accelerated and U.S. military advantages eroded.¹ By November of 2014, the DoD announced the Defense Innovation Initiative, a bold department-wide transformation intended to rethink business practices and operations in order to offset adversary advantages and restore sustainable, U.S. military dominance.² The initiative, led by then Deputy Secretary of Defense Robert Work, formed the basis of the Third Offset Strategy, a plan to leverage commercial innovation to develop a suite of new operational concepts that revitalize U.S. conventional deterrence.³ Under the specter of budget sequestration, the strategy recognized that the U.S. could not affordably match adversary competition by pursuing a conventional military build-up of traditional weapons systems that require open-ended budgets and decades to field.

Instead, the strategy sought to offset U.S. adversaries' growing space and cyberspace capabilities and fielded anti-access weapons by initiating an aggressive, technology innovation campaign to quickly fund and transition promising technologies

¹ U.S. Department of Defense, *Quadrennial Defense Review Strategy* (Washington DC: Government Printing Office, 2014), 17, 21, 27,

https://archive.defense.gov/pubs/2014_Quadrennial_Defense_Review.pdf (accessed February 2, 2020).

² U.S. Department of Defense, *Secretary of Defense Speech: Reagan National Defense Forum Keynote*.

³ "Innovation Strategies," National Defense Industrial Association, <https://www.ndia.org/policy/defense-innovation/innovation-strategies> (accessed January 7, 2020); U.S. Department of Defense, *Secretary of Defense Speech: Reagan National Defense Forum Keynote*.

while abandoning languishing concepts and programs.⁴ Emphasizing affordable, distributed, and resilient approaches to new warfighting concepts, Deputy Secretary Work focused the Third Offset on autonomous systems, collaborative human-machine teaming, machine learning, network-enabled weapons, and high-speed projectiles.⁵ The newly formed, Defense Innovation Unit Experiment (DIUx), teamed with industry in Silicon Valley, Austin, and Boston to begin tackling computer vision and machine learning projects to support autonomous systems and artificial intelligence (AI) portfolios.⁶

In May 2016, the Obama administration announced a series of U.S. AI initiatives to foster public dialogue on issues of AI and promote research that fosters government adoption of AI technologies to improve the lives of its citizens.⁷ In October 2016, the National Science and Technology Council (NSTC) released the National AI Strategy, a whole-of-government approach spanning six lines of effort that included security, ethics, safety, education, and economic issues.⁸ Perhaps in response to the U.S. government's increasing interest, China declared its intention to capture global leadership of AI development by 2030 in State Council Document No. 35, published in 2017.⁹ While

⁴ Kathleen H. Hicks et al., *Assessing the Third Offset Strategy* (Washington DC: Center for Strategic & International Studies, 2017), 3.

⁵ Kathleen H. Hicks et al., *Assessing the Third Offset Strategy*, 3.

⁶ "Innovation Strategies," National Defense Industrial Association, <https://www.ndia.org/policy/defense-innovation/innovation-strategies> (accessed January 7, 2020).

⁷ National Science and Technology Council, Networking and Information Technology Research and Development Subcommittee, *The Artificial Intelligence Research and Development Plan* (Washington DC: Government Printing Office, October 2016), v, https://www.nitrd.gov/pubs/national_ai_rd_strategic_plan.pdf (accessed February 5, 2020).

⁸ National Science and Technology Council, Networking and Information Technology Research and Development Subcommittee, *The Artificial Intelligence Research and Development Plan*, 6-22.

⁹ Chinese State Council, "Notice of the State Council Issuing the New Generation of Artificial Intelligence Development Plan," trans. Flora Sapio, Weiming Chen, and Adrian Lo, (Washington DC: Foundation for Law & International Affairs, 2017), 6, <https://flia.org/wp-content/uploads/2017/07/A-New-Generation-of-Artificial-Intelligence-Development-Plan-1.pdf> (accessed February 3, 2020).

China's AI development plan reflects many areas of common interest between U.S. and Chinese entities, some of China's proposed applications imply improved government surveillance in exchange for convenience.¹⁰ The plan proposes several initiatives aimed at improving AI-enabled judicial services, government administration, and social governance including pilot projects in evidence collection and case analysis.¹¹

With the transition from the Obama Administration to the Trump Administration, the DoD abandoned the Third Offset Strategy narrative, but retained and accelerated the portfolio of technology research focused on autonomy and AI intended to mitigate adversary advantages that contest U.S. capabilities across all domains—air, land, sea, space, and cyberspace—by targeting command and control networks, restricting operational reach with inexpensive missile systems, and intensifying competition below the threshold of armed conflict.¹² In concert with the 2018 NDS, the 2019 NDAA directed the DoD to develop an AI strategy, established the Joint AI Center in June 2018, and formed the National Security Commission for AI to explore whole-of-government approaches to integrate AI into national security.¹³ In his 2018 State of the Union address, President Trump emphasized the importance of AI research and development

¹⁰ Chinese State Council, "Notice of the State Council Issuing the New Generation of Artificial Intelligence Development Plan," 18-20, 28.

¹¹ Ibid, 18-20, 28.

¹² U.S. Department of Defense, *National Defense Strategy*, (Washington DC: Government Printing Office, 2018), 1-3, <https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf> (accessed December 23, 2020).

¹³ Terri Moon Cronk, "DoD Unveils Its Artificial Intelligence Strategy," U.S. Department of Defense, <https://www.defense.gov/Explore/News/Article/Article/1755942/dod-unveils-its-artificial-intelligence-strategy/> (accessed February 5, 2020); National Security Commission on Artificial Intelligence, *Interim Report* (Washington DC: Government Printing Office, November 2019), 4; *John S. McCain National Defense Authorization Act for Fiscal Year 2019*.

and subsequently signed Executive Order 13859 ordering support for the American Artificial Intelligence Initiative, a refresh of the 2016 National AI R&D Strategic Plan.¹⁴

Disappointment and Achievement

AI research, as it is known today, began with the first academic conference on the topic at Dartmouth College in 1956.¹⁵ AI research progressed steadily across a wide range of approaches and methods until expectations for the technology exceeded performance in the late 1980s that contributed to hype, disappointment, and a collapse in funding.¹⁶ At the time, AI researchers developed so-called expert machines that encoded large knowledge bases using elaborate rules-engines to search for solutions. The approach, which relied on software engineers to encode knowledge from domain experts, proved cumbersome and costly to update with new knowledge and logic. The subsequent AI winter, a period when AI research and interest declined, began to thaw around 2010 as artificial neural networks returned to the fore with successes in image recognition problems.¹⁷ By the time of Work's Third Offset strategy, AI interest and funding continued to grow, but awareness remained generally contained to the science and technology communities.

¹⁴ Executive Order 13859, "Maintaining American Leadership in Artificial Intelligence;" National Science and Technology Council, Networking and Information Technology Research and Development Subcommittee, *The Artificial Intelligence Research and Development Plan*, 2.

¹⁵ "J. McCarthy et al., "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," (August 31, 1955), <http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf> (accessed May 4, 2020).

¹⁶ Allen Newell, *Intellectual Issues in the History of Artificial Intelligence* (Pittsburgh: Carnegie-Mellon University, 1982), 5, 11-15, <https://apps.dtic.mil/dtic/tr/fulltext/u2/a125318.pdf> (accessed May 4, 2020); Chris Smith et al., "The History of Artificial Intelligence" (University of Washington History of Computing CSEP590A, December 2006), 1-4, <https://courses.cs.washington.edu/courses/csep590/06au/projects/history-ai.pdf> (accessed May 4, 2020).

¹⁷ National Security Commission on Artificial Intelligence, *Interim Report*, 9; Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence* (Santa Monica: RAND Corporation, 2019), 31.

Then, in March of 2016, DeepMind’s AlphaGo, a deep learning algorithm, bested Go world champion, Nine-Dan Grandmaster Lee Sedol, in four out of five games in a globally broadcast tournament.¹⁸ While perhaps less startling to scientists and researchers in the field of deep learning, the journals *Nature* and *Science* both heralded the win as a definitive breakthrough for AI, defying expectations that machine dominance of Go remained at least a decade away.¹⁹ For policy makers outside of the machine learning field, the win served as a Sputnik moment.²⁰ Played for thousands of years in Asia, Go, known as Wei-Chi in China, forms a core aspect of human identity.²¹ Unlike Chess, computational brute force approaches alone cannot defeat a professional human Wei-Chi player since the number of possible moves, $\sim 10^{56}$, exceeds the capabilities of today’s fastest supercomputers.²² Unsurprisingly, AlphaGo’s mastery of Wei-Chi, previously considered a uniquely human endeavor introduced and energized new companies, universities, and government policy makers to the swiftly growing field of AI. Three years after AlphaGo’s first victory, investment and expectations continued to climb, even as the first signs of academic caution began to appear. In his year-end

¹⁸ *AlphaGo*, directed by Greg Kohs, Moxie Pictures & Reel As Dirt, 2017, 1:31, <https://www.youtube.com/watch?v=WXuK6gekU1Y> (accessed January 26, 2020)

¹⁹ David Silver et al., “Mastering the Game of Go with Deep Neural Networks and Tree Search,” *Nature* 529 (January 28, 2016): 484-489; Science News Staff, “From AI to Protein Folding: Our Breakthrough Runners-Up,” *Science*, December 22, 2016, <https://www.sciencemag.org/news/2016/12/ai-protein-folding-our-breakthrough-runners> (accessed February 5, 2020).

²⁰ Kelley M. Sayler, *Artificial Intelligence and National Security*, CRS Report No. R45178 Version 7 (Washington DC: Congressional Research Service, 2019), 5; Georgia Perry, “The AI Cold War That Threatens Us All,” *Wired*, October 23, 2018, <https://www.wired.com/story/ai-cold-war-china-could-doom-us-all/> (accessed February 5, 2020); Henry A. Kissinger, “How the Enlightenment Ends,” *The Atlantic*, June 2018, <https://www.theatlantic.com/magazine/archive/2018/06/henry-kissinger-ai-could-mean-the-end-of-human-history/559124/> (accessed February 5, 2020); Graham Allison, “Is China Beating America to AI Supremacy?” *The National Interest*, December 22, 2019, <https://nationalinterest.org/feature/china-beating-america-ai-supremacy-106861> (accessed December 27, 2020).

²¹ *AlphaGo*, directed by Greg Kohs, Moxie Pictures & Reel As Dirt, 2017, 1:31, <https://www.youtube.com/watch?v=WXuK6gekU1Y> (accessed January 26, 2020)

²² David Silver et al., “Mastering the Game of Go with Deep Neural Networks and Tree Search,” *Nature* 529 (January 28, 2016): 484; Paul Scharre, *Army of None: Autonomous Weapons and the Future of War* (New York: W.W. Norton & Company, Inc., 2018), 125.

summary lecture, Lex Fridman, an MIT autonomous vehicle researcher, declared 2019 the year that AI criticism returned to the mainstream grabbing headlines reporting AI failures including two fatal Tesla autopilot accidents, Amazon's automated resume screening biases, and Microsoft's racist chatbots.²³

Notwithstanding the public failures, national security policy makers continue to accelerate funding for AI research in recognition of the fact that AI offers substantial opportunities for both coalition and adversary militaries alike. Future operating concepts designed to restore U.S. military advantages call for distributed weapons systems that leverage high levels of autonomy and machine advantages in speed and persistence to enable fast-paced, long-range lethal fires in heavily contested environments.²⁴ But imagining highly autonomous, lethal fires capable of operating in high-risk, anti-access environments is the most obvious of AI applications.

AI advancements are also poised to reshape John Boyd's Observe-Orient-Decide-Act (OODA) Loop. Reimagining the OODA Loop with AI, a common and appropriate avenue of contemplation, is often fundamentally misunderstood. At first glance, most

²³ Lex Fridman, "Deep Learning State of the Art (2020)," MIT Deep Learning Series, January 10, 2020, <https://www.youtube.com/watch?v=0VH1Lim8gL8> (accessed February 4, 2020); Jeffrey Dastin, "Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women," *Reuters*, October 9, 2018, <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G> (accessed May 4, 2020); Oscar Schwartz, "In 2016, Microsoft's Racist Revealed the Dangers of Online Conversation," *IEEE Spectrum*, November 25, 2019, <https://spectrum.ieee.org/tech-talk/artificial-intelligence/machine-learning/in-2016-microsofts-racist-chatbot-revealed-the-dangers-of-online-conversation> (accessed May 4, 2020); NTSB, *Collision Between a Sport Utility Vehicle Operating With Partial Driving Automation and a Crash Attenuator, Mountain View, California, March 23, 2018*, NTSB/HAR-20/01 (Washington DC: Government Printing Office, February 25, 2020), <https://www.nts.gov/investigations/AccidentReports/Reports/HAR2001.pdf> (accessed May 4, 2020); NTSB, *Collision Between a Car Operating With Automated Vehicle Control Systems and a Tractor-Semitrailer Truck Near Williston, Florida, May 7, 2016*, NTSB/HAR-17/02 (Washington DC: Government Printing Office, September 12, 2017), <https://www.nts.gov/investigations/AccidentReports/Reports/HAR1702.pdf> (accessed May 4, 2020);

²⁴ U.S. Joint Chiefs of Staff, *Capstone Concept for Joint Operations Joint Force 2030 (Unclassified)* (Washington DC: Joint Chiefs of Staff, June 18, 2019), 6-9.

believe that victory comes to those military commanders able to run through the OODA Loop faster than their opponent.²⁵ Those people incorrectly interpret the OODA Loop. The power of the OODA Loop is only unlocked for those who appreciate that the operational art of paralyzing and dislocating the enemy occurs not by simply making decisions faster than the enemy, but by anticipating the enemy's actions, distorting the enemy's reality, and causing the collapse of the enemy's OODA Loop through friendly decisive action.²⁶ When appropriately and safely applied, AI offers commanders a new means of improving the OODA Loop. Given those consequences, militaries that fail to embrace AI cannot risk conflict with militaries transformed by AI.²⁷ AI offers the potential to disrupt the enemy's decision-making process and achieve the pinnacle of warfighting: winning without fighting.²⁸

²⁵ Frans Osinga, *Science, Strategy, and War: The Strategic Theory of John Boyd* (New York: Routledge, 2007), 272-278.

²⁶ Frans Osinga, *Science, Strategy, and War: The Strategic Theory of John Boyd*, 272-278.

²⁷ Dr. Stephen K. Rogers, interview by author, January 23, 2020.

²⁸ Samuel B. Griffith, *Sun Tzu: The Art of War* (London: Oxford University Press, 1963), 78.

Chapter 3: Developing the Foundation of AI Literacy

Why Literacy?

The DoD remains poorly postured to adopt artificial intelligence (AI) technologies and benefit from AI's enabling characteristics in part due to a lack of widespread AI literacy across its workforce. Presently, islands filled with knowledgeable AI experts exist, though scattered across the DoD in places like DARPA, DIU, the JAIC, and the Services' research labs.¹ Aside from the proliferation of smart products enabled by machine learning algorithms, most of the DoD's workforce remains AI illiterate and unable to access these growing islands of AI expertise. Regardless, much of the DoD workforce and the US public remain uninformed about the highly technical scientific research reports or misinformed by the speculative predictions alerting society to the dangers of super-intelligent, lethal robots.²

When encountering disruptive technologies, workforce literacy promotes buy-in, dispels myths, and facilitates organizational adoption.³ Increasing literacy guards against hype by enabling users and managers to decide for themselves whether a particular AI solution represents a new opportunity capable of offering a decisive combat advantage or a new vulnerability for the joint force. By democratizing access to AI concepts and

¹ National Security Commission on Artificial Intelligence, *Interim Report* (Washington DC: Government Printing Office, November 2019), 22, 31, <https://www.nscai.gov/reports> (accessed February 7, 2020).

² U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, by Dr. Greg L. Zacharias (Maxwell AFB: Air University Press, March 2019), 82; Sydney J. Freedberg, Jr., "Should We Ban 'Killer Robots'? Can We?" *Breaking Defense*, March 11, 2019, <https://breakingdefense.com/2019/03/should-we-ban-killer-robots-can-we/> (accessed February 5, 2020).

³ Tatiana Sanches et al., "Education and Psychology Trends: Impact on Information Literacy," In *Information Literacy: Progress, Trends and Challenges*, ed. Luis Freeman (New York: Nova Science Publishers, 2018), 1-16; Emmett Lombard and Vishal Arghode, "Information Literacy and Organizational Theory," In *Information Literacy: Progress, Trends and Challenges*, ed. Luis Freeman (New York: Nova Science Publishers, 2018), 114-120; "What Exactly Is Information Literacy And What Role Does It Play In Education," USC Marshall School of Business, <https://librarysciencedegree.usc.edu/blog/what-exactly-is-information-literacy-and-what-role-does-it-play-in-education/> (accessed May 4, 2020).

tailoring workforce education, the DoD can bridge the widening gulf between expert AI islands and the general workforce.⁴

Founding lead of Google Brain, Andrew Ng tells companies interested in adopting AI to emphasize widespread workforce education.⁵ While the need for workforce development features prominently in the DoD's 2018 AI strategy, the current focus appears centered on attracting and retaining expert talent onto the DoD's AI expert islands rather than improving widespread literacy.⁶ In order to transform the DoD from a military with AI capabilities into an organization transformed by AI, the DoD must invest in workforce literacy.⁷ Failing to do so risks the growth of calcified skepticism, blind trust, or magical thinking that contributes to AI hype and incoherent investments capable of derailing DoD's AI adoption. The foundation for building workforce literacy begins with developing a common lexicon and plain language explanation. Yet, some of the challenges attributable to the development of a common lexicon stem from differences within the AI fields.⁸

Lexicon

Building a bridge for the Joint Force to access the AI expert islands begins with a common lexicon that avoids Service and career field jargon and biases. One challenge to

⁴ Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* (New York: Basic Books, 2015), xvi.

⁵ Andrew Ng, "AI Transformation Playbook: How to Lead Your Company into the AI Era," Landing AI, https://d6hi0znd7umn4.cloudfront.net/content/uploads/2019/11/LandingAI_Transformation_Playbook_11-19.pdf (accessed January 30, 2020), 3.

⁶ U.S. Department of Defense, *Summary of the 2018 Department of Defense Artificial Intelligence Strategy* (Washington DC: Government Printing Office, 2018), 12-14, <https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMARY-OF-DOD-AI-STRATEGY.PDF> (accessed November 7, 2019).

⁷ Andrew Ng, "AI Transformation Playbook: How to Lead Your Company into the AI Era," 5; National Security Commission on Artificial Intelligence, *Interim Report* (Washington DC: Government Printing Office, November 2019), 36-39, <https://www.nscai.gov/reports> (accessed February 7, 2020).

⁸ Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*, 46.

developing literacy stems from the lack of consensus regarding the definition of artificial intelligence; many researchers define AI in terms of the type of algorithm.⁹ Outside technical journals, AI definitions frequently compare machines' abilities to human intelligence to form the basis for defining AI. A side effect of this approach contributes to AI hype and misunderstanding as machine capabilities relative to human intelligence constantly improves.¹⁰ The *DoD's AI strategy*, the *National Security Commission on AI's Interim Report*, and the *National Artificial Intelligence R&D Strategic Plan* all define AI as some variation of the "ability of machines to solve problems and perform tasks that would otherwise require human intelligence."¹¹ RAND found little interest among government, industry, and academia to spend time developing a formal definition of AI despite some interviewees noting that a standardized vernacular within an organization improved team communication.¹² Despite a lack of consensus across researchers, Congress directed the DoD to develop a definition for use within the department, and seeded the discussion with a very broad definition in the FY2019 NDAA by defining artificial intelligence as including the following:

- (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.

⁹ Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence* (Santa Monica: RAND Corporation, 2019), 22, 149-152.

¹⁰ Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence*, 22.

¹¹ U.S. Department of Defense, *Summary of the 2018 Department of Defense Artificial Intelligence Strategy*, 5; National Security Commission on Artificial Intelligence, *Interim Report*, 7; National Science and Technology Council, Select Committee on Artificial Intelligence, *The National Artificial Intelligence Research and Development Strategic Plan: 2019 Update*, (Washington DC: Government Printing Office, June 2019), iv, <https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf> (accessed February 5, 2020).

¹² Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence*, 21-22.

- (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, designed to approximate a human 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.¹³

One of the most succinct definitions for AI, provided by the Air Force Chief Scientist in *Autonomous Horizons*, approaches the definition in the opposite direction by defining AI as “any machine that possesses intelligence, and intelligence is the ability to gather observations, create knowledge, and appropriately apply that knowledge to accomplish tasks.”¹⁴ Part of the confusion over lexicon stems from the explosion of machine learning, a specialized subfield of AI, currently dominating commercial information technology research and development.¹⁵ Recognizing that a straightforward lexicon provides a unifying and accessible foundation for literacy, the DoD should prioritize simplicity and stability as it develops its AI vernacular.

Autonomy and AI

Explanations of artificial intelligence frequently overlook autonomy and automation, specifically automating increasingly complex tasks, as a principle motivation.¹⁶ Similar to the definition of AI, definitions of autonomy and autonomous systems vary between communities carrying different nuances and connotations depending on the context. From the 2016 Defense Science Board study, “autonomy

¹³ John S. McCain National Defense Authorization Act for Fiscal Year 2019, Public Law 115-232, 115th Congress, 2nd Session (August 13, 2018), § 238, 64, <https://www.congress.gov/115/plaws/publ232/PLAW-115publ232.pdf> (accessed January 23, 2020).

¹⁴ U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, 270.

¹⁵ Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*, 8.

¹⁶ U.S. Department of Defense, Defense Science Board, *Summer Study on Autonomy* (Washington DC: Government Printing Office, June 2016), iii, 1, 4-6,

results from delegation of a decision to an authorized entity to take action within specific boundaries.”¹⁷ The study makes a point to differentiate autonomous systems capable of developing and selecting courses of actions based on knowledge from automated systems governed by prescribed rules that inhibit deviation.¹⁸

Paul Scharre, a senior fellow at the Center for New American Security and former Special Assistant to the Under Secretary of Defense for Policy who led the DoD's autonomy working group, offers an introductory model of autonomous concepts organized along three, orthogonal dimensions: task complexity, human supervision, and algorithm sophistication.¹⁹ To illustrate task complexity along the first dimension, Scharre compares simple tasks such as temperature regulation to extremely complex tasks like driverless cars negotiating congested streets.²⁰ Along the second dimension of his model, human supervision, Scharre describes a continuum between semi-autonomy, supervised autonomy, and full autonomy.²¹ During semi-autonomous operation, called human-in-the-loop, machines automatically accomplish prescribed functions, but then stop and wait for additional human input prior to continuing.²² Supervised-autonomy or human-on-the-loop operation allows the machine to sense, decide, and act without any additional human input.²³ When operating with humans-on-the-loop, machines do not stop and wait for human permission to continue, instead humans must monitor and

¹⁷ U.S. Department of Defense, Defense Science Board, *Summer Study on Autonomy*, 4.

¹⁸ Ibid, 4.

¹⁹ Paul Scharre, “Between A Roomba and a Terminator: What is Autonomy?” *War on the Rocks*, February 18, 2015, <https://warontherocks.com/2015/02/between-a-roomba-and-a-terminator-what-is-autonomy/> (accessed January 2, 2020).

²⁰ Paul Scharre, *Army of None: Autonomous Weapons and the Future of War* (New York: W.W. Norton & Company, Inc., 2018), 28.

²¹ Paul Scharre, *Army of None: Autonomous Weapons and the Future of War*, 28-30.

²² Ibid, 29.

²³ Ibid, 29-30.

intervene to change the machine's behavior.²⁴ For example, airline pilots, monitoring an autopilot, but capable of intervention, constitute supervised, human-on-the-loop autonomy in Scharre's model. Fully autonomous human-outside-the-loop operation transfers full control over to the machine in situations where a human is either unable or unwilling to intervene.²⁵ According to Scharre, differentiating between levels of machine autonomy depends upon context, not on the capability of the underlying automation scheme. Scharre illustrates this point using a Roomba vacuum cleaner. In Scharre's model, a human sitting on the couch while a Roomba vacuums the floor constitutes supervised, human-on-the-loop autonomy.²⁶ If, however, the human leaves the house for work while the Roomba continues to vacuum, then the relationship changes to human-out-of-the-loop autonomy even though nothing about the complexity of the task or the intelligence of the robot changed.²⁷

The final dimension of autonomy in Scharre's model refers to the complexity of the automation scheme. The complexity dimension ranges from simple, threshold-based systems on the low end to complex, goal-oriented, self-directed, and adaptive systems on the high end.²⁸ According to Scharre, high task complexity does not necessarily dictate applying an adaptive and complex AI algorithm using human-outside-the-loop automation. As an example, astronauts landed on the moon, an extremely complex task, using only simple, rule-based software programs governing linear feedback control systems executing semi- and supervised autonomous operations.

²⁴ Ibid, 29.

²⁵ Ibid, 30.

²⁶ Ibid, 30.

²⁷ Ibid, 30.

²⁸ Ibid, 31.

The human supervision dimension finds significant commonality with the DoD's policy directive governing the development of autonomous weapons systems and the Society of Automotive Engineers' (SAE) model of autonomy.²⁹ The DoD's policy directive matches Scharre's terminology to differentiate between weapon systems covered by the policy and those exempted based on the human-machine relationship. To extend the analogy, the SAE model defines six discrete levels of autonomy to describe the division of responsibility for the driving task between the driver and the car.³⁰ Defining discrete levels of autonomy may serve as a useful policy shorthand, but such models do little to advance discussion and development of the key attributes of autonomy and assist in the division of cognitive functions between human and machine.³¹ The 2012 Defense Science Board proposed replacing the DoD's autonomy levels with a three-view framework to assist developers specifically allocate cognitive functions to human or machine, recognizing that allocations may vary by mission phase and provide visibility into high-level system trades inherent in the design.³² The Air Force's recent *Autonomous Horizons* built upon the DSB's conclusions and developed three key attributes for autonomous systems: proficiency, trust, and flexibility.³³ *Autonomous Horizons* provides a detailed discussion about the properties of proficiency, the tenets of

²⁹ U.S. Department of Defense, *Autonomy in Weapon Systems*, DoD Directive 3000.09 (Washington DC: Government Printing Office, November 21, 2012 Incorporating Change 1, May 8, 2017), 13-15, <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/300009p.pdf> (accessed January 2, 2020).

³⁰ "Taxonomy and Definition for Terms Related to Driving Automation Systems for On-Road Motor Vehicles," SAE, https://www.sae.org/standards/content/j3016_201806/ (accessed March 26, 2020).

³¹ U.S. Department of Defense, Defense Science Board, *The Role of Autonomy in DoD Systems* (Washington DC: Government Printing Office, July 2012), 3; Paul Scharre, "Between A Roomba and a Terminator: What is Autonomy?"

³² U.S. Department of Defense, Defense Science Board, *The Role of Autonomy in DoD Systems*, 4.

³³ U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, 1-2, 24, 269.

trust, and principals of flexibility.³⁴ The study examines autonomous system flexibility in terms of task, peer, and cognitive flexibility to both generalize and simplify the Scharre model and avoid developing new levels of autonomy.³⁵

Narrow AI Versus General AI

To date, Google, Facebook, Apple, Amazon, IBM, and Microsoft, driven by market forces, lead commercial AI research and development with approaches exceling at product recommendation, targeted advertising, speech recognition, and decision-making assistance.³⁶ So-called narrow AI approaches, which encompass all current AI technologies, exhibit relatively low task flexibility capable of solving only narrowly defined problems and unable to contextualize outputs.³⁷ Unlike narrow AI, definitions for artificial general intelligence (AGI), the ability of machines to perform *all* the same intellectual tasks of humans including reasoning, analogic thinking, and metacognition, continue to use humans as the benchmark.³⁸ Predicting when AGI may become a reality remains a deeply speculative topic. Regardless of when AGI may be realized, incremental advancements in AI will continue to improve the intelligence of machines and challenge human concepts on the limitations of machines. In the case of autonomous weapons, Deputy Secretary Work commented that he would be “extremely careful in

³⁴ Ibid, 269.

³⁵ Ibid, 2.

³⁶ “Overview of Artificial Intelligence,” CRS In Focus, IF10608, October 27, 2017, 1, <https://crsreports.congress.gov/product/pdf/IF/IF10608> (accessed February 7, 2020).

³⁷ National Security Commission on Artificial Intelligence, *Interim Report*, 7; Kelley M. Sayler, *Artificial Intelligence and National Security*, CRS Report No. R45178 Version 7 (Washington DC: Congressional Research Service, 2019), 2, 35; Greg Allen and Taniel Chen, *Artificial Intelligence and National Security* (Cambridge: Belfer Center for Science and International Affairs, 2017), 8; U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, 139, 155; Paul Scharre, *Army of None: Autonomous Weapons and the Future of War*, 6-7, 231.

³⁸ National Security Commission on Artificial Intelligence, *Interim Report*, 7; U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, 57; Paul Scharre, *Army of None: Autonomous Weapons and the Future of War*, 6-7, 231.

trying to put general AI into an autonomous weapon.”³⁹ While not yet capable of producing AGI, many of the technologies that will enable increased weapon system autonomy to accomplish future DoD mission concepts closely align with the hallmarks of AGI, and a system that possesses all three key attributes of proficiency, trust, and flexibility might qualify as AGI. Accordingly, the imperative for improving DoD literacy increases to enable an informed Joint Force capable of meaningful contributions to the evolving policy framework guiding the safe and effective application of AI and autonomous systems.

Modern AI Family Tree

Before the last AI winter in the 1990s, expert systems, also known as rules engines, dominated the AI ecosystem.⁴⁰ These methods required programmers to symbolically encode vast quantities of knowledge into large databases and develop elaborate rule-based algorithms to search for knowledge, such as TurboTax and spell-checking software.⁴¹ The programmers had to update the knowledge base and rule engines to accommodate new and changing environments.⁴² Research, development, and commercial interest tanked as the cost to maintain staffs of domain experts and programmers to maintain the rules engine ballooned.⁴³

In contrast, machine learning, a subfield of AI, focuses on computer software and algorithms able to modify their own structure and output behavior based on new data. In the field of machine learning, programmers do not explicitly define output responses that

³⁹ Paul Scharre, *Army of None: Autonomous Weapons and the Future of War*, 98-99.

⁴⁰ U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, 130.

⁴¹ *Ibid*, 73.

⁴² *Ibid*, 72-73.

⁴³ *Ibid*, 130.

account for all the possible environmental situations. A rapidly growing discipline, machine learning draws heavily upon mathematics, statistics, computer science, and engineering to develop algorithms capable of making accurate predictions using statistical models trained on large datasets.⁴⁴ While the differences between machine learning and statistics remains an active debate, since both disciplines rely on overlapping mathematical underpinnings to interpret and use data, statistics focuses primarily on forming hypotheses and proving relationships between parameters for making inferences.⁴⁵ Machine learning, in contrast, focuses primarily on making predictions, largely avoiding inference, causality relationships, and model interpretability.⁴⁶

Regardless of the technique, all machine learning systems consist of a model representation, output evaluation, and an optimization method.⁴⁷ The representation consists of the model's structure and parameters for encoding knowledge. The evaluation component measures the effectiveness of the model's output, often quantified in terms of an error function. Finally, the optimization method applies an algorithm to adjust the parameters in the model's representation to minimize or maximize the model's evaluation function based on the model's goal. Determining how to adjust the model parameters refers generically to the credit assignment problem, an increasingly complicated problem for larger models with thousands of parameters in its representation.⁴⁸

⁴⁴ Kelley M. Sayler, *Artificial Intelligence and National Security*, 2.

⁴⁵ Danilo Bzdok et al., "Statistics Versus Machine Learning," *Nature Methods* 15, no. 4 (April 2018), 233-234, <https://www.nature.com/articles/nmeth.4642.pdf> (accessed February 5, 2020).

⁴⁶ Danilo Bzdok et al., "Statistics Versus Machine Learning," *Nature Methods*.

⁴⁷ Ryan Hefron, "RDT&E of Autonomous Systems" (U.S. Air Force Test Pilot School Short Course Charts, August 23, 2019); Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*, 240.

⁴⁸ U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, 132; Richard S. Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. (Cambridge: MIT Press,

At a high level, machine learning algorithms, grouped by the method of learning, typically fall into three categories: supervised, unsupervised, and reinforcement learning. For supervised learning, large, labeled datasets in which there exists an observed response for every measured prediction provide the developer a truth source for constructing the evaluation function to quantify the error or accuracy of the model's prediction. Developers divide the dataset into a training dataset and test dataset. During training, the optimization algorithm updates the model's parameters to minimize error based on the training data.⁴⁹ With the model's parameters frozen after training, the developer then checks the model's accuracy using only the test data. Supervised learning methods are typically applied to classification and regression problems, such as identifying objects in an image (classification) or predicting life expectancy based on demographic data (regression).⁵⁰ When labeled data do not exist, the more challenging problem of unsupervised learning cannot train a model using a truth source to fit the parameters. Instead, unsupervised learning uses clustering methods to discern patterns in the data and group unlabeled data into distinct categories.⁵¹ Unsupervised learning is often applied to anomaly detection and dimensionality reduction, a method of compressing the number of model parameters.⁵²

2018), 17; Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*, 101.

⁴⁹ U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, 71-72; Vijay Gadepally et al., *AI Enabling Technologies: A Survey* (Lexington, MA: MIT Lincoln Laboratory, 2019), 16-17, 20, <https://arxiv.org/ftp/arxiv/papers/1905/1905.03592.pdf> (accessed March 13, 2020).

⁵⁰ Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence*, 31; Gareth James et al., *An Introduction to Statistical Learning: With Applications in R*. (New York: Springer, 2017), 1, <https://link.springer.com/content/pdf/10.1007%2F978-1-4614-7138-7.pdf> (accessed February 12, 2020); Vijay Gadepally et al., *AI Enabling Technologies: A Survey*, 16-17.

⁵¹ Vijay Gadepally et al., *AI Enabling Technologies: A Survey*, 16-17.

⁵² Ryan Heffron, "RDT&E of Autonomous Systems," Vijay Gadepally et al., *AI Enabling Technologies: A Survey*, 10, 15, 17.

The last major category of machine learning, reinforcement learning (RL), differs by focusing on goal-oriented approaches to determine the actions of a decision-making agent interacting with its environment. The main components of a reinforcement learning problem are the policy, reward function, and value function.⁵³ The policy maps the agent's immediate behavior by translating the agent's perception of the environment into responses, which could be a simple look up table.⁵⁴ The reward function maps perceived states to desirability; the more desirable the state, the higher the agent's reward.⁵⁵ Value functions determine the long-term desirable state and provide a mechanism for predicting possible rewards in future states.⁵⁶ In order to maximize reward over the long-term, reinforcement learning methods must estimate the value function that predicts states of maximum value.⁵⁷

The explosion in AI research and products over the last decade, which contributed to the success of Netflix, Amazon Alexa, Google Images, and Tesla's autopilot, is largely based on the renewed viability of artificial neural networks (ANNs).⁵⁸ Neural networks, inspired by the interconnected neurons of the human brain, owe their origins to the perceptron, a simple logical operator conceived by two researchers McCullough and Pitts in the 1940s.⁵⁹ While the concept of ANNs existed for decades, the technology experienced modest advancement until the convergence of improved parallel computing

⁵³ Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction*, 6-7.

⁵⁴ *Ibid.*, 7.

⁵⁵ *Ibid.*

⁵⁶ *Ibid.*

⁵⁷ *Ibid.*

⁵⁸ Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence*, 28; U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, 6-7.

⁵⁹ *Ibid.*, 131-132.

and large data repositories generated from increasing network connectivity afforded by the Internet.⁶⁰

Due to the transformative effect ANNs have had on AI over the last decade, gaining a summary understanding of ANNs provides a significant improvement in AI literacy. ANNs refer to a general connectionist architecture for defining a model representation that can be applied in each of the categories (supervised, unsupervised, and RL).⁶¹ Artificial neural networks consist of vast numbers of artificial neurons, virtual nodes that accept multiple, weighted inputs to compute a single output value, arranged into arrays called layers. Generally, the output from each neuron in one layer is connected to each of the neurons' inputs in the next layer, but many variations on the architecture exist based on the specific application.⁶² Network nomenclature typically refers to the first layer as the input layer and the last layer as the output layer and designates the layers in between as hidden layers. Networks called deep learning networks simply have many hidden layers. Gaining literacy with recent advancements in connectionist learning algorithms, such as ANNs, remains a critical aspect of any AI literacy effort, but a robust AI literacy curriculum would be incomplete without inclusion of other learning algorithm techniques. Although beyond the scope of this paper, Pedro Domingos, a computer science professor and machine learning expert, thoroughly

⁶⁰ Ibid, 7; Vijay Gadepally et al., *AI Enabling Technologies: A Survey*, 25-26; National Science and Technology Council, Select Committee on Artificial Intelligence, *The National Artificial Intelligence Research and Development Strategic Plan: 2019 Update*, (Washington DC: Government Printing Office, June 2019), 12-13, <https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf> (accessed February 5, 2020).

⁶¹ U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, 131-132; Jurgen Schmidhuber, *Deep Learning in Neural Networks: An Overview* (Switzerland: University of Lugano, October 8, 2014) <https://arxiv.org/pdf/1404.7828.pdf> (accessed February 10, 2020).

⁶² Jurgen Schmidhuber, *Deep Learning in Neural Networks: An Overview*; Vijay Gadepally et al., *AI Enabling Technologies: A Survey*, 19-22.

explains and explores each of the five major machine learning approaches: symbolists, connectionists, evolutionaries, Bayesians, and analogizers.⁶³

⁶³ Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*, xvii.

Chapter 4: Enabling Acquisition Reforms

The 2018 National Defense Strategy fully reframed the 21st century strategic security environment towards Great Power Competition and prioritized acquisition efforts to build a more lethal force capable of exploiting rapid commercial information technology advances. In order to reap the benefits of artificial intelligence and leverage the speed and agility of commercial innovation, the DoD has challenged the role and importance of legacy systems.¹ The Defense Acquisition System (DAS) produces lethal and capable weapon systems, but its processes remain tethered to a Cold War mindset focused on large-scale, hardware-intensive weapons systems.² The hardware-centric mindset emphasizes generating firm, upfront requirements from warfighters to feed lengthy acquisition programs. While AI and autonomous systems rely on a robust and diverse stack of technologies, including modern computing hardware for advanced processing and data storage, a significant portion of AI advancement depends upon quick and effective software development. The adoption of AI and autonomous systems requires a culture that embraces DAS reforms, permitting rapid and iterative software capability development and deployment instead of defaulting to linear, hardware-centric waterfall processes.³

¹ Paul Mcleary, “SecDef Eyeing Moving Billions By Eliminating Offices, Legacy Systems,” *Breaking Defense*, February 5, 2020, <https://breakingdefense.com/2020/02/secdefs-review-is-in-and-hes-willing-to-shut-down-entire-offices/> (accessed May 1, 2020).

² Susanna V. Blume and Molly Parrish, *Make Good Choices, DoD: Optimizing Core Decisionmaking Processes for Great-Power Competition* (Washington DC: Center for a New American Security, November 2019), 5; National Security Commission on Artificial Intelligence, *Interim Report* (Washington DC: Government Printing Office, November 2019), 22-24, <https://www.nsc.ai.gov/reports> (accessed February 7, 2020).

³ U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage* (Washington DC: Government Printing Office, May 3, 2019), vii-15, https://media.defense.gov/2019/Apr/30/2002124828/-1/-1/0/SOFTWAREISNEVERDONE_REFACTORINGTHEACQUISITIONCODEFORCOMPETITIVEADVANTAGE_FINAL.SWAP.REPORT.PDF (accessed March 26, 2020).

As the U.S. began the war in Afghanistan in 2001, the still-deflating Internet Dot-Com bubble continued to bankrupt information technology companies. Apple's market share shrank, Microsoft fought the government over its alleged monopoly, and most consumers remained wary of Internet shopping. By 2007, the trend reversed, and Facebook, founded in 2004, dominated education-related social media with over 30 million users.⁴ Apple launched the first iPhone in 2007 and accelerated its mobile phone development and sales by joining with AT&T's 3G offering for unlimited data service.⁵ Google's Android OS debuted in 2008 enabling hardware companies, lacking their own software platforms, to rapidly enter the mobile phone market. Soon Android-enabled Samsung and Nokia phones exploded onto the market and contributed to an exponential increase in user-generated data, a new commodity that forms the basis of many commercial AI applications. During approximately the same period of commercial information technology transformation, the U.S. military focused its modernization on struggling, but traditional programs like the F-22, KC-46, F-18E/F, the Littoral Combat Ship, and the Ford Class Carrier. Juxtaposing the DoD's largest acquisition program, the F-35, the subject of constant OSD, Service, and Congressional oversight, with Silicon Valley's industry transformation serves to highlight the disparity in agility between U.S. information technology companies and the traditional U.S. military-industrial base.

Many of the DAS's cumbersome processes owe their roots to massive Cold War programs aimed at matching, exceeding, or offsetting Soviet capabilities.⁶ During the

⁴ Sarah Phillips, "A Brief History of Facebook," *The Guardian*, July 25, 2007, <https://www.theguardian.com/technology/2007/jul/25/media.newmedia> (accessed February 5, 2020).

⁵ Woyke, Elizabeth, *The Smartphone: Anatomy of an Industry* (New York: The New Press, 2014), 146.

⁶ Kathleen H. Hicks et al., *Assessing the Third Offset Strategy* (Washington DC: Center for Strategic & International Studies, 2017), 2.

Cold War, government funded research pushing new technologies into increasingly exotic weapon systems peaked. As the Cold War ended, defense budgets tightened and government research for cutting edge science and high-risk technology maturation dwindled, resulting in more and more groundbreaking technologies originating outside of government contracts.⁷ As the hub of innovation shifted from government labs to commercial and university labs, the DAS failed to adjust and keep pace with the accelerating tempo of information technology and software development in particular. Inertia in the DAS did not go unnoticed, however. As of 2002, RAND's list of publications recommending acquisition reforms topped 63 reports.⁸ The 2019 Defense Innovation Board's software report concluded that the DoD's software development process remains broken despite fifteen previous studies also dedicated to improving DoD software acquisition practices.⁹

Never Finished - Agile & DevSecOps

Developing increasingly sophisticated systems featuring the behaviors and attributes of autonomous systems demands an increasingly flexible software framework from which to host continuously evolving algorithms.¹⁰ Recognizing that successful AI grows from effective software development, the DoD's default hardware mindset, built on linear processes, becomes a major impediment to implementing processes that support the continuous, iterative nature of software development. The Defense Innovation

⁷ Susanna V. Blume and Molly Parrish, *Make Good Choices, DoD: Optimizing Core Decisionmaking Processes for Great-Power Competition*, 9.

⁸ J. Ronald Fox, *Defense Acquisition Reform 1960-2009: An Elusive Goal* (Washington DC: Center of Military History, U.S. Army, 2011), 231.

⁹ U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, i.

¹⁰ National Security Commission on Artificial Intelligence, *Interim Report* (Washington DC: Government Printing Office, November 2019), 22; U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, by Dr. Greg L. Zacharias (Maxwell AFB: Air University Press, March 2019), 269.

Board’s conclusions--that software is never done and an ongoing software project does not equate to poorly defined requirements—calls for the DoD to adopt Agile and DevSecOps approaches to software development.¹¹

Agile software development, born out of a manifesto published in 2001, sought to improve the relationship between the user and the developer by emphasizing a new software development mindset based on frequent deliveries of smaller, functional software releases instead of large-scale monolithic projects. Agile's manifesto emphasizes face-to-face collaboration and engagement with customers over cumbersome documentation and up front planning.¹² Agile differs significantly from the traditional waterfall management process by eschewing linear, serial processes and instead embracing an iterative approach building and adding capability in small, easily testable batches.¹³ Instead of delivering a shrink-wrapped software package after months or years of serial development with little interaction and end user engagement, Agile encourages developers to seek early feedback and embrace customer requirement changes instead of viewing change requests as a worrisome project risk.¹⁴ Agile practices tightly couple design, development, and verification to test functionality as early as possible to avoid large software branches.¹⁵ By delivering a Minimally Viable Product (MVP) for user

¹¹ U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, v-xiii, 6, 9-12.

¹² Caroline Mimbs Nyce, “The Winter Getaway That Turned the Software World Upside Down,” *The Atlantic*, December 8, 2017, <https://www.theatlantic.com/technology/archive/2017/12/agile-manifesto-a-history/547715/> (accessed February 7, 2020); “Manifesto for Agile Software Development,” <https://agilemanifesto.org/> (accessed February 7, 2020).

¹³ Caroline Mimbs Nyce, “The Winter Getaway That Turned the Software World Upside Down,” *The Atlantic*.

¹⁴ U.S. General Accounting Office, *Software Development: Effective Practices and Federal Challenges in Applying Agile Methods* (Washington DC: Government Printing Office, July 2012), 4-6, 8, 26, <https://www.gao.gov/assets/600/593091.pdf> (accessed March 8, 2020).

¹⁵ U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, 10.

demonstration as soon as possible, software developers can rapidly fix, update, and change the software based on early user feedback.¹⁶ Even though the MVP may lack must-have features, the Agile mindset places a premium on direct user feedback to correct deficiencies and adds features using a collaborative approach as both the users' and developers' mutual understanding improves. The resultant cooperative and iterative feedback-design-develop-release framework helps avoid lengthy, end-of-project requirement verification checklists to satisfy contracts. Agile's goal focuses on delivering functioning software not paperwork.

DevSecOps, now considered an industry best practice, refers to the mindset and set of tools, processes, and workflows that seek to unify software development, security, and operations into a continuous, integrated pipeline rather than distinct communities using separate processes.¹⁷ With DevSecOps, security is baked into the software design from the beginning and continuously monitored and improved throughout the software lifecycle instead of using compliance checklists or an overreliance on network scans during operation.¹⁸ Continuous Integration and Continuous Delivery (CI/CD) form two fundamental cornerstones of DevSecOps practices. CI/CD workflows enable software factories to scale up the scope of work across a development project, but avoid long integration, test, and delivery delays typical of traditional approaches.¹⁹ Traditional waterfall requirement processes emphasize thorough and stable requirements definition

¹⁶ U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, 7-11; Section 809 Panel, *Report of the Advisory Panel on Streamlining and Codifying Acquisition Regulations*, 2 (Washington DC: Government Printing Office, June 2018), 57, <https://discover.dtic.mil/section-809-panel/> (accessed February 25, 2020).

¹⁷ U.S. Air Force, "DoD Enterprise DevSecOps Initiative (Software Factory) v4.7," by Nicolas Chaillan, <https://software.af.mil/dsop/documents/> (accessed February 23, 2020).

¹⁸ "Manifesto," <https://www.devsecops.org/> (accessed February 23, 2020).

¹⁹ Martin Fowler, "Continuous Integration," <https://martinfowler.com/articles/continuousIntegration.html> (accessed February 23, 2020).

before beginning a project. After gathering and documenting detailed user requirements, the team writes an in-depth plan to map out the project assigning tasks to specific functional teams. As teams complete their assigned projects, integration teams build the software baseline and verification teams begin system-level testing to identify customer requirement deficiencies and initiate fixes. Major fixes typically generate software forks that serve to exacerbate mainline software build and integration challenges. Attempts to scale legacy software development approaches lacking continuous integration practices frequently suffer significant delays or collapse as the communication and collaboration overhead increases when the team tries to integrate code changes from multiple development teams back into a functioning mainline.²⁰

Continuous delivery extends continuous integration by flowing software as quickly as possible to customers, but preserves a human decision gate prior to deployment.²¹ Continuous deployment closes the loop entirely by leveraging increasing levels of automated testing and verification tools to enable large scale software deployment without human intervention.²² Continuous deployment deletes the notion of a release day; as soon as software passes its automated testing the changes flow to production.²³

Based on the Defense Innovation Board's software report, the DoD recognizes the challenges associated with software development, but changing the DoD's hardware-

²⁰ Max Rehkopf, "What Is Continuous Integration," <https://www.atlassian.com/continuous-delivery/continuous-integration> (accessed May 15, 2020).

²¹ Sten Pittet, "Continuous Integration vs. Continuous Delivery vs. Continuous Deployment," <https://www.atlassian.com/continuous-delivery/principles/continuous-integration-vs-delivery-vs-deployment> (accessed February 23, 2020); U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, 6-7.

²² Ibid, 6.

²³ Sten Pittet, "Continuous Integration vs. Continuous Delivery vs. Continuous Deployment;"

focused culture to a culture embracing Agile and DevSecOps software practices, even as part of hardware intensive programs, remains daunting. The DoD, spurred to action by some of the interim findings of the Section 809 panel, seems eager to continue reforms with the completion of the Defense Innovation Board's 2019 software report.²⁴ By September 2019, the DoD Chief Information Officer (CIO) released the DevSecOps Reference Design initiative aimed at implementing a DevSecOps mindset and proliferating CI/CD pipelines across the DoD.²⁵ In early January 2020, the Under Secretary for Defense Acquisition and Sustainment (USD A&S) released an interim policy aimed at simplifying software acquisition processes and enabling continuous integration and delivery.²⁶ The interim policy delegates the decision to use the software pathway to the component acquisition executive and permits exemption from the Joint Capabilities and Integration Development (JCIDS) process. Instead, the sponsoring organization and the program manager develop a capability need statement and user agreement that maps mission use cases, negotiates trade space for demonstration and test of the MVP, and identifies metrics to support the subsequent deployment of the Minimum Viable Capability Release (MVCR), the first iteration of mission-ready software.²⁷ While the policy does not require or prescribe set timelines for development,

²⁴ Section 809 Panel, *Report of the Advisory Panel on Streamlining and Codifying Acquisition Regulations*, 2; U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*.

²⁵ U.S. Department of Defense, DoD Chief Information Officer, *DoD Enterprise DevSecOps Reference Design Version 1.0* (Washington DC: Government Printing Office, August 12, 2019), 1-3. 15-17, <https://dodcio.defense.gov/Portals/0/Documents/DoD%20Enterprise%20DevSecOps%20Reference%20Design%20v1.0%20Public%20Release.pdf?ver=2019-09-26-115824-583> (accessed February 25, 2020).

²⁶ U.S. Department of Defense, Under Secretary of Defense for Acquisition and Sustainment, "Software Acquisition Pathway Interim Policy and Procedures," January 3, 2020, [https://www.acq.osd.mil/ae/assets/docs/USA002825-19%20Signed%20Memo%20\(Software\).pdf](https://www.acq.osd.mil/ae/assets/docs/USA002825-19%20Signed%20Memo%20(Software).pdf) (accessed February 25, 2020).

²⁷ Ibid.

the software pathway graphic depicts the policy's intent for speed by emphasizing delivery of the MVP using months as the unit for time.

The new software policy seeks to mainstream gains made on software-only projects demonstrated at unique organizations like the Defense Digital Service and the Air Force's Kessel Run.²⁸ The new policy widens the DoD's aperture for improving software acquisition beyond small scale, niche areas and promotes a DevSecOps approach to a wider array of programs, including traditional weapon systems. Shortly after the interim software policy publication, USD A&S also released new policy guidance for the rewrite of DoDI 5000.02, the instruction that governs DoD's acquisition operations.²⁹ The new instruction, retitled, "Operation of the Adaptive Acquisition Framework," outlines the transition plan for the new instruction and consolidates a series of policy reforms that comprise the new Agile Acquisition Framework (AAF).³⁰ The AAF restructures the legacy system into six acquisition pathways: Urgent Capability, Middle Tier, Major Capability, Software, Defense Business Systems, and Services. Under the legacy process, now known as the Major Capability Acquisition pathway, program managers (PMs) were encouraged to tailor acquisition strategies to the unique needs of their program.³¹ Aggressive and creative tailoring, however, often created

²⁸ U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, 31, 57-63.

²⁹ "DoD Rewrite of 5000 Series to Include a Software Acquisition Pathway," National Defense Industrial Association, <https://www.ndia.org/policy/recent-posts/2019/7/26/dod-rewrite-of-5000-series-to-include-a-software-acquisition-pathway>, July 26, 2019 (accessed February 4, 2020).

³⁰ U.S. Department of Defense, Under Secretary of Defense for Acquisition and Sustainment, *Operation of the Defense Acquisition System*, DoD Instruction 5000.02T (Washington DC, Government Printing Office, January 7, 2020, Incorporating Change 6, January 23, 2020), <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500002T.PDF?ver=2020-01-24-100028-310> (accessed February 4, 2020).

³¹ U.S. Department of Defense, Under Secretary of Defense for Acquisition and Sustainment, *Operational of the Adaptive Acquisition Framework*, DoD Instruction 5000.02 (Washington DC, Government Printing Office, January 23, 2020), 9,

significant friction as PMs tried to justify novel approaches that did not conform to traditional practices. Under the AAF, PMs now have a stable of ready-made tools, a common vocabulary, and an outlined set of approval authorities from which to assemble a custom acquisition program.

Enterprise Data

While the AAF's new software policy should certainly improve the DoD's software development efforts and by extension AI and autonomous system algorithm development, the AAF also stands to improve other AI-supporting technologies and systems. To match private industry's success, part of the DoD strategy for adoption includes ingraining a culture that leverages enterprise data across Service and organizational seams. Through the acquisition of services and business systems, enabled by the AAF, the DoD can foster the production and sustainment of well-curated collections of data accessible across the DoD enterprise. By transforming the DoD's current information architecture of disparate computing enclaves into a cloud-based enterprise platform, the DoD will better position itself to take advantage of large-scale data services and tools.

The Joint Enterprise Defense Infrastructure (JEDI) contract award to Microsoft in October 2019 is part of a strategy to transform the DoD's IT infrastructure from a fractured, poorly integrated collection of island enclaves into a modern, distributed cloud storage and computing environment.³² Part of the impetus for such a widespread policy,

<https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500002p.pdf?ver=2020-01-23-144114-093> (accessed February 4, 2020).

³² Andrew Eversden, "So What Problems Does JEDI Solve, Really?" *Federal Times*, October 30, 2019, <https://www.federaltimes.com/govcon/contracting/2019/10/30/so-what-problems-does-jedi-solve-really/> (accessed May 1, 2020).

moving the DoD to a common cloud, was to enable the necessary data pooling and sharing across the Services that will power data-hungry, supervised machine learning algorithms. To bridge the gap until JEDI implementation, Service CIO's are authorizing cloud storage in the interim.³³ Transforming the DoD's IT infrastructure from a locally managed, hardware focused endeavor to a platform as a service architecture ties in with the AAF's pathway for the acquisition of services enabling the DoD to shed duplicative and often ineffective methods of operating its IT enterprise.³⁴ Instead, the JEDI service will contribute to a solid foundation for improved DoD software development and deployment that provides an environment for unifying databases and fostering AI advancement.

As noted in Chapter 3, supervised machine learning methods such as artificial neural networks power industry achievements solving difficult classification and regression problems. Specifically, the success of applications like Alexa, Siri, Google Images, Amazon, and Netflix recommendations depends largely upon a cornucopia of data generated by individual user behaviors interacting with the Internet, smart phones, and the Internet of Things (IoT). Considering that every button click or status message on an Internet-connected device represents an opportunity to collect potential training data for machine learning, it is no surprise that an availability of user data fueled a

³³ U.S. Department of Defense, Chief Information Officer, "Updated Guidance on the Acquisition and Use of Commercial Cloud Computing Services," December 15, 2014, https://dodcio.defense.gov/portals/0/documents/cloud/dod%20cio%20-%20updated%20guidance%20-%20acquisition%20and%20use%20of%20commercial%20cloud%20services_20141215.pdf (accessed February 25, 2020); Department of the Navy, Chief Information Officer, "Acquisition and Use of Commercial Cloud Computing Services," May 15, 2015, <https://www.doncio.navy.mil/TagResults.aspx?ID=104> (accessed February 25, 2020).

³⁴ U.S. Air Force, "DoD Enterprise DevSecOps Initiative (Software Factory) v4.7," by Nicolas Chaillan; U.S. Air Force, Chief Software Officer, "Preferred Agile Framework," December 28, 2019, <https://software.af.mil/wp-content/uploads/2019/12/CSO-MFR-on-Agile-Frameworks-12282019.pdf> (accessed February 25, 2020).

learning algorithm boom across the information technology industry. Companies like Google, Facebook, Apple, Microsoft, and Amazon dominate the commercial AI field largely because they offer free or low-cost convenient services in exchange for user data.³⁵ When a user tags a friend in a Facebook image, the user provided Facebook with clean, high-fidelity training data at no cost. In other applications, acquiring clean, labeled data represents one of the largest barriers to entry for many machine learning applications. Whole industries have sprung up in the commercial sector offering services to dig manually through a company's data, organize, and label it, a unique challenge for the DoD.³⁶

Presently, the DoD lacks the culture, infrastructure, and policy needed to mirror industry's success in leveraging large data repositories and enforcing labeling at the point of production.³⁷ Over the last two decades, the Services pushed large scale efforts to eliminate paper-based processes and data storage in favor of digital methods. While storing personnel records in PDF files and mandating the use of electronic correspondence certainly reduces the DoD's paper consumption, policies that lack data curation enforcement fail to capitalize on the utility of the digital medium.³⁸ Widespread digital and data literacy as well as deep culture change are necessary for the DoD to achieve AI implementation.

³⁵ Rani Molla, "Why Your Free Software Is Never Free," *Vox*, January 29, 2020 (accessed May 15, 2020).

³⁶ "China's Success at AI Has Relied on Good Data," *The Economist*, January 2, 2020, Technology Quarterly, <https://www.economist.com/technology-quarterly/2020/01/02/chinas-success-at-ai-has-relied-on-good-data> (accessed February 4, 2020).

³⁷ Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence* (Santa Monica: RAND Corporation, 2019), 57-60; National Security Commission on Artificial Intelligence, *Interim Report*, 25-28, 34.

³⁸ Interviews with two U.S. Air Force machine learning experts, November 4, 2019 and February 21, 2020.

Sustainment

Identifying clean, labeled data to enable supervised learning methods for classification and prediction applications presents one of the largest barriers to entry. Airlines have already recognized the need for data-driven analytics to minimize aircraft downtime and schedule disruptions while maximizing profits by avoiding costly excessive maintenance and supply chain overhead.³⁹ Recognizing aircraft maintenance as one of the most straightforward areas for direct transfer from industry, the Air Force, Navy, DIU, and JAIC all quickly began working toward leveraging industry advances. In 2018, Air Force Materiel Command announced initiatives to leverage predictive maintenance tools demonstrated by Delta Airlines to improve C-5, B-1, and C-130 aircraft maintenance.⁴⁰ In the case of the B-1, General Pawlikowski, then AFMC commander, noted that the government originally purchased the data rights to the aircraft, but failed to make use of the data.⁴¹ Recognizing the tension between government and industry on data rights, General Pawlikowski acknowledged that some implementation efforts will require entirely organic resources.⁴² Organic approaches by the Services to adopt machine learning methods for predictive maintenance is enabled by the DoD's supply catalog and Service-specific databases that already track maintenance activity. With the exception of contractor-supported weapons systems, supply orders for organic maintenance activity use national stock numbers to track inventory using the Defense

³⁹ "7 Ways Airlines Use Artificial Intelligence and Data Science to Improve Operations," Altexsoft, July 10, 2018, <https://www.altexsoft.com/blog/datascience/7-ways-how-airlines-use-artificial-intelligence-and-data-science-to-improve-their-operations/> (accessed January 2, 2020).

⁴⁰ Marcus Weisgerber, "The U.S. Air Force Is Adding Algorithms to Predict When Planes Will Break," *Defense One*, May 15, 2018, <https://www.defenseone.com/business/2018/05/us-air-force-adding-algorithms-predict-when-planes-will-break/148234/> (accessed January 2, 2020);

⁴¹ Marcus Weisgerber, "The U.S. Air Force Is Adding Algorithms to Predict When Planes Will Break," *Defense One*. (accessed January 2, 2020).

⁴² Ibid.

Logistics Agency's database, FED LOG.⁴³ By constraining data entry to specific fields and logging supply and maintenance transactions, these databases give the Services a head start on adopting machine learning tools by avoiding data contracting services with manufacturers.⁴⁴

Intelligence

Similar to the maintenance community, the intelligence community generates tremendous volumes of data. Often unable to process all the data, the intelligence community increasingly looks to automated, machine learning tools to assist the human analyst. The first publicly known initiative, the Algorithmic Warfare Cross Functional Team, known as Project Maven, began in April 2017.⁴⁵ Motivated by overwhelmed intelligence analysts combing through hundreds of hours of video, the DoD sought to automate mundane tasks such as searching for objects of interests in otherwise uninteresting video.⁴⁶ Project Maven teamed with AI industry leader, Google, to label video and image data gathered from drones, train a neural network to classify the images,

⁴³ "Information Operations (J6)," Defense Logistics Agency, <https://www.dla.mil/HQ/InformationOperations/Offers/Products/LogisticsApplications/FEDLOG.aspx> (accessed February 24, 2020).

⁴⁴ "DISA Developed Application Chosen to Consolidate Several Air Force Aircraft Maintenance Systems," Defense Information Systems Agency, October 21, 2019, <https://disa.mil/NewsandEvents/2019/application-consolidate-Air-Force-maintenance-systems> (accessed February 25, 2020).

⁴⁵ U.S. Department of Defense, Under Secretary of Defense for Intelligence & Security, "Disruption in UAS: The Algorithmic Warfare Cross-Functional Team (Project Maven)," by Lieutenant General Jack Shanahan, <http://airpower.airforce.gov.au/APDC/media/Events-Media/RAAF%20AP%20CONF%202018/1130-1200-Shanahan-Disruption-in-UAS-The-AWCFT.pdf> (accessed January 12, 2020).

⁴⁶ Marcus Weisgerber, "The Pentagon's New Algorithmic Warfare Cell Gets Its First Mission: Hunt ISIS," *Defense One*, May 14, 2017, <https://www.defenseone.com/technology/2017/05/pentagons-new-algorithmic-warfare-cell-gets-its-first-mission-hunt-isis/137833/> (accessed January 12, 2020).

and enable downstream analytics.⁴⁷ Project Maven, heralded as an early success for the DoD, faced many challenges in the early stages including clear, uniform data labeling.⁴⁸

Security constraints in the intelligence community pose another challenge to adoption. Imagery data collected from overhead sources typically reside on highly classified, compartmentalized network enclaves that significantly restrict user access to large, labeled data sets. Similarly, obtaining security clearances and accrediting contractor facilities to process sensitive intelligence data presents high cost burdens on contractors and lengthy timelines for approval. High capital costs for security combined with low profit potential and few opportunities for algorithm reuse in a commercial setting, disincentivizes small, AI startup companies from pursuing DoD contracts in the intelligence enterprise.⁴⁹

The challenge facing the achievement of future Agile and DevSecOps approaches will be shifting the mindset for programs that manage the intersection of software development and hardware. Even though the future certainly holds more and more software-centric programs, developing and procuring systems capable of producing physical, kinetic effects, such as ships, aircraft, and armor, will remain a mainstay of

⁴⁷ Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence*, 27; Tom Simonite, “Pentagon Will Expand AI Project Prompting Protests at Google,” May 29, 2018, <https://www.wired.com/story/googles-contentious-pentagon-project-is-likely-to-expand/> (accessed January 12, 2020).

⁴⁸ Interview with two machine practitioners familiar with Project Maven, December 28, 2019 and January 29, 2020.

⁴⁹ Interview with two machine practitioners familiar with Project Maven, December 28, 2019 and January 29, 2020.

lethal force. Resolving friction between the warfighter requirements generation, contract fulfillment, and independent OT&E pinch points will likely intensify.⁵⁰

⁵⁰ U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, x-xi, 23, 39; Susanna V. Blume and Molly Parrish, *Make Good Choices, DoD: Optimizing Core Decisionmaking Processes for Great-Power Competition*, 6-7.

Chapter 5: Safety-Critical AI Development

Rushing the adoption of machine learning AI to accelerate autonomous system development risks pushing immature technologies into operations and, without adequate safety controls, could increase the potential for accidents. Armed with the Agile Acquisition Framework (AAF) and an improved software development enterprise, program managers, pressured by DoD and Service leadership to accelerate AI adoption, may unwittingly provide fertile ground for introducing new hazards into decision-making, training, support, and combat operations. Autonomous systems incorporating AI methods are not inherently unsafe, but the character of human interaction with these systems can lead to misplaced trust, unsafe practices, and the increased potential for accidents.¹ The system safety community defines an accident as an “undesired or unplanned loss, including a loss of human life or human injury, property damage, environmental pollution, or mission loss.”² The DoD currently emphasizes reliability based approaches to safety rather than embracing a systems-based safety methodology aimed at addressing the inherent complexities and dynamics between component, environment, and user psychology. Incorporating learning AI into increasingly sophisticated automated systems and DoD missions requires considering the whole socio-

¹ U.S. Department of Defense, Defense Innovation Board, *AI Principles: Recommendation on the Ethical Use of Artificial Intelligence by the Department of Defense* (Washington DC: Government Printing Office, October 31, 2019), 5, https://media.defense.gov/2019/Oct/31/2002204458/-1/-1/0/DIB_AI_PRINCIPLES_PRIMARY_DOCUMENT.PDF (accessed March 13, 2020).

² Nancy G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety* (Cambridge: MIT Press, 2011), 181.

technical system, deliberately deriving safety requirements, and enforcing safety constraints and controls to reduce the risk of accidents.³

Humans Versus Automated Systems

Automated control of safety-critical processes is nothing new, but learning AI adds complexity to accident prevention that requires developers and users rethink the relationship between humans and automation. Automation often creeps into our lives to boost profitability by cutting human labor costs and increasing productivity through improved process optimization.⁴ Companies seeking to boost profits frequently target dull, dirty, or dangerous jobs for automation.⁵ Repetitive tasks typical in manufacturing and logistics sectors translate easily into quantifiable rules that engineers use to program computers and robots. The automotive industry offers many examples of computer-controlled manufacturing replacing machinists, welders, and assembly workers with robots. Previously, limited task proficiency and flexibility confined automated systems to controlling processes through clearly, statically-defined logical rules. Increasingly capable automation enabled by learning algorithms expands the reach of the technology into a variety of industries otherwise thought immune to automation, including retail customer service, transportation, and healthcare. Militaries across the globe, seeking to gain a competitive advantage, have rushed to apply industry's AI advances to automated systems on the battlefield.

³ Nancy G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*, 82; Nancy G. Leveson and John P. Thomas, *STPA Handbook*, (March 2018), 1, https://psas.scripts.mit.edu/home/get_file.php?name=STPA_handbook.pdf (accessed March 28, 2020).

⁴ Mark Muro et al., *Automation and Artificial Intelligence: How Machines Are Affecting People and Places* (Washington DC: Brookings, January 2019), 13, https://www.brookings.edu/wp-content/uploads/2019/01/2019.01_BrookingsMetro_Automation-AI-Workforce_Report_Muro-Maxim-Whiton.pdf (accessed January 17, 2020).

⁵ Kelley M. Sayler, *Artificial Intelligence and National Security*, CRS Report No. R45178 Version 7 (Washington DC: Congressional Research Service, 2019), 27

Increased automation on the battlefield is a natural ally for militaries seeking to protect their forces from casualties and to extend their operational reach.⁶ The U.S. military uses automation across all domains to protect its forces and enhance lethality. Systems such as the RQ-4 remotely piloted aircraft (RPA), Aegis cruisers, and Patriot missile batteries all incorporate advanced automated control systems to enhance their effectiveness.

These systems' architectures, each compliant with the DoD's policy on autonomous weapons, require humans either *in* or *on* the decision-making loop to ensure that the system executes human intent. Most automated systems require human supervision to deal with edge and corner cases not specified in the system's original programming and to serve as an emergency backup able to take over if the automation fails. Humans, however, make poor monitors of automated systems due to the psychology of the interaction.⁷ For humans, monitoring automated processes makes for boring work that leads to further mental disengagement from the controlled process. Additionally, when automated systems fail, requiring human intervention to prevent an accident, operators typically require time to orient themselves to assess the situation and affect corrective action.⁸ The disengagement between human monitors and automated systems can lead to significant divergence between the human's mental model of the system and the system's actual state. To guard against overreliance, automated systems cannot rely on humans exclusively as a crutch to resolve automated system failures and shortcomings. Developers and operators will face the challenge of tailoring system

⁶ Paul J. Springer, *Outsourcing War to Machines: The Military Robotics Revolution* (Santa Barbara: Praeger, 2018), 74-114.

⁷ Nancy G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*, 275.

⁸ *Ibid.*, 276-278.

design and usage to balance the division of tasks between human and machine that engages humans in the controlled task and simultaneously leverages the enabling endurance and precision of machines.

Misplaced Trust

Human reliance on predictable and reliable automated systems can lead to overreliance and complacency via misplaced trust. Calibrating human trust to autonomous system performance remains an active area of research. In *Autonomous Horizons*, the USAF dedicates significant consideration to the development of the tenants of trust for autonomous systems.⁹ Overtrust and undertrust in autonomous systems requires careful consideration. On the one hand, overtrust contributes to complacency and overreliance typical of accidents involving aviation autopilots.¹⁰ For instance, while automated driver-assistance technologies have become increasingly capable, no vehicle today offers autonomy greater than SAE Level 2, requiring drivers to remain fully engaged in the driving task.¹¹ Despite warnings from the NTSB and manufactures, users tend to overtrust driver-assistance features based on their increasing capability and

⁹ U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, by Dr. Greg L. Zacharias (Maxwell AFB: Air University Press, March 2019), 77.

¹⁰ NTSB, *China Airlines Boeing 747-SP, N4522V, 300 Nautical Miles Northwest of San Francisco, California, February 19, 1985*, NTSB/AAR-86/03 (Washington DC: Government Printing Office, March 29, 1986) https://www.faa.gov/about/initiatives/maintenance_hf/library/documents/media/human_factors_maintenanc/e/china_airlines_boeing_747-sp.n4522v.300_nautical_miles_northwest_of_san_francisco.california.february_19.1985.pdf (accessed March 23, 2020); Nick Oliver et al., “The Tragic Crash of Flight AF447 Shows the Unlikely But Catastrophic Consequences of Automation,” *Harvard Business Review*, September 15, 2017, <https://hbr.org/2017/09/the-tragic-crash-of-flight-af447-shows-the-unlikely-but-catastrophic-consequences-of-automation> (accessed March 23, 2020).

¹¹ “Tesla Crash Investigation Yields 9 NTSB Safety Recommendation,” National Transportation Safety Board, February 25, 2020, <https://www.nts.gov/news/press-releases/Pages/NR20200225.aspx> (accessed March 7, 2020).

reliability.¹² On the other hand, undertrust typically leads to the automated systems falling into disuse and may contribute to accidents resulting from situations where humans became saturated or confused by numerous or noisy inputs.¹³ For non-learning systems, standardized procedures, user training, and human-centered design form many of the primary socio-technical tools developers and operators use to guard against overreliance, complacency, and misplaced trust, manifesting as overtrust and undertrust.¹⁴ Learning algorithms compound complications in this dynamic by adapting their output behavior based on the accumulation of experience in the form of data. The novel and notably surprising, non-human solutions presented by deep learning algorithms to the games Go (Wei-Chi) and several Atari games foreshadow challenges for developing human trust in algorithms that lack cognitive congruence with humans.¹⁵

DARPA is presently working to advance many aspect of AI including algorithm development and explainable AI initiatives.¹⁶ A DARPA project aimed at developing algorithms for autonomous air-to-air combat supports both efforts. The program seeks to develop AI agents capable of executing air-to-air visual-range basic fighter maneuvers (BFM)–dogfighting.¹⁷ The program also aims to calibrate pilot trust in algorithm

¹² “Driver Errors, Overreliance on Automation, Lack of Safeguard, Led to Fatal Tesla Crash,” National Transportation Safety Board, September 12, 2017, <https://www.nts.gov/news/press-releases/Pages/PR20170912.aspx> (accessed March 7, 2020).

¹³ Paul Scharre, *Army of None: Autonomous Weapons and the Future of War* (New York: W.W. Norton & Company, Inc., 2018), 166.

¹⁴ Nancy G. Leveson, *CAST Handbook: How to Learn More from Incidents and Accidents*, (2019), 8, <http://sunnyday.mit.edu/CAST-Handbook.pdf> (accessed February 25, 2020).

¹⁵ *AlphaGo*, directed by Greg Kohs, Moxie Pictures & Reel As Dirt, 2017, 1:31, <https://www.youtube.com/watch?v=WXuK6gekU1Y> (accessed January 26, 2020); Lex Fridman, “Deep Learning State of the Art (2020),” MIT Deep Learning Series, January 10, 2020, <https://www.youtube.com/watch?v=0VH1Lim8gL8> (accessed February 4, 2020); U.S. Air Force, Office of the Chief Scientist, *Autonomous Horizons: The Way Forward*, 84.

¹⁶ Kelley M. Sayler, *Artificial Intelligence and National Security*, 31.

¹⁷ Graham Warwick, “DARPA Automated Dogfighting to Develop Pilot Trust in AI in Combat,” *Aviation Week and Space Technology*, April 20-May 3, 2020, 36-37.

performance with the ultimate aim of promoting synergy between human and machine rather than cultivating an adversarial human-machine relationship.¹⁸ Researchers are examining cognitive congruence from both the perspective of human attitudes towards AI and adapting AI to humans. Dr. Javorsek, program manager of the Alpha Dogfight program, plans to measure pilot trust by measuring and quantifying the degree to which the human pilot delegates control of the BFM fight to the AI in favor of accomplishing human-specific battle management tasks.¹⁹ At present most explainable AI research focuses on helping the human adapt to the autonomous system. Dr. Ayanna Howard, a roboticist at Georgia Tech, notes that historically designers focused on controlling and accommodating human behavior when designing automated safety-critical processes.²⁰ She emphasizes that most safety systems incorporating robots center on controlled environments to encourage predictable human-machine dynamics. However, she notes that robots capable of adapting to humans may provide options for safety-critical human-machine solutions, but such solutions have only been demonstrated in a laboratory environment.²¹

A Systems-Based Approach

In order to foster safe and effective interactions between humans and autonomous systems supporting military missions, the DoD must continue to develop and apply deliberate processes for resolving or mitigating automated system safety hazards and risks. As the DoD moves to accelerate adoption of AI learning algorithms using rapid

¹⁸ Daniel Javorsek, e-mail message to author, May 16, 2020.

¹⁹ Graham Warwick, “DARPA Automated Dogfighting to Develop Pilot Trust in AI in Combat,” 36-37.

²⁰ Ayanna Howard, interview by Lex Fridman, January 17, 2020, AI Podcast <https://www.youtube.com/watch?v=J21-7AsUcgM> (accessed April 11, 2020).

²¹ Ayanna Howard, interview by Lex Fridman, January 17, 2020.

software acquisition methods, the DoD must cultivate a safety culture that recognizes and appreciates the complexity introduced by learning AI, and it must embrace an accident prevention framework capable of dealing with the associated increase in system complexity. Dr. Nancy Leveson, an MIT professor and researcher, contributed to developing improved models of accident causation and new methodologies for safety analysis and accident prevention. Her models and methods leverage systems theory and emphasize analysis of system interactions and dynamics as a whole system.

Approaches like Dr. Leveson's Systems-Theoretic Accident Model and Processes (STAMP) frames accident causality and prevention in terms of system states and interactions bounded by a set of top-down system safety constraints designed to avoid hazards that contribute to accidents.²² STAMP eschews event-chain methodologies that focus on breaking the cascade of failures leading to an accident. Legacy models such as the Domino and Swiss Cheese theories conceptualize accident causation as a chain of events initiated by a single root-cause.²³ Thus, mishap prevention strategies based on these models emphasize recognizing and stopping the chain of events leading to an accident rather than addressing the systemic pressures and dynamics that promote the conditions for an emergent accident. Event-chain models also tend to accentuate a culture of blame and liability since human error remains one of the easiest causes to pinpoint.²⁴ Accident investigations based on event-chain models trace a series of events backward in time from the proximate cause of the mishap and typically stop when the

²² Nancy G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*, 73.

²³ Ibid, 16-18, 34, 91.

²⁴ Ibid, 16-18.

investigation reaches the first human in the chain.²⁵ Continuing to apply legacy mishap prevention safety practices that emphasize root cause analysis based on event chains will fail to mitigate the novel hazards that will arise from machine learning AI. Dr. Leveson suggests that assigning blame for accidents does little to prevent future accidents.²⁶ Instead, investigations need to widen their scope beyond the proximate causes of the mishap and deemphasize blame to capture systemic pressures and interactions that contributed to the system's migration into a hazardous state from which an accident emerged. Emphasizing a blameless approach to system safety melds well with Agile software development approaches that prize solutions over blame.²⁷ However, investigations and reporting on accidents involving automated systems indicate a trend in the opposition direction.²⁸

Reliability Versus Safety

In general, policy and regulation aimed at improving safety commonly centers on reliability analysis and improvements to mitigate the risk of accidents.²⁹ Too often, designers and users equate reliability with safety, thinking that fewer accidents will result from increased component or system reliability.³⁰ Restated using the autonomy vocabulary from Chapter 3, developers frequently rely on improved autonomous system

²⁵ Ibid, 21-47.

²⁶ Ibid, 57.

²⁷ U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage* (Washington DC: Government Printing Office, May 3, 2019), 10, S79, S108, S192, https://media.defense.gov/2019/Apr/30/2002124828/-1/-1/0/SOFTWAREISNEVERDONE_REFACTORINGTHEACQUISITIONCODEFORCOMPETITIVEADVANTAGE_FINAL.SWAP.REPORT.PDF (accessed March 26, 2020).

²⁸ Ian Bogost, "Can You Sue a Robocar?" *The Atlantic*, March 20, 2018, <https://www.theatlantic.com/technology/archive/2018/03/can-you-sue-a-robocar/556007/> (accessed February 25, 2020).

²⁹ Nancy Leveson, "Are You Sure Your Software Will Not Kill Anyone?" *Communications of the ACM* 63, no. 2 (February 2020), 26.

³⁰ Nancy G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*, 57, 64, 73, 173.

task proficiency to advance safety. However, improving proficiency or reliability alone is insufficient to enhance system safety. Reliability refers to the rate of failure for components or systems under specified conditions. Reliability measures, such as mean time between failure, quantify the frequency of component or system failures and form the basis of failure modes and effects analysis (FMEA) and system accident predictions such as probability of loss of aircraft (PLOA). Despite the widespread inappropriate assumption of equivalence between safety and reliability, a reliable system does not equate to a safe system.³¹ In some cases, increasing reliability may even reduce system safety.³² The confusion between safety and reliability may stem from the significant attention that avoiding component failure receives in traditional engineering disciplines.³³ Material, mechanical, electrical, and semiconductor hardware failure rates tend to dominate safety analysis. Safety review boards often focus on the potential for single-point failures, placing undue emphasis on reliability figures when assessing the risk of integrated system accidents.

Dr. Leveson cites many examples of accidents caused by interactions rather than by component failures. In one example, an auto ferry, not suffering from any component failures, capsized killing 193 people after unanticipated interactions between vessel and harbor designs, scheduling and operations, and crew dynamics led to a hazardous system state.³⁴ In 1993, an Airbus A320 landing in heavy rains and crosswinds with no failed systems, but unable to decelerate, ran off the runway killing one crew member and one passenger. Anticipating the crosswind landing, the pilot correctly landed the aircraft in a

³¹ Ibid, 7.

³² Ibid, 11-12.

³³ Ibid, 8.

³⁴ Ibid, 13.

banked attitude to minimize the aircraft's lateral drift over the runway. With the banked attitude, the aircraft touched down on only one of the main landing gear. To prevent the dangerous possibility of thrust reverser deployment in-flight, the A320's flight control software locked out thrust reverser deployment until it detected weight on both of the main landing gear.³⁵ The aircraft continued its ground run on one wheel for several seconds, thus inhibiting activation of the thrust reversers. The A320 flight controls performed exactly as designed without failure yet the system migrated into a hazardous state from which an accident resulted. Similarly, the Mars Polar Lander (MPL) crashed into the Martian surface after the descent engine cut off prematurely when the landing legs deployed. The probe's normal deployment of the landing legs during descent generated a noisy transient signal on each of the leg's touch down sensors. The descent control software interpreted the transient signal as contact with the Martian surface and shutdown the descent engine, and the spacecraft crashed into the planet's surface.³⁶ In this scenario, none of the individual components of the probe failed, but their interactions contributed to a hazardous system state that led to the crash.³⁷ Since the probe's components all functioned without failure and as programmed, the proximate cause of the crash was an insufficiently robust logic scheme to detect the surface of the planet and reject false detections. Improved component reliability or sensor redundancy would not have avoided the crash.

³⁵ R.D. Hawkins et al., "The Principles of Software Safety Assurance" (31st International System Safety Conference, 2013), 3, <https://www-users.cs.york.ac.uk/~rhawkins/papers/HawkinsISSC13.pdf> (accessed April 18, 2020).

³⁶ JPL Special Review Board, *Report on the Loss of the Mars Polar Lander and Deep Space 2 Missions* (Jet Propulsion Laboratory, March 22, 2000), 26 https://spaceflight.nasa.gov/spacenews/releases/2000/mpl/mpl_report_1.pdf (accessed April 11, 2020).

³⁷ Nancy G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*, 8.

In each of these examples, no variety of bottom-up component reliability or fault tree methods, analyzing each component in isolation, could predict these hazards or protect these complex systems from mishap.³⁸ The accidents emerged from complex component, system, and environmental interactions. Applying a systems theory approach to accidents reframes safety as an emergent property that results from the interactions of the elements of the system.³⁹ Dr. Leveson's systems-based approach views safety as a control problem itself.⁴⁰ Rather than focusing on reducing failures through reliability improvements or pinning mishap prevention solely on the detection and arrest of proximate causes, Dr. Leveson reframes the design problem as a top-down safety control problem. Instead of breaking mishap chains to avoid an accident, users and operators, developers and engineers collaborate to identify hazardous system states, describe system safety constraints, and then design and implement specific controls to prevent the system from migrating into a hazardous states. When analyzing system hazards and designing system controls and constraints to guard against unsafe actions, Dr. Leveson emphasizes the importance of considering both technical and non-technical aspects of the system. She emphasizes this point in her detailed analysis of the 1994 friendly fire shootdown of two U.S. Army UH-60 Blackhawks in Iraq. Dr. Leveson's analysis of the unsafe control actions that led to the accident included detailed consideration of technical, engineered components as well as behavior-shaping mechanisms that influenced human decision making.⁴¹ Most safety investigations focus on operations and the proximate accident

³⁸ Ibid, 61-67.

³⁹ Ibid, 67.

⁴⁰ Ibid.

⁴¹ Ibid, 103-167.

causes and fail to thoroughly analyze interactions across the entire developmental and operational socio-technical system.⁴²

Software System Safety

Prior to roughly 1980, engineers typically avoided software in safety-critical control applications and instead relied on hardware-based, analog, electro-mechanical control systems to govern safety-critical processes.⁴³ Unlike software-based controllers, hardware controllers offered engineers devices with a finite and manageable set of physical states characterized by well understood and quantified failure modes and reliability metrics. This enabled exhaustive testing and allowed engineers to confidently apply these controllers to safety-critical processes and devise specific procedures accounting for a finite set of failure modes. However, the increasing flexibility and decreasing costs offered by software drove the incorporation of software-based controllers in safety-critical applications. While the increased flexibility afforded by software permitted developers to devise increasingly sophisticated engineered systems, the resultant architectures became too difficult for humans to intellectually manage.⁴⁴ Software introduces an exponential increase in the number of possible states of the system making exhaustive developmental testing impossible and leaving potentially hazardous system states for the user to discover during operation.⁴⁵ Dr. Leveson refers to this predicament as software's "curse of flexibility."⁴⁶ With the introduction of software,

⁴² Ibid, 82, 103.

⁴³ Nancy Leveson, "Are You Sure Your Software Will Not Kill Anyone?" 25.

⁴⁴ James Somers, "The Coming Software Apocalypse," *The Atlantic*, September 26, 2017, <https://www.theatlantic.com/technology/archive/2017/09/saving-the-world-from-code/540393/> (accessed April 10, 2020).

⁴⁵ Nancy Leveson, "Are You Sure Your Software Will Not Kill Anyone?" 27-28.

⁴⁶ Nancy G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*, 50.

a piece of silicon can become either an aircraft autopilot or lawn sprinkler controller each with vastly different safety consequences.⁴⁷ However, when contemplating the contribution of software to system safety, software is an abstraction. Software does not contain energy, it cannot catch fire or maim people or damage property.⁴⁸ Software safety is a misleading term since software, divorced from its physical hardware can neither be safe or unsafe.⁴⁹

In 1991, a numerical round-off error contributed to the inability of a Patriot missile to shoot down an incoming SCUD missile that killed 28 Americans.⁵⁰ At the time of the accident, the Patriot's software stored time as an integer count of tenths of seconds that required multiplication by 0.1 to convert time into seconds when calculating an incoming missile's trajectory to support tracking and interception.⁵¹ Because binary cannot exactly represent 0.1, the round-off error associated with the 24-bit representation of 0.1 is 0.000000095.⁵² At the time of the SCUD launch, the Patriot had been in operation for approximately 100 hours, resulting in a 0.34 second error in the Patriot's time computation ($0.000000095 \times 100 \times 60 \times 60 \times 10 = 0.34$).⁵³ The Patriot could not accurately track the 1,676 meter per second SCUD because the 0.34 second error meant that the missile flew approximately 500 meters outside of the radar's track. In the Patriot example, when developers translated the tracking algorithm into software requirements,

⁴⁷ James Somers, "The Coming Software Apocalypse."

⁴⁸ Nancy Leveson, "Are You Sure Your Software Will Not Kill Anyone?" 26.

⁴⁹ Ibid.

⁵⁰ U.S. General Accounting Office, *Patriot Missile Defense: Software Problem Led to System Failure at Dhahran, Saudi Arabia* (Washington DC: Government Printing Office, February 1992), 1-4, <https://www.gao.gov/products/IMTEC-92-26> (accessed April 8, 2020).

⁵¹ Douglas N. Arnold, "The Patriot Missile Failure," <http://www-users.math.umn.edu/~arnold/disasters/patriot.html> (accessed April 10, 2020).

⁵² Douglas N. Arnold, "The Patriot Missile Failure."

⁵³ Ibid.

they failed to consider the integrated hardware-software system's limitations. Porting software from one application to another, known as software reuse, a common cost-saving measure, can contribute to accidents when software is considered in isolation and without consideration for the combined hardware-software system.⁵⁴

In most instances, discipline engineers design control systems and translate the design into software requirements for programmers to implement in software code.⁵⁵ Software programmers, often lacking depth in the engineering discipline that provided the software's requirements, typically focus their effort on coding to specified requirements. For decades, writing software required the dedication of uniquely talented individuals capable of thinking like computers to translate requirements into code, a perspective that often further divorces programmers from the design they are implementing in code.⁵⁶ Writing software code can be so frustrating and challenging that programmers become preoccupied with simply getting their programs to run that they have little capacity left over to contribute to feedback on system design.⁵⁷

In the case of the Mar Polar Lander, engineers *verified* that the software met its requirement specification, but failed to *validate* that the resultant design would function as intended. While nuanced definitions of verification and validation vary across technical disciplines, the following themes persist across a majority of fields. On the one hand, verification refers to the process of ensuring that a system meets its specified build-

⁵⁴ Nancy Leveson, "Are You Sure Your Software Will Not Kill Anyone?" 26.

⁵⁵ Nancy G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*, 48.

⁵⁶ James Somers, "The Coming Software Apocalypse."

⁵⁷ Ibid.

to requirements.⁵⁸ Validation, on the other hand, refers to how well the realized system performs its intended task.⁵⁹ Verification determines whether or not the system was built correctly; validation evaluates whether or not the correct system was built.⁶⁰ According to DoDD 5000.59, the governing directive for modeling and simulation, verification refers to determining whether a model meets the developer's specification and conceptualization, whereas validation refers to the degree to which the model matches its intended real world application.⁶¹ In the MPL example, the JPL team verified that the flight control software was coded as specified to shut down the engine when a touchdown signal was received. However, the team failed to specify, develop, and integrate logical protections that guard against spurious signals. The subsequent investigation revealed that noise from the touchdown sensors was a known issue, but systemic pressure to move quickly while keeping costs low drove the team to curtail system level validations.⁶²

Modern software tools, including improved debugging and automated testing utilities and scripts, help to speed up specification verification and reduce build-to errors, but do little to prevent inherent design shortfalls and account for the increasing complexity of engineered systems. Recent efforts to expand model-based design tools and encourage thorough, logical planning before typing a single line of code aim to

⁵⁸ Avner Engel, *Verification, Validation, and Testing of Engineered Systems* (Hoboken: John Wiley & Sons, Inc., 2010), 16-17; C. Warren Axelrod, *Engineering Safe and Secure Software Systems* (Boston: Artech House, 2013), 49.

⁵⁹ Avner Engel, *Verification, Validation, and Testing of Engineered Systems*, 16-17.

⁶⁰ Ibid.

⁶¹ U.S. Department of Defense, Under Secretary of Defense of Research & Engineering, *DoD Modeling and Simulation (M&S) Management*, DoDD 5000.59 (Washington DC: Government Printing Office, August 8, 2007, Incorporating Change 1, October 15, 2018), 8, <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/500059p.pdf> (accessed April 10, 2020).

⁶² JPL Special Review Board, *Report on the Loss of the Mars Polar Lander and Deep Space 2 Missions*, 6; Nancy G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*, 85-86.

ameliorate these challenges.⁶³ Engineers using model-based design tools prototype their algorithms and control systems using discipline-specific computational tools capable of generating autocode output for implementation. These approaches avoid time-consuming and error-prone requirements translation processes between engineers and software developers. Model-based methods also offer engineers and scientists simulation environments to check the capability of the design by testing use cases. Promoting deliberate and thorough planning that is logically complete is another means of reducing software development errors.⁶⁴ One such tool, TLA+, uses formal mathematical and logical notation to describe an algorithm helping developers examine and validate the design before typing a single line of code.⁶⁵

With the DoD's recent acquisition policy reforms formalizing incorporation of Agile software development approaches, the DoD must integrate a robust system safety approach to mitigate the novel safety risks that emerge from software development supporting learning AI algorithms hosted in automated weapons systems. In fact, control of safety-critical systems using software from Agile methods remains an open area of research.⁶⁶ On the surface, Agile software development methods' emphasis on rapid prototyping and cultural resistance to formal requirements analysis seems ill-suited to software development efforts integrating learning algorithms in safety-critical applications.⁶⁷ However, Agile and DevSecOps practices need not contradict the integrated, multidisciplinary safety requirements analysis dictated by the systems-based

⁶³ James Somers, "The Coming Software Apocalypse."

⁶⁴ Ibid.

⁶⁵ Ibid.

⁶⁶ Rashidah Kasauli, et al., "Safety-Critical Systems and Agile Development: A Mapping Study," (44th Euromicro Conference on Software Engineering and Advanced Applications, 2018), 470-477, <https://arxiv.org/pdf/1807.07800.pdf> (accessed April 13, 2020).

⁶⁷ *Autonomy*, directed by Alex Horwitz, A Car & Driver Film, Haven Entertainment, 2019, 1:20.

approach called for by Dr. Leveson. By acknowledging and embracing the inherent flexibility of software, developers, users, testers, and operators can cooperate to build a robust socio-technical system, rooted in Agile and DevSecOps practices, that accounts for the design and enforcement of top-down system safety controls. Instead of expanding the requirements and specification waterfall to address every permutation of system states, shown to be an impossibility, developers can replace the waterfall with continuous integration and continuous delivery such that prototype software is frequently returned to the customer for operational, validation testing with the recognition that feedback from the field will necessarily require additional iterations to improve the system's safety for continuous delivery.

Chapter 6: Conclusion

Today, the DoD continues its struggle to adopt Artificial Intelligence (AI) technologies despite increasing pressure from executive policy makers and Congressional leaders. AlphaGo's 2016 defeat of Lee Sedol and China's subsequent 2017 declaration to dominate AI research and development served as a Sputnik moment for U.S. policy makers, spooking them into action over fears of losing an AI-arms race.¹ Despite growing Congressional and Presidential attention and increases in government funding for AI and autonomy technologies, RAND found that the DoD was poorly postured to take advantage of AI technologies advanced in the last decade in the commercial sector.²

Much of the struggle stems from the DoD's lack of widespread familiarity and literacy with the emerging technologies. Already algorithms shape our perception of reality and constrain our choices across a variety of sectors and applications ranging from entertainment choices to financial product offerings. As the technologies continue to mature, the mechanics of their function grow increasingly opaque to the average user. The trend is particularly problematic for the DoD, an organization seeking to leverage machine advantages to offset adversary anti-access threats that erode U.S. operational reach and coalition decision-making. Part of the challenge of developing widespread AI literacy stems from the lack of a unified AI vocabulary, accessible to non-experts

¹ Georgia Perry, "The AI Cold War That Threatens Us All," *Wired*, October 23, 2018, <https://www.wired.com/story/ai-cold-war-china-could-doom-us-all/> (accessed February 5, 2020); Henry A. Kissinger, "How the Enlightenment Ends," *The Atlantic*, June 2018, <https://www.theatlantic.com/magazine/archive/2018/06/henry-kissinger-ai-could-mean-the-end-of-human-history/559124/> (accessed February 5, 2020); Graham Allison, "Is China Beating America to AI Supremacy?" *The National Interest*, December 22, 2019, <https://nationalinterest.org/feature/china-beating-america-ai-supremacy-106861> (accessed December 27, 2020).

² Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence* (Santa Monica: RAND Corporation, 2019), 106, 129; Interview with a Defense Innovation Unit employee, December 28, 2019.

throughout the Department. As RAND noted, however, even among AI practitioners and experts there is little agreement on the definition of AI and whether or not developing one is worth it. For a massive organization like the DoD, charged with integrating the game-changing capabilities of a complex and disruptive emergent technology, a common lexicon establishes a necessary cultural cornerstone upon which to build literacy. Establishing a baseline level of workforce literacy and understanding, as well as targeted education for key stakeholders, would help accelerate organizational adoption by dispelling myths and fears to increase confidence and familiarity with the applications and limitations of AI and autonomy. Without concentrating on improving widespread workforce literacy for learning algorithms and autonomous system concepts, the DoD will continue to struggle to integrate these disruptive technologies as a fearful or oblivious workforce either neglects or resists the transformative tools.

But literacy alone will not transform the DoD into an AI-ready institution. Widespread adoption and integration of AI and autonomous systems require deliberate investments in hardware and software platforms to support machine learning algorithms and cultivate productive automated systems. The DoD's sclerotic acquisition system, geared toward avoiding manufacturing risks on large-scale, capital-intensive major weapons system acquisitions, has for too long applied a one-size fits all approach to gathering warfighter requirements, securing definitive funding portfolios, and establishing time-certain program schedules each with federated communities scattered across OSD and the Services focused on requirements, development, testing, support, training, and operations. The problem at the root of the scattered, federated approach is that it prevents the collection, curation, and aggregation of data. The serial, stove-piped

approaches will not work for developing learning AI algorithms that leverage vast, enterprise-level data to feed learning algorithms.³ Additionally, software constitutes an essential building block of AI systems, and the 2019 Defense Innovation Board's (DIB) software report declared the DoD software development broken.⁴ Recently, the DoD made significant operational changes to the Defense Acquisition System. At the start of 2020, the Undersecretary for Defense for Acquisition and Sustainment (USD A&S) released interim policy guidance specifically for software development to stimulate Agile software management approaches within the DoD and the traditional industrial base. The software policy, along with the rewrite of DODI 5000.02, the instruction governing defense acquisition, marked a major step toward challenging legacy acquisition practices, thawing the frozen middle, and encouraging new ways of rapidly integrating and fielding disruptive technologies that give the warfighter a decisive advantage.⁵ The recent software policy in particular indicates that the DoD acknowledges the DIB's finding that "software is never done," thereby enabling program managers to embrace Agile and DevSecOps practices capable of delivering improved capabilities at the speed of relevance through continuous integration and continuous delivery.⁶

³ Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence*, 57-60.

⁴ U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage* (Washington DC: Government Printing Office, May 3, 2019), i, https://media.defense.gov/2019/Apr/30/2002124828/-1/-/1/0/SOFTWAREISNEVERDONE_REFACTORINGTHEACQUISITIONCODEFORCOMPETITIVEADVANTAGE_FINAL.SWAP.REPORT.PDF (accessed March 26, 2020).

⁵ U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, 26.

⁶ U.S. Department of Defense, Defense Innovation Board, *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, ix; U.S. Department of Defense, *National Defense Strategy*, (Washington DC: Government Printing Office, 2018), 10, <https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf> (accessed December 23, 2020).

Arming PMs with the new Agile Acquisition Framework (AAF) and encouraging iterative software development aimed at continuously delivering updated software to the warfighter presents new challenges and risks. While rapid acquisition frameworks and iterative software development projects using early user testing are not inherently unsafe, these methods introduce the potential for new process seams that careless or unscrupulous managers and developers could exploit to shortcut safety and rush immature systems into the field. Low AI literacy across the DoD could exacerbate the temptation to speed decision making. Decision makers who lack sufficient familiarity with AI and autonomy technologies may not detect inflated promises or appreciate nuanced safety pitfalls when introducing learning algorithms into safety-critical applications. Too often hardware-focused legacy acquisition practices lead developers to pursue reliability improvements to boost safety. But reliability-focused methods using bottom-up analysis of isolated components cannot deal with the myriad complex system interactions and can lead to accidents otherwise deemed improbable. Already the complexity of today's weapon systems deranges system-level accident predictions based on bottom-up reliability approaches, therefore confounding accident prevention measures.⁷ Thus, developers and operators need to collaborate across the entire socio-technical system to ensure that systemic pressures in one area do not drive the entire system into a hazardous state and unnecessarily risk increased accident potential.⁸ Failing to adopt systems-based approaches capable of contending with increasing complexity tempts disaster by ignoring decades of experience with automated systems in

⁷ Nancy G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety* (Cambridge: MIT Press, 2011), 62-63.

⁸ Nancy G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*, 75-82.

aviation and overlooks recent driverless car disasters involving learning algorithms that highlight the technology's limitations and the human tendency to overtrust automation.

Given the shortage of machine learning experts, industry and university machine learning programs have sprung up everywhere. Using the AAF's acquisition pathway for services, the DoD could work to resolve its literacy shortfall by contracting small scale online subscriptions or instruction from industry and academic institutions in order to develop an organic, DoD literacy curriculum.⁹ Regardless of the approach, the DoD's literacy curriculum should ensure basic understanding of the algorithms, capabilities and applications, long-term sustainment considerations, and limitations. With improved workforce literacy, individuals working seemingly unrelated projects to AI and autonomy will recognize the importance of data collection and curation. Since massive quantities of labeled data powers learning algorithms, reinforcing the importance of accurately collecting data at the point of production will enable machine learning experts and data scientists to improve algorithm utility and applicability across the Joint Force. In fact, the need for centralized data curation across the DoD demands a joint approach free of Service parochialism. Feeding learning algorithms with data from across the Joint Force necessitates centralized data curation for the benefit of national security. Breaking down data sharing barriers and encouraging data pooling remains an essential task for DoD leaders.

Sustainment, a joint function common to all Services, offers an excellent example from which the DoD can model its literacy education. Unlike other joint functions,

⁹ Chris Telley, "Info Ops Officer Offers Artificial Intelligence Roadmap," *Breaking Defense*, July 11, 2017, <https://breakingdefense.com/2017/07/info-ops-officer-offers-artificial-intelligence-roadmap/> (accessed May 3, 2020).

sustainment and logistics remains a largely unclassified endeavor with few security barriers. Furthermore, sustainment offers concrete examples familiar to nearly every Soldier, Sailor, Airman, and Marine. And unlike many of the other joint functions, the DoD already has a head start on the collection and curation of data from which to leverage the power of learning algorithms.

Regardless of whether or not the DoD's path for AI adoption is smooth, learning algorithms will persist and become more prominent. They are changing and influencing our choices and reshaping our world. Militaries who understand AI will leverage algorithms to master the art of winning without fighting.¹⁰ It is a false dilemma to suggest that the DoD's slow adoption of AI technologies places the U.S. at a disadvantage. The U.S. should avoid taking China's bait and needlessly pouring money into an AI arms race. Instead, the U.S. should patiently and deliberately invest in widespread AI literacy education. Seeding workforce literacy and interest with the fundamental concepts of learning algorithms will cement sustainable AI and autonomy adoption across the DoD at a faster rate than top-down and peripheral approaches. Only with a knowledgeable workforce will acquisition reforms enable the safe, rapid, responsive, and continuous development, testing, and operation of weapon systems capable of restoring U.S. access to contested environments and meaningfully assist decision makers sifting through the cacophony of disinformation and noise. With widespread workforce literacy underpinning a competent Joint Force, leaders will be able to confidently rely on the decisions of commanders and managers at all echelons to safely

¹⁰ Samuel B. Griffith, *Sun Tzu: The Art of War* (London: Oxford University Press, 1963), 78.

and effectively employ future AI and autonomous systems and maintain advantage over peer competitors.

Bibliography

- Allen, Greg and Taniel Chen. *Artificial Intelligence and National Security*. Cambridge: Belfer Center for Science and International Affairs, 2017.
- Allison, Graham. "Is China Beating America to AI Supremacy?" *The National Interest*, December 22, 2019, <https://nationalinterest.org/feature/china-beating-america-ai-supremacy-106861> (accessed December 27, 2020).
- AlphaGo*. Directed by Greg Kohs. Moxie Pictures & Reel As Dirt, 2017
<https://www.youtube.com/watch?v=WXuK6gekU1Y> (accessed January 26, 2020).
- Altexsoft. "7 Ways Airlines Use Artificial Intelligence and Data Science to Improve Operations." July 10, 2018: <https://www.altexsoft.com/blog/datascience/7-ways-how-airlines-use-artificial-intelligence-and-data-science-to-improve-their-operations/> (accessed January 2, 2020).
- Arnold, Douglas N. "The Patriot Missile Failure." <http://www-users.math.umn.edu/~arnold/disasters/patriot.html> (accessed April 10, 2020).
- Autonomy*. Directed by Alex Horwitz. A Car & Driver Film, Haven Entertainment, 2019.
- Axelrod, C. Warren. *Engineering Safe and Secure Software Systems*. Boston: Artech House, 2013.
- Blume, Susanna V. and Molly Parrish. *Make Good Choices, DoD: Optimizing Core Decisionmaking Processes for Great-Power Competition*. Washington DC: Center for a New American Security, November 2019.
- Bogost, Ian. "Can You Sue a Robocar?" *The Atlantic*, March 20, 2018.
<https://www.theatlantic.com/technology/archive/2018/03/can-you-sue-a-robocar/556007/> (accessed February 25, 2020).
- Brose, Christian. "The New Revolution in Military Affairs: War's Sci-Fi Future." *Foreign Affairs* 98, no. 3 (May/June 2019): 122-124.
- Bzdok, Danilo, Naomi Altman, and Martin Krzywinski. "Statistics Versus Machine Learning." *Nature Methods* 15, no. 4 (April 2018)
<https://www.nature.com/articles/nmeth.4642.pdf> (accessed February 5, 2020).
- Chinese State Council. "Notice of the State Council Issuing the New Generation of Artificial Intelligence Development Plan." Translated by Flora Sapio, Weiming Chen, and Adrian Lo. Washington DC: Foundation for Law & International Affairs, 2017 <https://flia.org/wp-content/uploads/2017/07/A-New-Generation-of-Artificial-Intelligence-Development-Plan-1.pdf> (accessed February 3, 2020).
- Congressional Research Service. "Overview of Artificial Intelligence." CRS In Focus, IF10608, Congressional Research Service, October 27, 2017,
<https://crsreports.congress.gov/product/pdf/IF/IF10608> (accessed February 7, 2020).
- Cronk, Terri Moon. "DoD Unveils Its Artificial Intelligence Strategy." U.S. Department of Defense, <https://www.defense.gov/Explore/News/Article/Article/1755942/dod-unveils-its-artificial-intelligence-strategy/> (accessed February 5, 2020).

- Dastin, Jeffrey. "Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women." *Reuters*, October 9, 2018, <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G> (accessed May 4, 2020).
- Domingos, Pedro. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. New York: Basic Books, 2015.
- Economist staff. "China's Success at AI Has Relied on Good Data." *The Economist*, January 2, 2020, Technology Quarterly, <https://www.economist.com/technology-quarterly/2020/01/02/chinas-success-at-ai-has-relied-on-good-data> (accessed February 4, 2020).
- Engel, Avner. *Verification, Validation, and Testing of Engineered Systems*. Hoboken: John Wiley & Sons, Inc., 2010.
- Eversden, Andrew. "So What Problems Does JEDI Solve, Really?" *Federal Times*, October 30, 2019 <https://www.federaltimes.com/govcon/contracting/2019/10/30/so-what-problems-does-jedi-solve-really/> (accessed May 1, 2020).
- Fowler, Martin. "Continuous Integration." <https://martinfowler.com/articles/continuousIntegration.html> (accessed February 23, 2020).
- Fox, J. Ronald. *Defense Acquisition Reform 1960-2009: An Elusive Goal*. Washington DC: Center of Military History, U.S. Army, 2011.
- Fridman, Lex. "Deep Learning State of the Art (2020)." MIT Deep Learning Series, January 10, 2020, <https://www.youtube.com/watch?v=0VH1Lim8gL8> (accessed February 4, 2020).
- Gadepally, Vijay, Justin Goodwin, Jeremy Kepner, Albert Reuther, Hayley Reynolds, Siddharth Samsi, Jonathan Su, and David Martinez. *AI Enabling Technologies: A Survey*. Lexington, MA: MIT Lincoln Laboratory, 2019, <https://arxiv.org/ftp/arxiv/papers/1905/1905.03592.pdf> (accessed March 13, 2020).
- Gady, Franz-Stefan. "Elsa B. Kania on Artificial Intelligence and Great Power Competition: On AI's Potential, Military Uses, and the Fallacy of an AI Arms Race." *The Diplomat*, December 31, 2019, <https://thediplomat.com/2020/01/elsa-b-kania-on-artificial-intelligence-and-great-power-competition/> (accessed January 1, 2020).
- Griffith, Samuel B. *Sun Tzu: The Art of War*. London: Oxford University Press, 1963.
- Hawkins, R.D., I. Habli, and T.P. Kelly. "The Principles of Software Safety Assurance." 31st International System Safety Conference, 2013, <https://www-users.cs.york.ac.uk/~rhawkins/papers/HawkinsISSC13.pdf> (accessed April 18, 2020).
- Hefron, Ryan. "RDT&E of Autonomous Systems." U.S. Air Force Test Pilot School Short Course Charts, August 23, 2019.
- Hicks, Kathleen H., Andrew Hunter, Jesse Ellman, Lisa Samp, and Gabriel Coll. *Assessing the Third Offset Strategy*. Washington DC: Center for Strategic & International Studies, 2017.

- Horowitz, Michael C. and Lauren Kahn. “The AI Literacy Gap Hobbling American Officialdom.” *War on the Rocks*, January 14, 2020, <https://warontherocks.com/2020/01/the-ai-literacy-gap-hobbling-american-officialdom/> (accessed January 22, 2020).
- Howard, Ayanna. Interview by Lex Fridman, January 17, 2020, AI Podcast <https://www.youtube.com/watch?v=J21-7AsUcgM> (accessed April 11, 2020).
- James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An Introduction to Statistical Learning: With Applications in R*. New York: Springer, 2017, <https://link.springer.com/content/pdf/10.1007%2F978-1-4614-7138-7.pdf> (accessed February 12, 2020).
- JPL Special Review Board. *Report on the Loss of the Mars Polar Lander and Deep Space 2 Missions*. Jet Propulsion Laboratory, March 22, 2000, https://spaceflight.nasa.gov/spacenews/releases/2000/mpl/mpl_report_1.pdf (accessed April 11, 2020).
- Kasauli, Rashidah, Eric Knauss, Benjamin Kanagwa, Agneta Nilsson, and Gul Calikli. “Safety-Critical Systems and Agile Development: A Mapping Study.” 44th Euromicro Conference on Software Engineering and Advanced Applications, 2018, 470-477, <https://arxiv.org/pdf/1807.07800.pdf> (accessed April 13, 2020).
- Kissinger, Henry A. “How the Enlightenment Ends.” *The Atlantic*, June 2018, <https://www.theatlantic.com/magazine/archive/2018/06/henry-kissinger-ai-could-mean-the-end-of-human-history/559124/> (accessed February 5, 2020).
- Leveson, Nancy G. “Are You Sure Your Software Will Not Kill Anyone?” *Communications of the ACM* 63, no. 2 (February 2020): 25-28.
- _____. *CAST Handbook: How to Learn More from Incidents and Accidents*. (2019), <http://sunnyday.mit.edu/CAST-Handbook.pdf> (accessed February 25, 2020).
- _____. *Engineering a Safer World: Systems Thinking Applied to Safety*. Cambridge: MIT Press, 2011.
- Leveson, Nancy G. and John P. Thomas. *STPA Handbook*. (March 2018), https://psas.scripts.mit.edu/home/get_file.php?name=STPA_handbook.pdf (accessed March 28, 2020).
- Lombard, Emmett and Vishal Arghode. “Information Literacy and Organizational Theory.” In *Information Literacy: Progress, Trends and Challenges*, edited by Luis Freeman, 114-120. New York: Nova Science Publishers, 2018.
- McLeary, Paul. “SecDef Eyeing Moving Billions By Eliminating Offices, Legacy Systems.” *Breaking Defense*, February 5, 2020, <https://breakingdefense.com/2020/02/secdefs-review-is-in-and-hes-willing-to-shut-down-entire-offices/> (accessed May 1, 2020).
- Molla, Rani. “Why Your Free Software Is Never Free.” *Vox*, January 29, 2020 (accessed May 15, 2020).
- Muro, Mark, Robert Maxim, Jacob Whiton, and Ian Hathaway. *Automation and Artificial Intelligence: How Machines Are Affecting People and Places*. Washington DC:

Brookings, January 2019, https://www.brookings.edu/wp-content/uploads/2019/01/2019.01_BrookingsMetro_Automation-AI-Workforce_Report_Muro-Maxim-Whiton.pdf (accessed January 17, 2020).

National Defense Industrial Association. “DoD Rewrite of 5000 Series to Include a Software Acquisition Pathway.” National Defense Industrial Association, July 26, 2019 <https://www.ndia.org/policy/recent-posts/2019/7/26/dod-rewrite-of-5000-series-to-include-a-software-acquisition-pathway> (accessed February 4, 2020).

_____. “Innovation Strategies.” National Defense Industrial Association, <https://www.ndia.org/policy/defense-innovation/innovation-strategies> (accessed January 7, 2020).

National Science and Technology Council. Networking and Information Technology Research and Development Subcommittee. *The Artificial Intelligence Research and Development Plan*. Washington DC: Government Printing Office, October 2016, https://www.nitrd.gov/pubs/national_ai_rd_strategic_plan.pdf (accessed February 5, 2020).

_____. Select Committee on Artificial Intelligence. *The National Artificial Intelligence Research and Development Strategic Plan: 2019 Update*. Washington DC: Government Printing Office, June 2019, <https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf> (accessed February 5, 2020).

National Security Commission on Artificial Intelligence. *Interim Report*. Washington DC: Government Printing Office, November 2019, <https://www.nscai.gov/reports> (accessed February 7, 2020).

National Transportation Safety Board. *China Airlines Boeing 747-SP, N4522V, 300 Nautical Miles Northwest of San Francisco, California, February 19, 1985*. NTSB/AAR-86/03, Washington DC: Government Printing Office, March 29, 1986, https://www.faa.gov/about/initiatives/maintenance_hf/library/documents/media/human_factors_maintenance/china_airlines_boeing_747-sp.n4522v.300_nautical_miles_northwest_of_san_francisco.california.february_19.1985.pdf (accessed March 23, 2020).

_____. “Driver Errors, Overreliance on Automation, Lack of Safeguard, Led to Fatal Tesla Crash.” National Transportation Safety Board, September 12, 2017 <https://www.nts.gov/news/press-releases/Pages/PR20170912.aspx> (accessed March 7, 2020).

_____. “Tesla Crash Investigation Yields 9 NTSB Safety Recommendation.” National Transportation Safety Board, February 25, 2020 <https://www.nts.gov/news/press-releases/Pages/NR20200225.aspx> (accessed March 7, 2020).

_____. *Collision Between a Car Operating With Automated Vehicle Control Systems and a Tractor-Semitrailer Truck Near Williston, Florida, May 7, 2016*. NTSB/HAR-17/02. Washington DC: Government Printing Office, September 12, 2017, <https://www.nts.gov/investigations/AccidentReports/Reports/HAR1702.pdf> (accessed May 4, 2020).

- _____. *Collision Between a Sport Utility Vehicle Operating With Partial Driving Automation and a Crash Attenuator*, Mountain View, California, March 23, 2018. NTSB/HAR-20/01. Washington DC: Government Printing Office, February 25, 2020, <https://www.nts.gov/investigations/AccidentReports/Reports/HAR2001.pdf> (accessed May 4, 2020);
- Newell, Allen. *Intellectual Issues in the History of Artificial Intelligence*. Pittsburgh: Carnegie-Mellon University, 1982
<https://apps.dtic.mil/dtic/tr/fulltext/u2/a125318.pdf> (accessed May 4, 2020).
- Ng Andrew. “AI Transformation Playbook: How to Lead Your Company into the AI Era.” Landing AI,
https://d6hi0znd7umn4.cloudfront.net/content/uploads/2019/11/LandingAI_Transformation_Playbook_11-19.pdf (accessed January 30, 2020).
- NOVA. “Look Who’s Driving.” Season 46, episode 19 (originally aired October 23, 2019).
- Nurkin, Tate and Stephen Rodriguez. *A Candle in the Dark: U.S. National Security Strategy for Artificial Intelligence*. Washington DC: Atlantic Council, 10 December 2019.
- Nyce, Caroline Mimbs. “The Winter Getaway That Turned the Software World Upside Down.” *The Atlantic*, December 8, 2017
<https://www.theatlantic.com/technology/archive/2017/12/agile-manifesto-a-history/547715/> (accessed February 7, 2020).
- Oliver, Nick, Thomas Calvard, and Kristina Potočník. “The Tragic Crash of Flight AF447 Shows the Unlikely But Catastrophic Consequences of Automation.” *Harvard Business Review*, September 15, 2017, <https://hbr.org/2017/09/the-tragic-crash-of-flight-af447-shows-the-unlikely-but-catastrophic-consequences-of-automation> (accessed March 23, 2020).
- Osinga, Frans. *Science, Strategy, and War: The Strategic Theory of John Boyd*. New York: Routledge, 2007.
- Perry, Georgia. “The AI Cold War That Threatens Us All.” *Wired*, October 23, 2018, <https://www.wired.com/story/ai-cold-war-china-could-doom-us-all/> (accessed February 5, 2020).
- Phillips, Sarah. “A Brief History of Facebook.” *The Guardian*, July 25, 2007, <https://www.theguardian.com/technology/2007/jul/25/media.newmedia> (accessed February 5, 2020).
- Pittet, Sten. “Continuous Integration vs. Continuous Delivery vs. Continuous Deployment.” Atlassian. <https://www.atlassian.com/continuous-delivery/principles/continuous-integration-vs-delivery-vs-deployment> (accessed February 23, 2020).
- Rehkopf, Max. “What Is Continuous Integration.” Atlassian.
<https://www.atlassian.com/continuous-delivery/continuous-integration> (accessed May 15, 2020).

- Roff, Heather M. "The Frame Problem: The AI 'arms race' Isn't One." *Bulletin of the Atomic Scientists* 75, no. 3 (2019): 95-98.
- Sanches, Tatiana, Carlos Lopes, and Maria da Luz Antunes. "Education and Psychology Trends: Impact on Information Literacy." In *Information Literacy: Progress, Trends and Challenges*, edited by Luis Freeman. New York: Nova Science Publishers, 2018.
- Sayler, Kelley M. *Artificial Intelligence and National Security*. CRS Report No. R45178 Version 7. Washington DC: Congressional Research Service, 2019.
- Scharre, Paul. "Between A Roomba and a Terminator: What is Autonomy?" *War on the Rocks*, February 18, 2015, <https://warontherocks.com/2015/02/between-a-roomba-and-a-terminator-what-is-autonomy/> (accessed January 2, 2020).
- _____. "Killer Apps: The Real Dangers of an AI Arms Race." *Foreign Affairs* 98, no. 3 (May/June 2019): 135-138.
- _____. *Army of None: Autonomous Weapons and the Future of War*. New York: W.W. Norton & Company, Inc., 2018.
- Schmidhuber, Jurgen. *Deep Learning in Neural Networks: An Overview*. Switzerland: University of Lugano, October 8, 2014, <https://arxiv.org/pdf/1404.7828.pdf> (accessed February 10, 2020).
- Schwartz, Oscar. "In 2016, Microsoft's Racist Revealed the Dangers of Online Conversation." *IEEE Spectrum*, November 25, 2019, <https://spectrum.ieee.org/tech-talk/artificial-intelligence/machine-learning/in-2016-microsofts-racist-chatbot-revealed-the-dangers-of-online-conversation> (accessed May 4, 2020).
- Science News Staff. "From AI to Protein Folding: Our Breakthrough Runners-Up." *Science*, December 22, 2016, <https://www.sciencemag.org/news/2016/12/ai-protein-folding-our-breakthrough-runners> (accessed February 5, 2020).
- Section 809 Panel. *Report of the Advisory Panel on Streamlining and Codifying Acquisition Regulations*, 2. Washington DC: Government Printing Office, June 2018, <https://discover.dtic.mil/section-809-panel/> (accessed February 25, 2020).
- Silver, David, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepe, and Demis Hassabis. "Mastering the Game of Go with Deep Neural Networks and Tree Search." *Nature* 529 (January 28, 2016): 484-489.
- Simonite, Tom. "Pentagon Will Expand AI Project Prompting Protests at Google." May 29, 2018, <https://www.wired.com/story/googles-contentious-pentagon-project-is-likely-to-expand/> (accessed January 12, 2020).
- Smith, Chris, Brian McGuire, Ting Huang, and Gary Yang. "The History of Artificial Intelligence." History of Computing CSEP590A. University of Washington, December 2006 <https://courses.cs.washington.edu/courses/csep590/06au/projects/history-ai.pdf> (accessed May 4, 2020).

- Society of Automotive Engineers. "Taxonomy and Definition for Terms Related to Driving Automation Systems for On-Road Motor Vehicles." Society of Automotive Engineers, https://www.sae.org/standards/content/j3016_201806/ (accessed March 26, 2020).
- Somers, James. "The Coming Software Apocalypse." *The Atlantic*, September 26, 2017, <https://www.theatlantic.com/technology/archive/2017/09/saving-the-world-from-code/540393/> (accessed April 10, 2020).
- Springer, Paul J. *Outsourcing War to Machines: The Military Robotics Revolution*. Santa Barbara: Praeger, 2018.
- Sutton, Richard S. and Andrew G. Barto. *Reinforcement Learning: An Introduction*. 2nd ed. Cambridge: MIT Press, 2018.
- Sydney J. Freedberg, Jr., "Should We Ban 'Killer Robots'? Can We?" *Breaking Defense*, March 11, 2019, <https://breakingdefense.com/2019/03/should-we-ban-killer-robots-can-we/> (accessed February 5, 2020).
- Tarraf, Danielle C., William Shelton, Edward Parker, Brien Alkire, Diana Gehlhaus Carew, Justin Grana, Alexis Levedahl, Jasmin Leveille, Jared Mondschein, James Ryseff, Ali Wyne, Dan Elinoff, Edward Geist, Benjamin N. Harris, Eric Hui, Cedric Kenney, Sydne Newberry, Chandler Sachs, Peter Schirmer, Danielle Schlang, Victoria M. Smith, Abbie Tingstad, Padmaja Vedula, and Kristin Warren. *The Department of Defense Posture for Artificial Intelligence*. Santa Monica: RAND Corporation, 2019.
- Telley, Chris. "Info Ops Officer Offers Artificial Intelligence Roadmap." *Breaking Defense*, July 11, 2017 <https://breakingdefense.com/2017/07/info-ops-officer-offers-artificial-intelligence-roadmap/> (accessed May 3, 2020).
- Trent, Stoney and Scott Lathrop. "A Primer on Artificial Intelligence for Military Leaders." *Small Wars Journal*, <https://smallwarsjournal.com/jrnl/art/primer-artificial-intelligence-military-leaders> (accessed February 25, 2020).
- Turing, A.M., "Computing Machinery and Intelligence." *Mind* 49 (1950): 433-460, <https://www.csee.umbc.edu/courses/471/papers/turing.pdf> (accessed May 4, 2020).
- U.S. Air Force. Chief Software Officer. "Preferred Agile Framework." December 28, 2019, <https://software.af.mil/wp-content/uploads/2019/12/CSO-MFR-on-Agile-Frameworks-12282019.pdf> (accessed February 25, 2020).
- _____. "DoD Enterprise DevSecOps Initiative (Software Factory) v4.7." By Nicolas Chaillan, <https://software.af.mil/dsop/documents/> (accessed February 23, 2020).
- _____. Office of the Chief Scientist. *Autonomous Horizons: The Way Forward*. By Greg L. Zacharias. Maxwell AFB: Air University Press, March 2019.
- U.S. Congress. *John S. McCain National Defense Authorization Act for Fiscal Year 2019*. Public Law 115-232, 115th Congress, 2nd Session (August 13, 2018), § 238, <https://www.congress.gov/115/plaws/publ232/PLAW-115publ232.pdf> (accessed January 23, 2020).

- U.S. Department of Defense. *Autonomy in Weapon Systems*. DoD Directive 3000.09. Washington DC: Government Printing Office, November 21, 2012, Incorporating Change 1, May 8, 2017, <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/300009p.pdf> (accessed January 2, 2020).
-
- _____. Chief Information Officer. “Updated Guidance on the Acquisition and Use of Commercial Cloud Computing Services.” December 15, 2014, https://dodcio.defense.gov/portals/0/documents/cloud/dod%20cio%20-%20updated%20guidance%20-%20acquisition%20and%20use%20of%20commercial%20cloud%20services_20141215.pdf (accessed February 25, 2020).
-
- _____. Chief Information Officer. *DoD Enterprise DevSecOps Reference Design Version 1.0*. Washington DC: Government Printing Office, August 12, 2019, https://dodcio.defense.gov/Portals/0/Documents/DoD%20Enterprise%20DevSecOps%20Reference%20Design%20v1.0_Public%20Release.pdf?ver=2019-09-26-115824-583 (accessed February 25, 2020).
-
- _____. Defense Information Systems Agency. “DISA Developed Application Chosen to Consolidate Several Air Force Aircraft Maintenance Systems.” Defense Information Systems Agency, October 21, 2019 <https://disa.mil/NewsandEvents/2019/application-consolidate-Air-Force-maintenance-systems> (accessed February 25, 2020).
-
- _____. Defense Innovation Board. *AI Principles: Recommendation on the Ethical Use of Artificial Intelligence by the Department of Defense*. Washington DC: Government Printing Office, October 31, 2019, https://media.defense.gov/2019/Oct/31/2002204458/-1/-1/0/DIB_AI_PRINCIPLES_PRIMARY_DOCUMENT.PDF (accessed March 13, 2020).
-
- _____. Defense Innovation Board. *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*. Washington DC: Government Printing Office, May 3, 2019, https://media.defense.gov/2019/Apr/30/2002124828/-1/-1/0/SOFTWAREISNEVERDONE_REFACTORINGTHEACQUISITIONCODEFORCOMPETITIVEADVANTAGE_FINAL.SWAP.REPORT.PDF (accessed March 26, 2020).
-
- _____. Defense Logistics Agency. “Information Operations (J6),” Defense Logistics Agency, <https://www.dla.mil/HQ/InformationOperations/Offers/Products/LogisticsApplications/FEDLOG.aspx> (accessed February 24, 2020).
-
- _____. Defense Science Board. *Summer Study on Autonomy*. Washington DC: Government Printing Office, June 2016.
-
- _____. Defense Science Board. *The Role of Autonomy in DoD Systems*. Washington DC: Government Printing Office, July 2012.

-
- . *National Defense Strategy*. Washington DC: Government Printing Office, June 2008, <https://archive.defense.gov/pubs/2008NationalDefenseStrategy.pdf> (accessed February 2, 2020).
-
- . *National Defense Strategy*. Washington DC: Government Printing Office, 2018, <https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf> (accessed December 23, 2020).
-
- . *Quadrennial Defense Review Strategy*. Washington DC: Government Printing Office, 2014, https://archive.defense.gov/pubs/2014_Quadrennial_Defense_Review.pdf (accessed February 2, 2020).
-
- . *Secretary of Defense Speech: Reagan National Defense Forum Keynote*. Washington DC: Government Printing Office, 15 November 2014, <https://www.defense.gov/Newsroom/Speeches/Speech/Article/606635/> (accessed February 2, 2020).
-
- . *Summary of the 2018 Department of Defense Artificial Intelligence Strategy*. Washington DC: Government Printing Office, 2018, <https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMARY-OF-DOD-AI-STRATEGY.PDF> (accessed November 7, 2019).
-
- . Under Secretary of Defense for Acquisition and Sustainment. “Software Acquisition Pathway Interim Policy and Procedures.” January 3, 2020, [https://www.acq.osd.mil/ae/assets/docs/USA002825-19%20Signed%20Memo%20\(Software\).pdf](https://www.acq.osd.mil/ae/assets/docs/USA002825-19%20Signed%20Memo%20(Software).pdf) (accessed February 25, 2020).
-
- . Under Secretary of Defense for Acquisition and Sustainment. *Operation of the Defense Acquisition System*. DoD Instruction 5000.02T. Washington DC, Government Printing Office, January 7, 2020, Incorporating Change 6, January 23, 2020, <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500002T.PDF?ver=2020-01-24-100028-310> (accessed February 4, 2020).
-
- . Under Secretary of Defense for Acquisition and Sustainment. *Operational of the Adaptive Acquisition Framework*. DoD Instruction 5000.02. Washington DC, Government Printing Office, January 23, 2020, <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500002p.pdf?ver=2020-01-23-144114-093> (accessed February 4, 2020).
-
- . Under Secretary of Defense for Intelligence & Security. “Disruption in UAS: The Algorithmic Warfare Cross-Functional Team (Project Maven).” By Lieutenant General Jack Shanahan, <http://airpower.airforce.gov.au/APDC/media/Events-Media/RAAF%20AP%20CONF%202018/1130-1200-Shanahan-Disruption-in-UAS-The-AWCFT.pdf> (accessed January 12, 2020).
-
- . Under Secretary of Defense of Research & Engineering. *DoD Modeling and Simulation (M&S) Management*. DoDD 5000.59. Washington DC: Government Printing Office, August 8, 2007, Incorporating Change 1, October 15, 2018), <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/500059p.pdf> (accessed April 10, 2020).

- U.S. General Accounting Office, *Patriot Missile Defense: Software Problem Led to System Failure at Dhahran, Saudi Arabia*. Washington DC: Government Printing Office, February 1992, <https://www.gao.gov/products/IMTEC-92-26> (accessed April 8, 2020).
- _____. *Software Development: Effective Practices and Federal Challenges in Applying Agile Methods*. Washington DC: Government Printing Office, July 2012, <https://www.gao.gov/assets/600/593091.pdf> (accessed March 8, 2020).
- U.S. House Armed Services Committee. “Statement of Dr. Eric Schmidt.” *Promoting DoD’s Culture of Innovation*, April 17, 2018, <https://docs.house.gov/meetings/AS/AS00/20180417/108132/HHRG-115-AS00-Wstate-SchmidtE-20180417.pdf> (accessed May 2, 2020).
- U.S. Joint Chiefs of Staff. *Capstone Concept for Joint Operations Joint Force 2030 (Unclassified)*. Washington DC: Joint Chiefs of Staff, June 18, 2019.
- _____. *Charter of the Joint Requirements Oversight Council (JROC) and Implementation of the Joint Capabilities Integration and Development System (JCIDS)*. CJCSI 5123.01H. Washington DC: Joint Chiefs of Staff, August 31, 2018, <https://www.jcs.mil/Portals/36/Documents/Library/Instructions/CJCSI%205123.01H.pdf?ver=2018-10-26-163922-137> (accessed March 1, 2020).
- U.S. Navy. Chief Information Officer. “Acquisition and Use of Commercial Cloud Computing Services.” May 15, 2015, <https://www.doncio.navy.mil/TagResults.aspx?ID=104> (accessed February 25, 2020).
- U.S. President. Executive Order 13859. “Maintaining American Leadership in Artificial Intelligence.” *Code of Federal Regulations*, title 3 (February 11, 2019) <https://www.govinfo.gov/content/pkg/FR-2019-02-14/pdf/2019-02544.pdf> (accessed February 1, 2020).
- USC Marshall School of Business. “What Exactly Is Information Literacy And What Role Does It Play In Education.” University of Southern California, <https://librarysciencedegree.usc.edu/blog/what-exactly-is-information-literacy-and-what-role-does-it-play-in-education/> (accessed May 4, 2020).
- Warwick, Graham. “DARPA Automated Dogfighting to Develop Pilot Trust in AI in Combat.” *Aviation Week and Space Technology*. April 20-May 3, 2020.
- Weisgerber, Marcus. “The Pentagon’s New Algorithmic Warfare Cell Gets Its First Mission: Hunt ISIS.” *Defense One*, May 14, 2017, <https://www.defenseone.com/technology/2017/05/pentagons-new-algorithmic-warfare-cell-gets-its-first-mission-hunt-isis/137833/> (accessed January 12, 2020).
- _____. “The U.S. Air Force Is Adding Algorithms to Predict When Planes Will Break.” *Defense One*, May 15, 2018, <https://www.defenseone.com/business/2018/05/us-air-force-adding-algorithms-predict-when-planes-will-break/148234/> (accessed January 2, 2020).
- Woyke, Elizabeth. *The Smartphone: Anatomy of an Industry*. New York: The New Press, 2014.