



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

MBA PROFESSIONAL REPORT

ARTIFICIAL INTELLIGENCE: THE BUMPY PATH THROUGH DEFENSE ACQUISITION

December 2017

By: Eric J. Ehn

Advisors: John Dillard
Robert Mortlock

Approved for public release. Distribution is unlimited.

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, 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 this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE December 2017	3. REPORT TYPE AND DATES COVERED MBA professional report		
4. TITLE AND SUBTITLE ARTIFICIAL INTELLIGENCE: THE BUMPY PATH THROUGH DEFENSE ACQUISITION			5. FUNDING NUMBERS	
6. AUTHOR(S) Eric J. Ehn				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB number ___N/A___.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) The use of artificial intelligence systems is ready to transition from basic science research and a blooming commercial industry to strategic implementation in the Defense Acquisition system. The purpose of this research is to determine the problems awaiting artificial intelligence (AI) systems inherent to defense acquisition. AI is a field of scientific study focused on the construction of systems that can act rationally, behave humanly, and adapt. To achieve AI behavior takes AI essentials, which consider mobility, system perspective, and algorithms. Unfortunately, AI essentials are under addressed in the concept of operations that fuels the Joint Capabilities Integration and Development System. Influences to the concept of operations analyzed in this research include strategic documentation, joint technology demonstrations, and exercises that aim to capture technology-based lessons learned. Failure to address AI essentials causes problems in defense acquisition: system requirements are impossible to define; transition of AI technology fails; testing cannot be evaluated with confidence; and life cycle planning is at best a guess. To address these issues, the Department of Defense needs improved planning, acquisition personnel training, and AI-supported acquisition processes to achieve cost, schedule, and performance goals.				
14. SUBJECT TERMS artificial intelligence, AI, autonomous, autonomic, acquisition, validated requirements, technology demonstration, technology transition, requirements analysis, test and evaluation			15. NUMBER OF PAGES 145	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

**ARTIFICIAL INTELLIGENCE: THE BUMPY PATH THROUGH DEFENSE
ACQUISITION**

Eric J. Ehn, Captain, United States Air Force

Submitted in partial fulfillment of the requirements for the degree of

MASTER OF BUSINESS ADMINISTRATION

from the

**NAVAL POSTGRADUATE SCHOOL
December 2017**

Approved by: John Dillard, Col (Ret)

Robert Mortlock

Keith Snider
Academic Associate
Graduate School of Business and Public Policy

THIS PAGE INTENTIONALLY LEFT BLANK

ARTIFICIAL INTELLIGENCE: THE BUMPY PATH THROUGH DEFENSE ACQUISITION

ABSTRACT

The use of artificial intelligence systems is ready to transition from basic science research and a blooming commercial industry to strategic implementation in the Defense Acquisition system. The purpose of this research is to determine the problems awaiting artificial intelligence (AI) systems inherent to defense acquisition. AI is a field of scientific study focused on the construction of systems that can act rationally, behave humanly, and adapt. To achieve AI behavior takes AI essentials, which consider mobility, system perspective, and algorithms. Unfortunately, AI essentials are under addressed in the concept of operations that fuels the Joint Capabilities Integration and Development System. Influences to the concept of operations analyzed in this research include strategic documentation, joint technology demonstrations, and exercises that aim to capture technology-based lessons learned. Failure to address AI essentials causes problems in defense acquisition: system requirements are impossible to define; transition of AI technology fails; testing cannot be evaluated with confidence; and life cycle planning is at best a guess. To address these issues, the Department of Defense needs improved planning, acquisition personnel training, and AI-supported acquisition processes to achieve cost, schedule, and performance goals.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	BACKGROUND	1
B.	PURPOSE	2
C.	RESEARCH QUESTIONS	3
1.	Primary Research Question	3
2.	Secondary Research Questions.....	4
3.	Scope.....	5
4.	Organization of the Study	5
II.	LITERATURE REVIEW	7
A.	DEFINING AI.....	7
1.	History of AI Definitions	7
2.	The Spectrum of Autonomous Behavior	9
3.	Defining Modern AI Systems.....	11
B.	HOW ARTIFICIAL INTELLIGENCE WORKS	14
1.	What are AI Systems Made of?	14
2.	Algorithms 101	20
3.	Basic AI Techniques	23
C.	STATE OF THE AI COMMUNITY	38
1.	The Domain of AI Research.....	38
2.	AI Industry Leaders	46
3.	Commercial AI Products.....	49
D.	CHAPTER II SUMMARY.....	52
III.	JCIDS.....	53
A.	JCIDS PRIMER.....	53
1.	Inputs	54
2.	Generating Requirements	55
3.	Requirements Documents	55
B.	METHOD	56
C.	ANALYSIS	62
1.	AI in Strategic Documents	62
2.	AI in Exercises.....	72
D.	CHAPTER SUMMARY.....	81
IV.	DAS.....	83
A.	INTRODUCTION.....	83

B.	DAS PRIMER	83
C.	DAS FAILURES	89
	1. System Requirements Generation	90
	2. Technology Transition.....	94
	3. Test and Evaluation	100
	4. Life Cycle Sustainment.....	102
D.	CHAPTER SUMMARY.....	105
V.	CONCLUSION	107
A.	INTRODUCTION.....	107
B.	SECONDARY RESEARCH QUESTIONS.....	107
C.	PRIMARY RESEARCH QUESTION.....	110
D.	RECOMMENDATIONS.....	112
E.	RECOMMENDATIONS FOR FUTURE RESEARCH.....	115
	LIST OF REFERENCES	117
	INITIAL DISTRIBUTION LIST	125

LIST OF FIGURES

Figure 1.	Matrix Displaying Categorized Definitions of AI. Source: Russell & Norvig (2010).....	13
Figure 2.	Representations of the Basic Elements of an AI system. Source: Jones (2009).....	16
Figure 3.	Relationship between Three AI System Classes and Attributes. Source: Truszkowski et al. (2009).....	19
Figure 4.	Basic Computer Algorithm. Source: Frenzel (1987).....	21
Figure 5.	AI Algorithm. Source: Frenzel (1987).....	22
Figure 6.	“Tower of Hanoi” Decision Space. Source: Jones (2009).....	24
Figure 7.	First Two Moves for the Eight Game. Source: Jones (2009).....	27
Figure 8.	Static State Machine for a Simple AI Enemy	29
Figure 9.	Example of a Semantic Network. Source: Jones (2009).....	30
Figure 10.	Process for Mutation in Genetic Algorithms. Source: Jones (2009).....	34
Figure 11.	Singular Neuron in a Neural Network. Source: Jones (2009).....	36
Figure 12.	Possible Training Aid for Neural Networks. Source: Jones (2009).....	37
Figure 13.	Overview of JCIDS. Source: CJCS (2012).....	54
Figure 14.	Illustration of a Basic DAS Process, Complete with JCIDS Connections. Source: DOD (2017).....	84
Figure 15.	Incrementally Deployed Software Intensive Program. Source: DOD (2017).....	87
Figure 16.	System Requirements Generation Process. Source: Haskins et al. (2011).....	92
Figure 17.	GAO Chart Displaying Requirements Growth. Source: Sullivan (2015).....	93
Figure 18.	High-Tech Marketing Model. Source: Moore (2014).....	95

Figure 19.	TDTS Stage Gates Process. Source: Kropas-Hughes, Rutledge, & Sarmiento (2008).	98
Figure 20.	Flow Diagram from Requirement to Test Report. Source: DOD (2012).	101
Figure 21.	Life Cycle Management throughout the DAS. Source: DAU (2013).	103

LIST OF TABLES

Table 1.	Sample Evaluation Chart Displaying Evaluation Criteria	59
Table 2.	Criterion Weighting	60
Table 3.	Scoring Legend	61
Table 4.	Analysis of the Joint Unmanned Systems Integration Roadmap.....	64
Table 5.	Analysis of <i>The U.S. Army Robotics and Autonomous Systems Strategy</i>	65
Table 6.	Analysis of Air Force <i>Technology Horizons</i>	67
Table 7.	Analysis of <i>A Cooperative Strategy for 21st Century Seapower</i>	69
Table 8.	Analysis of the 2014 DIA <i>Innovation Strategic Plan</i>	70
Table 9.	Analysis of <i>Marine Corps Operating Concept</i>	72
Table 10.	Analysis of <i>UQ 16 Future Force Design II Final Report</i>	74
Table 11.	MIX16 Analysis.....	75
Table 12.	Analysis of Red Flag Exercise 2016.....	77
Table 13.	Analysis of Tech Warrior.....	78
Table 14.	Analysis of ThunderDrone.....	79
Table 15.	Analysis of DARPA’s Cyber Grand Challenge.....	81
Table 16.	CONOPs Composite Score	82

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

3-D	Three-Dimensional
ACTUV	Anti-Submarine Warfare Continuous Trail Vessel
AI	Artificial Intelligence
AFRS	ALIS' Anomaly and Failure Reporting System
ALIS	Autonomic Logistic Information System
ARCIC	Army Capabilities Integration Center
ARFL	Air Force Research Laboratories
ASW	Anti-Submarine Warfare
CBA	Capabilities Based Assessment
CDD	Capabilities Development Document
CER	Concept of Exploration Refinement
CEO	Chief Executive Officer
CGC	Cyber Grand Challenge
CJCS	Chairman of the Joint Chiefs of Staff
CONOPs	Concept of Operations
COO	Chief Operating Officer
COTS	Commercial Off-The-Shelf
CPD	Capabilities Production Document
DARPA	Defense Advanced Research Projects Agency
DAS	Defense Acquisition System
DAU	Defense Acquisition University
DCR	Design Change Request/Document Change Request
DIA	Defense Intelligence Agency
DOD	Department of Defense
DODD	Department of Defense Directive
DODI	Department of Defense Instruction
DOTmLPPF-P	Doctrine, Organization, Training, Materiel, Leadership and Education, Personnel, Facilities, and Policy
EMD	Engineering and Manufacturing Development
FAIR	Facebook Artificial Intelligence Research Group
FDD	Full Deployment Decision
FD	Full Deployment
FPV	First Person View
GAO	General Accounting Office/Government Accountability Office
GPS	Global Positioning System

IBM	International Business Machines Incorporated
ICD	Initial Capabilities Document
INCOSE	International Council on Systems Engineering
IPPD	Integrated Product and Process Development
IPT	Integrated Project Team
IOC	Initial Operational Capability
JCIDS	Joint Capabilities Integration and Development System
JEON	Joint Emergent Operational Need
JROC	Joint Requirements Oversight Council
JUON	Joint Urgent Operational Need
KPP	Key Performance Parameter
KSA	Key System Attribute
LCMP	Life Cycle Management Plan
MAGTF	Marine Air-Ground Task Force
MDD	Material Development Decision
MIX16	Marine MAGTF Integrated Exercise 2016
MLP	Multilayer Perceptron
MRL	Manufacturing Readiness Level
MSA	Material Solutions Analysis
NASA	National Aeronautics and Space Administration
NDIA	National Defense Industrial Association
O&S	Operations and Support
ONR	Office of Naval Research
OPNAV	Office of the Chief of Naval Operations
OT&E	Operational Test and Evaluation
OV-1	Operational Viewpoint
PD	Production and Deployment
PM	Program Manager
PPBE	Planning, Programming, Budgeting, and Execution
R&D	Research and Development
RFP	Request for Proposal
S&T	Science and Technology
T&E	Test and Evaluation
TDTS	Technology Development and Transition Strategy
TEMP	Test and Evaluation Master Plan

TMRR	Technology Maturation and Risk Reduction
TRL	Technology Readiness Level
TTPs	Tactics, Techniques, and Procedures
UAS	Unmanned Aerial System
UAV	Unmanned Aerial Vehicle
UGV	Unmanned Ground Vehicle
UON	Urgent Operational Need
UQ	Unified Quest
UX	User Experience
USSOCOM	United States Special Operations Command
Wi-Fi	Wireless Networking

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

There are many people I want to thank for their support throughout this project. It certainly was a grind to condense the seminal aspects of a complex Defense Acquisition system and determine how a new crop of artificially intelligent systems is expected to interact with it. To those that helped me along the way, thank you.

I want to thank my wife, Abbey, for her support throughout this project. Her purchase of Amazon AI products to manage our in-house lighting situation was the seed that inspired this project. Her ability to manage our family throughout this project, while I was tucked away in a library, and, in fact, for the entirety of our military journey has given me the time and energy to advance Air Force mission needs as far as my talents can take them. I hope her efforts have led to a significant and helpful assessment of DOD artificial intelligence.

In addition, I also want to thank Col (Ret) John Dillard, my lead advisor. His willingness to shepherd this effort empowered me to take on a daunting subject. He pushed me to think deeply about key problems, offered me a map to follow when I felt lost, and presented me an opportunity to see firsthand the readiness level of AI research projects as they approach transition to DOD Service Research Laboratories.

Furthermore, I am grateful to Professor Robert Mortlock, who offered his advice long before he signed on as an advisor. His knowledge of the Defense Acquisition System and ability to assess effects of dynamic inputs to the DAS have served me well during this project and will again throughout my career.

To the rest of my instructors and professors, Air Force support staff, and colleagues who were always there with competent instruction and open minds to hear our ideas, I thank you. You all were more motivating and supportive than I deserve. Thank you for the friendship and comradery.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. BACKGROUND

On 14 February 2017, a collection of program managers (PMs) from Naval Postgraduate School (NPS), Defense Advanced Research Projects Agency, and Leidos, the prime contractor for the Anti-Submarine Warfare (ASW) Continuous Trail Unmanned Vessel (ACTUV) System, gathered with mission-requirement holders from around the Navy to discuss the acquisition future of the ACTUV. The crux of the effort was to assess the System's technology maturity, develop operational requirements through the Joint Capabilities Integration and Development System (JCIDS), and outline a funding and fielding path. The transition from its current state as a collection of basic science applications to a verified and validated system that meets Navy operational demands is a long process. While ACTUV was designed with anti-submarine applications in mind, at its core, it is a modular open system vehicle running an artificial intelligence (AI) brain that can be adapted to vast array of military and many other non-military applications. DARPA has published many capabilities for ACTUV, but one stands out. DARPA writes, "autonomous compliance with maritime laws and conventions for safe navigation, autonomous system management for operational reliability, and autonomous interactions with an intelligent adversary" (Littlefield, n.d.). The stated behavior of completely autonomous operation creates a new paradigm that removes humans from decision-making, rendering ACTUV a self-governing vehicle.

As the discussions evolved from the amazing capabilities of the ACTUV ship itself, an idea emerged that DARPA's effort with the ACTUV ship was nearly complete. In fact, the ship was christened and began on-water testing 7 April 2016 (Defense Advanced Research Projects Agency [DARPA], 2016). If we follow the acquisition timeline, once the results from ACTUV are collected and published, the next step would be to continue research at the basic science level or to push forward toward Office of Naval Research level projects (Department of Defense [DOD], 2017, p. 8). Warfighters should be paying attention to the results of demonstrators like ACTUV and considering how they could be used to enhance mission capability. These uses should match with

expected capability gaps and enter the Joint Capabilities Integration and Development System (JCIDS), which analyses future resource needs and defines what requirements a system should have to fill those gaps (Chairman of the Joint Chiefs of Staff [CJCS], 2012 p. 1). Similarly, the various resource owners of the Navy, specifically OPNAV Staff (N9I and N96) and Office of Naval Research (Dillard, 2017) for ACTUV, should be preparing the budgetary demand to fill the needs outlined in the JCIDS process. When the budgetary system, Defense Acquisition System (DAS), and JCIDS process are managed well, emerging technologies are developed into a viable defense system, delivered to the warfighter on time and at the cost planned by the project team. When managed poorly, the synergies between these systems are lost, often delivering low quality systems, behind schedule, and overrun funding profiles.

The next step for ACTUV depends on innumerable factors, but what is most evident is that its AI core technologies are ready to be matched with validated requirements produced inside of the JCIDS. Department of Defense (DOD) basic science initiatives and commercial entities such as IBM, Google, and Facebook are engaged in publishing peer reviewed scientific journals from their research and development departments that are pushing ahead the capabilities of AI and its various uses. Commercial products complete with cloud-based AI algorithms are cropping up, some specifically targeting insertion into fundamentally military domains. AI technology is ready to begin its journey from the technology maturation and risk reduction phase of DAS through to system fielding, but there remains an unanswered question (DOD, 2017, p. 8). Are the DAS and other acquisition support processes ready to incorporate AI technology? While, the ACTUV ship is not the focus of this project, it clearly portrays a situation that will be more and more common across service level and joint acquisition. AI systems are ready now and the DOD needs to be ready to harness this technology.

B. PURPOSE

The purpose of this research is to gain an understanding of what AI should mean to professionals in the DAS and JCIDS environments and anticipate what obstacles are expected to obstruct the flow of AI through the DAS. From an acquisition standpoint, AI

is not much more dynamic than what we know as software intensive (Department of Defense, 2017, p. 9), but behaviorally AI performs much differently. Its uses are so foreign to most people that AI's potential use in our daily military life is an abstraction. AI has the ability to function as the mind of a system or to optimize production schedules, which can work to eliminate the very person trying to procure it. The steps to developing AI are the same as already defined in the Department of Defense Instruction (DODI) 5000.02, but due to behavioral difference AI may work to amplify the problems already found within those acquisition steps. By outlining the challenges that AI will have throughout the DAS, it is possible to establish program structures that allow for quality management of AI systems.

C. RESEARCH QUESTIONS

To properly analyze the problems DOD will have getting a handle on AI technology, this research must ask pointed questions. These questions aim the research. They determine what information surrounding AI technology is crucial and what should be excluded. The overall effort of this research is to answer the primary research question. To do this properly, secondary research questions have been developed to ensure the entirety of the primary question is covered.

1. Primary Research Question

This project seeks to answer the following primary question:

What problems are AI based systems expecting to encounter as they transition from basic science to executable program?

In general, military acquisition transitions technologies through the DAS and engages with a validated output from the JCIDS in order to field a military system. AI is a unique technology that can both improve the end items purchased and increase the efficiency of the procurement process, but brings with it the current complexities of software acquisition and more. Military AI solutions are destined to travel through the DAS, and just as systems engineering seeks to flesh out risks early in the project, it is wise to consider the risks that AI brings to defense acquisition.

2. Secondary Research Questions

Before the primary research question can be sufficiently answered, this project must address three additional questions:

- What does AI mean for DOD acquisition and industry today?
- How well does the joint concept of operations account for AI technologies and how does that impact the JCIDS?
- If poor AI requirements are transitioned from JCIDS to DAS, what problems will AI systems encounter?

Understanding what AI is and how it behaves, in terms that a PM can understand is necessary to interpret the potential concerns inside of the JCIDS and DAS. If the maturity of AI thinking inside of the JCIDS does not keep pace with the improvements in the technology, the result will be poor AI requirements, which generally leads to poor program performance. Additionally, DAS often is tailored to the specific projects and industry dynamics go a long way toward determining what instruments inside of DAS to use. By defining in general terms the landscape of the AI industry, it can help determine what will factor into AI programmatic decisions

All defense programs start with a validated need, which is an output of the JCIDS process. JCIDS has many sources that provide direction including geo-political, strategic national interests, and even personnel constraints. A significant input to the JCIDS comes from strategic documentation found inside of national security and DOD agencies. Additionally, industry analysis, threat analysis, warfighter opinion, and imagination all play a role in influencing JCIDS efforts. The joint concept of operations (CONOPs) is the synthesis of these influences and supports the Capabilities Based Assessment (CBA) inside of JCIDS (CJCS, 2012, p. A-B-2). Additionally, the CONOPs lends its elements through the JCIDS and embeds in the requirement documents that inform DAS (CJCS, 2012, p. B-10). By analyzing CONOPs, this research seeks to determine if AI systems and technology are well represented in the efforts that activate JCIDS and lead to validated requirements.

Once a validated need is created inside the JCIDS, the DAS performs several processes to manage the system throughout its life cycle. These processes include systems requirements generation, technology transition, validation and verification of design, and life cycle sustainment. DAS is often criticized for its inability to deliver systems that meets cost, performance, and schedule agreements. Historically, these DAS functions have struggled with software-intensive system acquisition, therefore, using research performed on software acquisition performance can reveal problems that may be similar for AI.

3. Scope

This research focuses on assessing the landscape of AI study and technology, and applying it to the JCIDS and DAS. When analyzing JCIDS processes, this research focuses on the CBA and its ability to output validated needs given a robust joint CONOPs. Next, this research will analyze seminal processes inside of DAS that have historically failed to meet cost, schedule, and performance requirements during software-intensive system acquisition. This project will not address significant portions of JCIDS and DAS, and will not deal with the planning, programming, budgeting, and execution (PPBE) construct as it relates to defense acquisition. At the end of this research, defense acquisition personnel should be able to understand what AI is and how it behaves. Additionally, it should be clear how well the CONOPs represents AI and the consequences associated with transitioning poor AI requirements to DAS.

4. Organization of the Study

Chapter II, “Literature Review” provides a snapshot of AI today. It works to provide a general understanding of the scientific field and technology that is AI, the spectrum of behaviors expected from AI systems, what composition of an AI system, and the identities of AI industry leaders. The reader should be able to understand a working definition of AI systems, a general sense of AI technology readiness, and the emerging industry surrounding AI.

Next, Chapter III, “JCIDS,” examines the ability for DOD processes to develop requirements for AI applications. Requirements developments starts at a strategic level,

directing military resources to achieve present and future military needs. The JCIDS clarifies strategic direction, identifying capability gaps and validating needs (CJCS, 2012, p. 2). This chapter outlines how the JCIDS builds validated requirement documents, then focuses on grading AI elements in the joint CONOPs. The reader should leave this section with an understanding of CONOPs AI maturity and its influence on validated requirements headed for the DAS.

Chapter IV, “DAS,” focuses on the DAS and the processes that PMs use to manage system acquisition. The DAS is defined by DOD regulation, and gives direction for management of systems engineering, financial management, and contracting efforts (DOD, 2017, p. 51). This chapter analyzes the general process for developing and purchasing defense systems and the seminal areas inside of the DAS where software-intensive systems have struggled. The reader should leave this section understanding the consequences that poorly defined AI requirements would have on program cost, performance, and schedule.

Chapter V, “Conclusion,” integrates the ideas uncovered from the research in order to answer the secondary research questions and then the primary research question. Next, it makes recommendations based on the research that should help to prepare JCIDS and DAS for success with AI systems. Lastly, Chapter V proposes future areas of research that will generate more comprehensive information about the definition of AI requirements and how to meet cost, schedule, and performance during system fielding.

II. LITERATURE REVIEW

AI literature review is essential to defense acquisition personnel. AI will ultimately affect both the processes that acquisition systems follow and the systems that defense acquisition fields. The documents studied reveal, the history of defining AI and applies a modern assessment of what exactly AI is. Further these documents outline the basic components of AI and review the primary AI techniques. The final aspects uncovered by this documentation are the domain of AI research, key visionaries from the industry, and a view of emerging commercial AI products. The key takeaways from this section of research are a simple working definition of AI, an understanding of how AI is accomplished, and knowledge of the industry that will ultimately bring the technology into being.

A. DEFINING AI

1. History of AI Definitions

Even before the invention of the field of study known as AI, it was postulated that the original punch card computers could lead to AI (Crevier, 1995, p. 24). While many people contributed to the rise of this field, one of the best known is Alan Turing and his Turing Test. Author Daniel Crevier outlines the life work of Turing and his test.

Turing was a mathematical savant who, in 1950, outlined exactly how a computer, capable of operating in many different modes, could follow a sequenced set of steps and solve a nearly infinite number of problems (Crevier, 1995, pp. 23–25). Through the course of this effort, before a solid definition for AI appeared, Turing outlined a basic test that would determine the Intelligence of a machine. Crevier (1995) summarizes the key distillation from Turing’s 1950 paper, *Computing Machinery and Intelligence*, like this:

Suppose a machine was capable of answering any question you might put to it just as a human would. In fact, suppose you were communicating through a terminal with two hidden parties and couldn’t tell by questioning them which was human and which was a computer. Wouldn’t

you have to grant the computer this evasive quality we call intelligence.
(p. 24)

From the logic of this test, Turing predicted that thinking machines, capable of imitating human behaviors would arise by the year 2000 (Crevier, 1995, p. 24). Turing may have been off by a few years, but the logic proposed in the *Turing Test* has been the framework and aim point for much of the development of AI for nearly 70 years. This idea paints a compelling picture of the capability a system would need to exhibit to be considered AI. Unfortunately, Crevier's insights to Alan Turing do not spell out exactly the aspects someone should be looking for in an AI system. Many other scientists with similar thoughts and attempts to give computers life, pushed the science along until the scientific field of AI was established in 1956, at a conference at Dartmouth College (Russell & Norvig. 2010, p. 17). Professor John McCarthy PhD, along with other prominent AI researchers used the conference to establish the field of AI, apart from operations research and mathematics, to pursue duplicating human faculties like creativity and self-improvement in machines (Russell & Norvig. 2010, p. 18). AI was born, but still not easily defined.

Herbert A. Simon is widely considered a father of the conceptual thinking for AI research. In his 1985 address to NPS, he presented some of his founding thoughts about the definition of AI from the perspective of its early years. Simon (1985) starts by clarifying that intelligence can have two common definitions:

It can refer to information of significance to military operations and to the means for securing or analyzing it; second it can refer to the faculty of the human mind and brain that enables us to think and learn. It is the second meaning that was intended by the inventors of the label *artificial intelligence*. (p. 11)

Clearly this research is focused on Simon's second definition. Simon also explained, "we can say that artificial intelligence has been exhibited by a computer when it has done something that would have required intelligence in a man or a woman" (Simon, 1985, p. 12). While this statement is another good litmus test for whether or not a computer has intelligence, more information is required to define exactly what AI is.

Fortunately, Simon continued and outlined what he believes makes a system intelligent, whether that is biological or mechanical. Simon (1985) stated that

to be capable of thought and intelligence is that the system be a physical symbol system: that is, that it be able to input (read) symbols, output (write) symbols, create structures of symbols related in various ways, store symbols and symbol structures in memory, compare symbol structures for identity or difference, and branch (adapt its behavior) on the basis of the outcomes of such comparisons. (p. 14)

Simon later explains symbols as any pattern built out of any medium like ink or chalk, and including spectrums humans cannot perceive like magnetism, electricity, or neuron patterns (1985, p. 14). His point is that anything that can be observed, stored, recalled and related can work as a symbol and anything that can do those processes and change behavior based on those processes, is intelligent. With this, Simon offers the first usable definition of Artificial Intelligence, as something mechanical that can process symbols as a human would.

Unfortunately, these early attempts to focus the world on AI cause a contextual paradigm (Simon, 1985, p. 12). As the technology from a certain time period matures it moves slowly, but surely, from a confusing amalgamation of basic sciences and algorithms to an accepted and understood technology (Simon, 1985, p. 12). Moving forward to current time, not only is a computer chess game no longer thought of as an AI system, it has solidified itself as an outdated windows game. It still lives on the spectrum of AI, given that it is a program reading symbolled inputs from a user and responding with a chess move, but is no longer thought of as AI. Many other systems have met this same context based fate, so it is important to define a spectrum of AI as well.

2. The Spectrum of Autonomous Behavior

In 2009, the National Aeronautics and Space Administration (NASA) published *Autonomous and Autonomic Systems: with Applications to NASA Intelligent Spacecraft Operations and Exploration Systems*, which outlines the spectrum of AI. It starts on the less intelligent end of the spectrum and defines *automatic*. NASA offers, “automated processes simply replace routine manual processes with software/hardware ones, which

follow a step by step sequences that may still have a human in the loop” (Truskowski et al., 2009, p. 9). They utilize a computer program that combs through a database and outputs some analysis as example (Truskowski et al., 2009, p. 10). A program like this certainly replaces arduous human actions and is behaving a routine set of steps to accomplish the desired output. The system is mechanical in nature and by the NASA standard is an example of automated process, rather than AI.

The next level of intelligence on the spectrum is an *autonomous* system. NASA writes that, “autonomous processes...have the more ambitious goal of eliminating human processes” (Truskowski et al., 2009, p. 9). This important distinction between replacing human actions and replacing humans completely is what differentiates automatic and autonomous. NASA adds two important characteristics inside a system, “self-governance” and “self-direction” that add to the definition of autonomous (Truskowski et al., 2009, p. 10). These two traits imbue the system with a responsibility to adjust itself to meet the goals laid out for the system. There are many examples used by NASA to outline this point of “self-governance.” They describe a flight software program that can monitor key spacecraft health and safety data, understand what factors are a concern to the spacecraft, and independently take corrective actions necessary to maintain spacecraft health (Truskowski et al., 2009, p. 11). This clarifies that a system using “self-governance” to completely eliminate human processes, like interpretation of data, is an autonomous system. This zone of system behavior is still very much a part of AI, but there is another level to the modern spectrum.

Autonomy is the next generation of mechanical intelligence, which NASA calls “self-management” (Truskowski et al., 2009, p. 10). This ability includes the precepts for autonomous behavior, but takes them a step further and allows the system to prioritize for itself the autonomous processes to undertake (Truskowski et al., 2009, p. 11). Just as autonomous systems have two defining characteristics, NASA outlines eight characteristics that apply to Autonomic systems, divided into four self-managing properties and four enabling properties (Truskowski et al., 2009, p. 11). They list (Truskowski et al., 2009, p. 11)

- self-configuring

- self-healing
- self-optimizing
- self-protecting
- self-aware
- self-situated
- self-monitoring
- self-adjusting

It is the last enabling property that implies the most about autonomic systems, by taking what is already considered an intelligent system and allowing it to manage its priorities or branch into new uses for its allowable behavior. While NASA does not propose an example of what an autonomic system would look like, it does specify that the system would be performing a process similar to autonomous systems and then reflecting on both its mission and performance to aim for better results (Truskowski et al., 2009, p. 11). It is the autonomic definition from NASA that works most to update the definition of AI for this research, given the context of time and the consideration that autonomous systems are deployed in DOD today.

A final note concerning the AI spectrum is the possibility of using adjustable levels of autonomy. This NASA explains this as using operational context to define the level of autonomy used by the system (Truskowski et al., 2009, p. 17). This is useful when complete autonomy may not be desired or possible for certain scenarios (Truskowski et al., 2009, p. 17). Adjustment like this can be made by the system itself or dictated by the human controlling the system (Truskowski et al., 2009, p. 17). In DOD applications this would be very useful to limit the scope of a system or eliminate risk for untested scenarios.

3. Defining Modern AI Systems

Still, these references do not fully encompass the current field of AI as it stands today, due to their narrow focus on system attributes. There is more to AI than simply the software interfaces and processing ability. Webster's dictionary does a wonderful job of isolating two key elements of modern AI. They define AI as (Artificial Intelligence, n.d.):

1. a branch of computer science dealing with the simulation of intelligent behavior in computers

2. the capability of a machine to imitate intelligent human behavior

Both definitions are helpful from the perspective of defense systems, in that, they outline both the system that can be built and the field of study dedicated to achieving a capable AI system. Defense personnel will have to be versed in both definitions to understand the weapon systems produced based on a validated requirement, and the scientific efforts that go into making the system possible.

In their 2010 book, *Artificial Intelligence a Modern Approach*, authors Stuart Russell and Peter Norvig support and expand on Webster's definition of the field of AI. They explain AI as an effort to explain how "a mere handful of matter can perceive, understand, predict, and manipulate a world far larger and more complicated than itself" (Russell & Norvig, 2010, p. 1). They add, "[AI] goes further still: it attempts not just to understand, but also to *build* intelligent entities" (Russell & Norvig, 2010, p. 1). From Russell and Norvig's definitions we can see that for complete understanding, one must regard the science that an AI system is derived from and the behaviors that an AI product is capable of.

The only error made in these definitions is the statement that the field lives within computer science. While an AI system can often be found inside of a computer and many techniques for building AI system utilize computer science, it is actually a multidisciplinary science. Russell and Norvig correct this misconception, stating that the field of AI study includes, but is not limited to, Philosophy, Mathematics, Economics, Neuroscience, Psychology, Computer Engineering, Control Theory and Cybernetics, and Linguistics (2010, p. 5-16).

Russell and Norvig offer another level of granularity to this research by introducing a matrix that defines AI in two directions. They state that systems in the AI field are capable of thinking or acting and these thoughts or actions can be done either rationally or humanly (Russell & Norvig, 2010, p. 2). This means if a system is thinking humanly or rationally it is an AI system. Also, if it is acting humanly or rationally, it is an AI system. Under each category the authors use various definitions from around field of

AI to provide depth to each category (Russell & Norvig, 2010, p. 2). Figure 1 shows the summary of their work.

<p>Thinking Humanly</p> <p>“The exciting new effort to make computers think . . . <i>machines with minds</i>, in the full and literal sense.” (Haugeland, 1985)</p> <p>“[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . .” (Bellman, 1978)</p>	<p>Thinking Rationally</p> <p>“The study of mental faculties through the use of computational models.” (Charniak and McDermott, 1985)</p> <p>“The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)</p>
<p>Acting Humanly</p> <p>“The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990)</p> <p>“The study of how to make computers do things at which, at the moment, people are better.” (Rich and Knight, 1991)</p>	<p>Acting Rationally</p> <p>“Computational Intelligence is the study of the design of intelligent agents.” (Poole <i>et al.</i>, 1998)</p> <p>“AI . . . is concerned with intelligent behavior in artifacts.” (Nilsson, 1998)</p>

Figure 1. Matrix Displaying Categorized Definitions of AI.
Source: Russell & Norvig (2010).

The many definitions from Figure 1 outline the behaviors that we expect AI systems to exhibit and it helps us to add valuable adjectives to a defense acquisition working definition. Figure 1 also shines light to the notion that capturing a solid definition for AI in a simple statement is perhaps constraining to the overall AI effort.

Still, a definition is required as a baseline to perform any type of analysis pertaining to defense acquisition. Keep in mind this is not intended to be the best definition of AI, but a definition aimed at building the knowledge and framework necessary to work an AI system through the defense acquisition. At this point, we can aggregate the critical components and work out a usable definition of AI for defense systems:

Artificial Intelligence for Defense Systems: *a multidisciplinary scientific field that aims to study human intelligence, and the attempt to bring rational and human behaviors into a system that can think, act, and self-manage.*

While this may be a somewhat simplistic and concise for such complicated systems, it does the job of grounding defense acquisition personnel with the necessary terminology and scope necessary to field AI systems.

B. HOW ARTIFICIAL INTELLIGENCE WORKS

While a working definition for AI is important, it is an incomplete understanding for acquisition personnel. A simplistic sense of “how” AI works is also necessary for success in fielding systems. Very few people are familiar with the underpinnings of software architecture and the complexity of forming heuristics into code. Even systems engineers struggle to update themselves with the intricacies that have emerged in the software world, and the added complexity that AI algorithms bring to basic software. For this reason, a presentation of the basic tenets of AI are necessary. *AI essentials* will be the term used throughout this research to refer to the various attributes and factors that make AI systems work.

1. What are AI Systems Made of?

The modern structure of AI systems has many formulations across the field of AI. Some take a narrow focus on the embedded AI algorithms themselves that can live exclusively in a cloud. This look can be very effective and is the type of AI that is most often associated with search engines. In section *AI Algorithms 101* this research aims to explain how a simple AI algorithm works, but this sight picture is too narrow to explain how a completed AI system works. The AI algorithm is only but a piece of the larger context concerning many other interactions required for effective AI. In the book *Artificial Intelligence: A Systems Approach*, M. Tim Jones offers a simple analysis of an AI system.

While AI algorithms are important, “no algorithm is useful in isolation” (Jones, 2009, p. 13). This quote outlines that there is more to an AI system, even one that lives

online, than simply the AI algorithm. Jones (2009) paints the AI algorithm as the core of the system (p. 14), the instrument that allows the system to achieve rational and human behaviors. Jones (2009) labels systems engaging with their environment “sensing” (p. 15) and creating outputs “effecting” (p. 15). From these labels he created supplementary components to the AI algorithm called *sensors* and *effectors* (Jones, 2009, p.15). These two components, not necessarily physical, give the AI algorithm an, “understanding of the environment and also a way to manipulate the environment” (Jones, 2009, p.13).

Another aspect that must be considered for an effective AI System is the *environment*. Jones (2009) sees the environment as the object that makes an AI systems practical or grounds the system in the real world (p. 13). Forgetting about the specific environment would mean missing functionality and success, for any system developed. The final component that Jones references is *practice* (p. 13). While he doesn’t define what practice is, it infers knowledge of two points; learning is a behavior embedded in AI algorithms and time improves AI algorithm capabilities. Practice is not only beneficial, it is necessary, for AI systems because designers are incapable of predicting all possible situations a system may encounter in a given environment, designers are incapable of predicting the changes to the environment in the future, and designers are at times incapable of programming possible solutions to problems a system might encounter (Russell & Norvig. 2010, p. 693). Practice includes this missing knowledge and the time function that allows a system to evolve into the necessary state to meet the environment. It is very important to understand that practice is not a component that is usually associated with mechanical systems, but rather more common when referring to personnel. Figure 2 displays how Jones (2009) orients these pieces of an AI System to function together, and how they aim to engage the environment and practice (p.14).

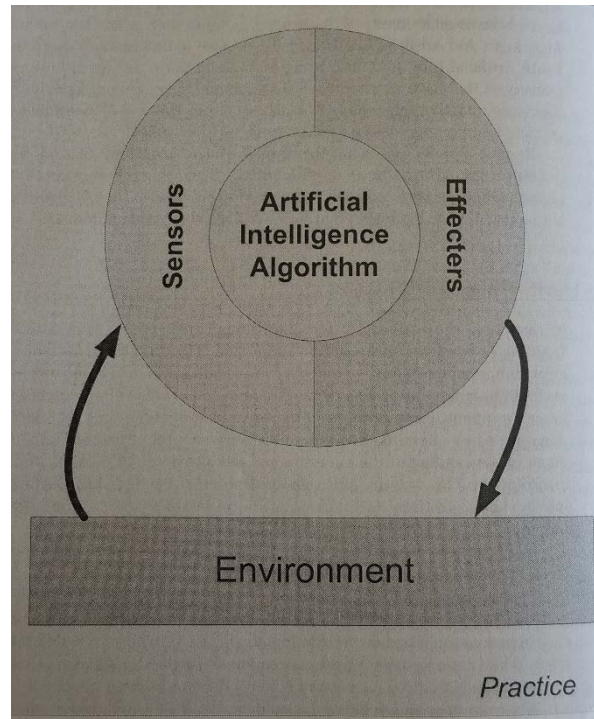


Figure 2. Representations of the Basic Elements of an AI system.
Source: Jones (2009).

AI has seen extensive use in games, dating back to the first computer games and currently in championship level chess matches, so they make for a great example when considering the proposed components of AI systems. Consider an online chess match, a zero-sum game, with a standard chess board and all of the appropriate pieces laid out on a computer screen. An AI system is cast as one of the players and lives exclusively in the computer software. An AI robot is the other player and is engaging the game through the same tools a person would use to manipulate a computer. The embedded AI system lives in the environment of a computer, complete with chess rules and computer function constraints. It must sense using mathematical calculations of pixels on the screen that are illuminated. Using the AI algorithm, it can interpret the position of the pieces based on the information retrieved from the sensors, and compute the logical next play. In the case of IBM's Deep Blue computer, it evaluates the chess board using four functions: material value [piece value], position, king safety, and tempo [how long it takes to gain an advantage] (IBM, 2001). Once it is satisfied of an optimized play, it must translate the

decision into a form the computer system can understand using the effectors. The computer system processes the play and adjusts the board.

Now it is the AI robot's turn, however its environment is vastly different. It is outside of the computer and is free from the influence of screen based pixel mathematics, but rather must interpret images and deal with gravity. It must be able to sense the image of the screen to understand the game. This requires some form of sensor that can translate the color differences on the screen into usable code for the AI algorithm. Once the necessary information is passed to the AI algorithm it too must decide on the best chess play given the options, but also must process how to engage the mouse or keyboard to input that decision. Once the proper play is determined, it must affect the computer using a physical interaction, based on understood mechanics and real world physics.

In each case the seminal principles are the same. There is an understood environment that requires sensing, processing, and effecting to achieve the desired goal of winning a chess game. While the physical components may be different, the system components that require consideration do not change. Additionally, the designer of each system will most likely not be capable of including every possible combination of chess moves and responses. Because of this it is better to define the roles of the AI system and let it practice, to include losing at chess, until it better understands how the game of chess should best be played. This creates a new paradigm for defense systems and would be digital chess champions alike; they will be completely built and verified by testing, but operational capability may be years away due to the needs of practice.

This example illustrates both the general makeup of AI systems and that AI systems in different environments will be composed of different things. NASA's book *Autonomous and Autonomic Systems: with Applications to NASA Intelligent Spacecraft Operations and Exploration Systems*, brings more granularity to M. Tim Jones' model for AI systems, in the way of attributes. These are the specific constructs that apply to AI systems, given different environments. They use the example of a specific type of AI system called Intelligent Agents, but for this portion of research it is working as a standard AI system (Truskowski et al., 2009, p. 17). NASA writes there are, "three

important classes: *software agents, robots, and immobots*” (Truskowski et al., 2009, p. 17). Each of these types of systems requires a different make up and attributes to achieve the sensing, AI algorithm functionality, and effecting aspect for AI systems.

A software agent is a completely digital system that pursues goals for their human owners (Truskowski et al., 2009, p. 17). Their specific environment allows them to be mobile across digital platforms and have a distributed array of digital sensors and digital actuators (Truskowski et al., 2009, p. 18). While their uses are wide, NASA uses the example of an information locator that receives a task, prioritizes gathered information digitally, and uses its effectors to present the information to its owner (Truskowski et al., 2009, p. 17).

A robot is a mobile system that pursues goals of the owner in the physical world (Truskowski et al., 2009, p. 18). While it does still need binary language to function, it is more concerned with measurements and using actuators to make physical changes (Truskowski et al., 2009, p. 18). NASA uses the example of the Sojourner robot that independently retrieved information from a database on Earth and took physical measurements of the environment on Mars (Truskowski et al., 2009, p. 18). A key difference from immobots, is that the focus of robots is outward, toward the environment at large (Truskowski et al., 2009, p. 18)

Immobots are physical systems that that manage a distributed network of physical sensors (Truskowski et al., 2009, p. 18). These are inwardly focused AI systems, generally with the goal to monitor and maintain general health of the overall system (Truskowski et al., 2009, p. 18). NASA uses the example of a modern factory floor, with sensors across the floor to monitor and control the efficiency of factory machines (Truskowski et al., 2009, p. 18). Figure 3 depicts the three classes and their biases in attributes.

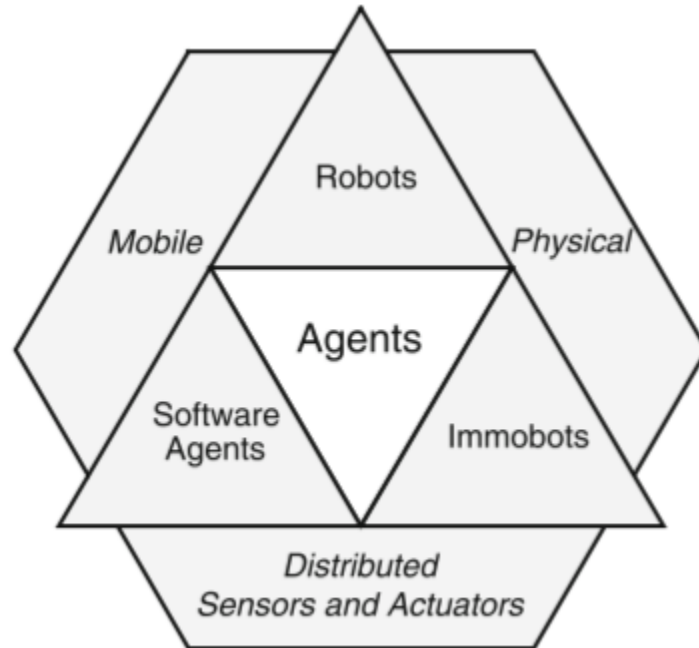


Figure 3. Relationship between Three AI System Classes and Attributes.
Source: Truskowski et al. (2009).

This may sound like the research has already moved past the question of this section, but the differences in environment define the classification for the system. NASA (Truskowski et al., 2009) offers the notion of attributes to help determine the appropriate class of AI system. The listed attributes are (Truskowski et al., 2009, p. 18):

- Purpose
- Domain of expertise
- Nature of sensors and actuators
- Mobility
- Physical or virtual
- How the domain is divided between systems
- How systems negotiate and cooperate
- Degree of cooperation
- Degree of individual identity

Consideration of these attributes and the result of those considerations, define how an AI system will accomplish its goals inside of a given environment. This method for defining the classes of AI, is one of many inside of the AI field, but serves well to inform acquisition personnel about the perspectives and mobility for AI systems. AI

components [sensors, AI algorithm, effectors, environment, practice] and AI classes [software agents, robots, immobots] are essential to AI concepts working through JCIDS and DAS.

2. Algorithms 101

At the heart of an AI system is the AI algorithm. Algorithms 101 seeks to define what an algorithm at the heart of an AI system does and simply how it functions. Noteworthy, is that an AI algorithm is a very different monster from a conventional computer algorithm. The difference between these two processes is stark, so for this research, basic algorithms or will serve as the definition for a common computer process and AI algorithm will serve to denote the combination of algorithms and knowledge bases.

In the book *Crash Course in Artificial Intelligence and Expert Systems*, author Louis E. Frenzel (1987) outlines exactly what a basic algorithm is when talking about computer programs. He writes that

conventional computer programs are based on an algorithm, a clearly defined, step-by-step procedure for solving a problem. It may be a mathematical formula or a clearly defined sequential procedure that will lead to a solution. The algorithm is converted into a computer program, a sequential list of instructions or commands that tell the computer exactly what operations to carry out. (p. 4)

From this definition, we can understand that algorithms are simply representations of the steps necessary to get to a desired solution. They are the sequences that allow a calculator to input into memory a number, recall that number, and apply operations to that number. Frenzel (1987) also explains that a basic algorithm uses data such as numbers, letters, or words to solve the problem using the pre-programmed operations (p. 4). Figure 4 displays a simple algorithmic processing nature of a conventional computer. Frenzel's definition is a clear, but equally telling is what it implies an algorithm cannot do. An algorithm cannot break the sequencing defined for it, nor can it compile the words, numbers, and symbols into anything other than what they are defined as.

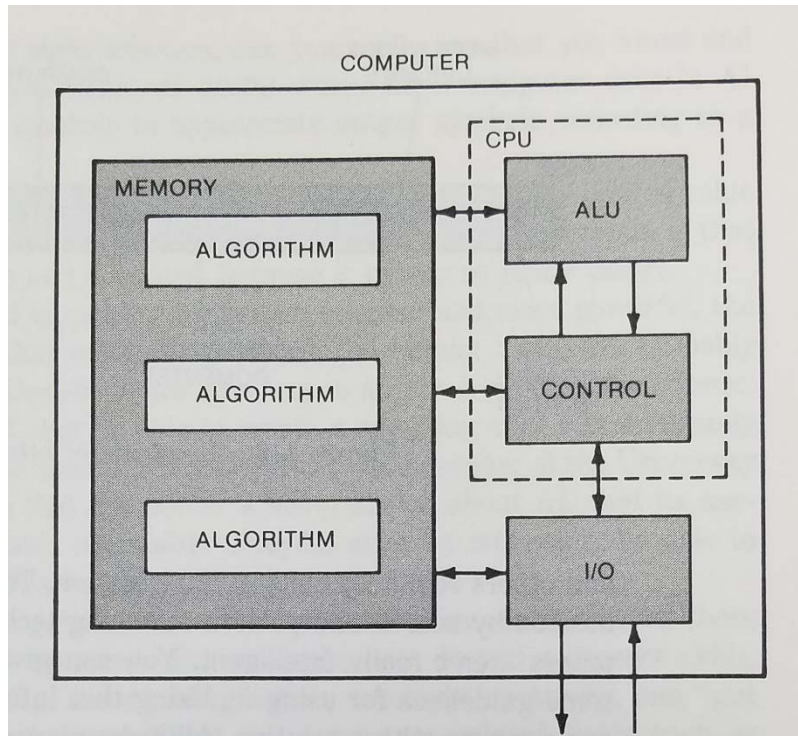


Figure 4. Basic Computer Algorithm. Source: Frenzel (1987).

In AI computing, an *AI algorithm* is used. To define how an AI algorithm works it is important to consider its two main parts: A particular set of basic algorithms and a *knowledge base* (Frenzel, 1987, p. 4). It is the cooperation between these two things that allow for abstraction in a computer.

It is important to not think of the knowledge base as an algorithm. It is a collection of symbols that gets subjected to algorithmic processes (Frenzel, 1987, p. 4). Frenzel (1987) calls this symbolic representation and manipulation (p. 4). He writes that “in AI a symbol is a letter, word, or number that is used to represent objects, processes, and their relationships” (p. 4). The magic here is that the things represented in symbol form take on a greater variety than those recognized in conventional computing. Objects can be ideas, events, statements of fact, or any other abstract concept that computers don’t usually use to operate. When the computer is holding all of these objects in its memory it is building the knowledge base (Frenzel, 1987, p. 4). A knowledge base is more useful than a sequence of steps, because it undergoes comparisons and contrasts,

resulting in an understanding of relationships between the symbols. Some of the symbols do contain algorithms, but the knowledge base is primarily the resulting information from symbol comparison.

Specific basic algorithms allow for the complex comparisons and contrasts necessary for knowledge bases to function. Search and pattern matching algorithms are the two key algorithms that Frenzel points out as most helpful (1987, p. 4). He writes, “the AI software searches the knowledge base looking for specific conditions or patterns. It looks for match ups that satisfy the criteria set up to solve the problem” (Frenzel, 1987, p. 4). This process of hunting delivers a best case response to the question, given the limits of the information stored in the knowledge base (Frenzel, 1987, p. 4). Figure 5 demonstrates the relationship between a knowledge base and AI specific basic algorithms. In this way the computer is using a simplified logic to solve problems in a manner that can be considered human like. It appears to behave intelligently, even if the response does not answer the question.

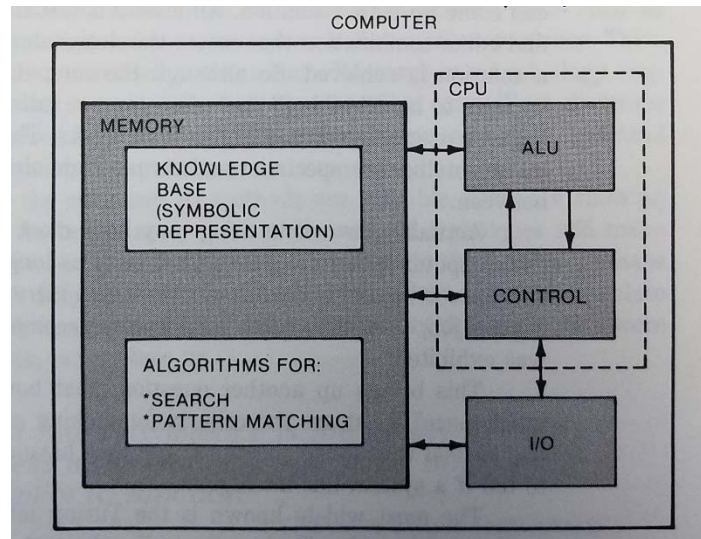


Figure 5. AI Algorithm. Source: Frenzel (1987).

The relationship between the knowledge base and the AI specific algorithms is what makes the AI processes different. While Frenzel shares two basic algorithms that

marry well with knowledge bases, there are actually several basic techniques that work leverage a knowledge base.

3. Basic AI Techniques

Without getting too deep into the specific mathematical or logical representations that empower the interaction between knowledge bases and algorithms, it is important to analyze the techniques that allow an AI algorithm to function. These are the mechanisms that designers use to make the knowledge base and algorithms work together. This can be thought of as the architecture for the AI algorithm. These techniques can be built completely in the software domain or a certain kind of hardware arrangement that enables the desired functionality.

The book *Artificial Intelligence: A Systems Approach*, by M. Tim Jones is again helpful. Mr. Jones (2009) explains seven general techniques that encompass the world of AI algorithms (pp. 21-268). Inside of each technique are groupings of several to hundreds of specific, and unique, AI techniques. It seems, no two authors on the subject of AI group the specific techniques in this same way or label the overarching techniques with the same words. After review of several publications, Mr. Jones' publication appears to be an approximate mean. It offers a central view of AI techniques. For the purpose of edification of defense acquisition personnel, the centrist, but still comprehensive, view is best.

a. Uninformed Search

Artificial Intelligence: A Systems Approach, starts with the technique of *uninformed search*. Uninformed search, or blind search is an AI algorithm that, “enumerates a problem space from an initial position in search of a goal position (or solution)” (Jones, 2009, p. 21). The problem space consists of a number of potential actions that the computer can take, and the associated consequences of those actions (Jones, 2009, p.22). Some of these consequences are that new actions are then available to the problem space and some of the consequences are that no more actions can be made. The goal for search is not necessarily to find a singular correct output, but rather a sequence of operations that transitions the system from the start position to the goal state

(Jones, 2009, p. 22). Search strategies are numerous and specific search techniques are available to fit many different environments and different system attributes.

One effective way to analyze the logic of uninformed search is through the use of tree diagrams (Jones, 2009, p. 27). Figure 6 shows a tree diagram that outlines the problem space of a game called “Towers of Hanoi.” Figure 6 shows the starting position, goal position, and all of the various possible permutations that can happen in the game. In this game, there are disks of different sizes that can be moved and stacked, given the rules until the goal position of cascading disks (smallest to largest) in a specific position is achieved (Jones, 2009, p.24). In the decision tree each move is analyzed for its resulting position and potential next moves available to the game (Jones, 2009, p.13).

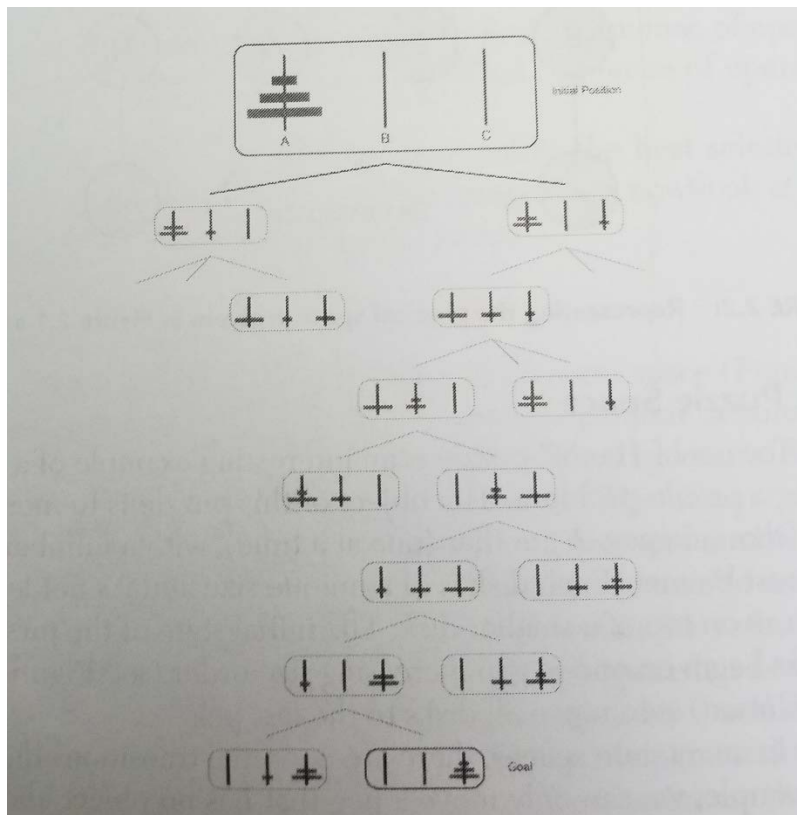


Figure 6. “Tower of Hanoi” Decision Space. Source: Jones (2009).

Uninformed search has two primary means for evaluating the various actions inside of a given construct. These are “Generate and Test” and “Random Search” (Jones, 2009, p. 31). Generate and test is an approach that that generates a potential solution and tests it against the stated goal (Jones, 2009, p. 31). If the solution is found then the search ends, but if the solution is not found an iteration occurs with a new proposed solution (Jones, 2009, p. 31). This is essentially, the algorithm defining end states to see if they are the goal solution, then using the rules to figure out how it got there from the preceding step. It works backward step by step until it reaches the starting position.

Random search works the other direction. This blind search uses the rules to develop a random next possible state (Jones, 2009, p. 31). Each of these states has the rules applied to it to see if further actions can be taken, and the goal state to see if an acceptable course of action is already found (Jones, 2009, p. 31). The algorithm works onward creating the arms of the tree diagram. There are more specific and technical approaches to uninformed search, but at its core it can work forward or backwards through a problem space, generating a logical sequence of steps that leads from a starting position to a goal position.

b. Informed Search

Informed search works in a very similar way to uninformed search, either starting at the goal state or working ahead from the starting conditions (Jones, 2009, p. 49). The difference is that informed search uses a heuristic (Jones, 2009, p. 49). A heuristic is a rule of thumb that uses simple criteria or methods to discriminate correctly between good and bad choices (Pearl, 1985, p. 3). A search algorithm can use a heuristic to increase efficiency while searching, determining the quality of actions in a decision space (Jones, 2009, p. 48). In a tree diagram, the heuristic produces a strategy of which branches to investigate first, as they are most likely to hold the series of actions that lead to the goal.

The “Eight Puzzle” offers a great example of where heuristics can be used to make search more efficient (Jones, 2009, p. 60). An eight puzzle game consists of a grid of nine squares holding eight numbers, 1–8. The numbers can be in any grid square to start, but the goal position is a specific placement of each number inside of the grid in

chronological order; starting with 1 in the upper left grid square and 8 in the middle on the bottom row. The 9th grid space should be empty. The game rule is that you may only move numbers adjacent to the empty space, into the empty space (Jones, 2009, p. 31). This is actually a common children's game, often using a picture and little sliding tiles. Figure 7 shows how the game can progress given a starting position and working through the first two moves. Similar to uninformed search, a decision tree can show the possible moves and the resulting possible actions if each of the decisions is made. The decision tree is actually much longer than in Figure 7. Still, we can understand that a heuristic reduces the possible decision space by applying a decision criterion for the computer.

Two common heuristics used in the Eight Puzzle are the **number of tiles out of place at each decision** and the **distance of each tile to its ending position** (Jones, 2009, p. 60–61). Normally, heuristics are a human capacity, but using these two can quantify the decision space making it relevant for computer use. Using these heuristics, the computer can reduce and eliminate the decision tree arms worth searching down. Each branch of tree that doesn't require investigating based on the heuristic can be skipped (Jones, 2009, p. 60–61). Perhaps the branch increases the number of tiles out of place and is skipped, or it increases the average distance of tiles from their goal position and is skipped. Both of these approaches make informed search superior to uninformed search. Heuristics, however, require a simplified model of the problem at hand and do not necessarily output a perfect understanding of the problem (Pearl, 1985, p. 115). The key is that heuristics take problem knowledge into account to help guide the search within the problem space (Jones, 2009, p. 49).

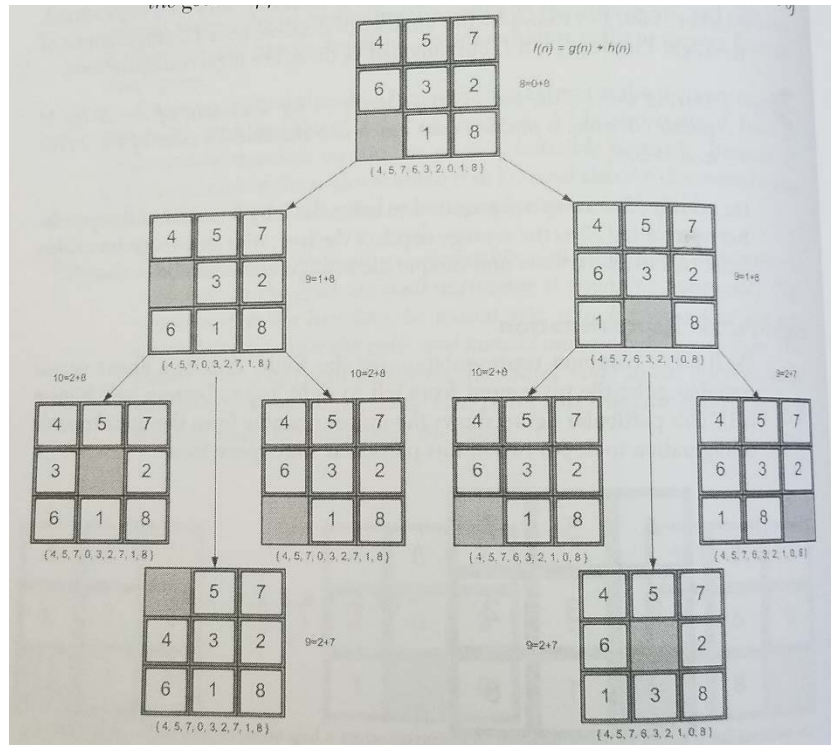


Figure 7. First Two Moves for the Eight Game. Source: Jones (2009).

c. Game-Based AI

Just as search functions have common algorithms used in them, games have a series of algorithm techniques that are commonly used. These techniques include types for zero sum games like checkers and chess and increase in complexity to fit modern video games like first-person shooters and strategy games (Jones, 2009, p. 89). Many of these algorithms build on the tenets of analyzing potential moves found in search algorithms, but go further and apply pruning techniques to better achieve decisions (Jones, 2009, p. 139). Additionally, they must consider a decision space that changes between when one decision is made and the next (Jones, 2009, p. 90).

The *minimax algorithm* is a classic game technique that expands on the informed search techniques. These are particularly effective in alternating move games where success is mutually exclusive; examples are checkers, tic-tac-toe, and chess (Jones, 2009, p. 92). At each potential move the algorithm establishes a goodness of move value for the player, depending on how much it furthers the problem space toward the various goal

positions (Jones, 2009, p. 92). Each decision's value is called "ply" and the total value for a decision tree is the maximum ply found in that decision tree (Jones, 2009, p. 92). The ideal move would advance toward a decision tree's maximum ply value. Once the alternating move is made, the problem space changes and a new series of plies and maximums is calculated (Jones, 2009, p. 92). In this incremental fashion the search function can work its way to a maximization, hopefully victory, given the possible moves and alternative moves made.

Many of these types of algorithms leverage pruning. Pruning is eliminating decision arms that would lead to disaster on the next move (Jones, 2009, p. 101). Imagine a move in chess where you position the king into harm's way. This is an invalid move, and the entire decision tree resulting from that move can be eliminated (Jones, 2009, p. 101). As humans we find a task like that simple, but it requires an algorithm to be built that understands the problem space and can understand the outcomes of various moves.

Modern video games have proven very useful for algorithm development. They have a rich environment already programmed and can translate visually the impacts of certain algorithms (Jones, 2009, p. 122). These applications can bring to life Japan's favorite plumber or allow practicing dangerous real world talents in a safe space (Jones, 2009, p. 122). A key element that makes this possible are the *pathfinding algorithms*. These algorithms analyze potential paths from point A to point B in the game environment, framing different choices inside the potential problem space (Jones, 2009, p. 123). Each choice brings with it possible advantages, disadvantages, constraints, and resource costs (Jones, 2009, p. 123). Analyzing this type of action using classic search algorithms already discussed is too time consuming and costly on computer resources to meet the real time demands for a realistic feeling environment. To handle this type of problem, static state machine algorithms and other recursive logic paths are used (Jones, 2009, p. 130). Figure 8 shows a static state machine logic path for a common enemy in a shooting game. The enemy patrols a path in the environment, confined to two points and a few basic actions. If one of the triggers is met, the enemy can leave the predetermined path and engage the player at a new location. In this way, the algorithm can disregard all

possible ways to move around the environment and efficiently represent the capabilities of an enemy on a digital platform.

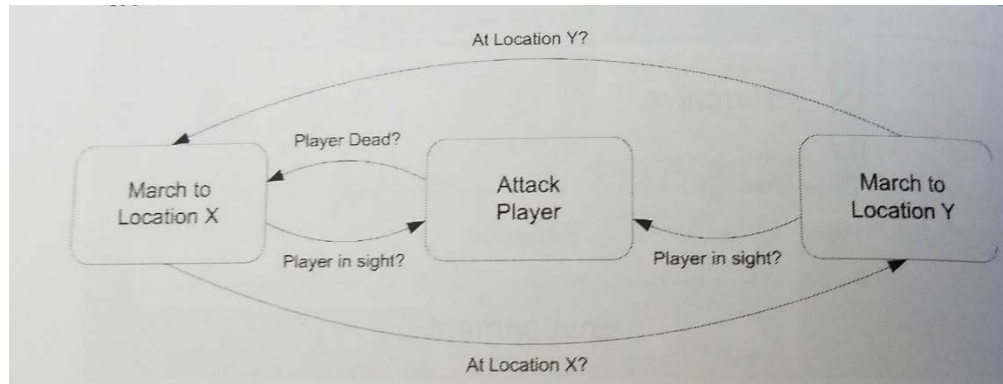


Figure 8. Static State Machine for a Simple AI Enemy

This same game based algorithm approach, when scaled larger can represent entire armies or squadron level air-to-air tactics (Jones, 2009, p. 133). Remember this is a very simplified example, in reality a game has many logical connections, leveraging high level scripts, script interpreters, and a game engine providing a virtual physics (Jones, 2009, p. 132).

d. Knowledge Representation

Knowledge representation is the theory and practice of storing different types of knowledge in a computer system (Jones, 2009, p. 143). In AI, knowledge is used to enable intelligent entities to make intelligent decisions about their environment (Jones, 2009, p. 144). It creates the knowledge base that builds up a portion of the AI algorithm.. Jones (2009) explains that knowledge can be knowing “that fire is hot (and should avoided), or that water in certain cases can be used to douse fire to make it passable” (p.144). Storing knowledge like this using semantic networks and framing, allows the AI algorithm to make decisions about how to effectively navigate its environment (Jones, 2009, p. 144).

There are four types of knowledge and many techniques that can be used to store this knowledge in AI algorithms. The key forms of knowledge to know are declarative,

procedural, analogous, and meta-knowledge (Jones, 2009, p. 144). Declarative forms are statements of fact that can be accessed by an algorithm (Jones, 2009, p. 144). This is usually displayed as logic and is flexible because it has the potential to be used in ways beyond the original intent (Jones, 2009, p. 144). Less flexible is procedural knowledge, which takes the form of a series of steps that help lead to achieving a goal (Jones, 2009, p. 144). An example would be production steps. Analogous knowledge, the ability to know through association, and meta-knowledge, knowledge about knowledge, supplement the first two forms to create a robust knowledge base (Jones, 2009, p. 144).

The tools to translate these forms of knowledge are wide, but two very common approaches are semantic networks and first order logic. Semantic networks are a digital relationship diagrams between any number of objects (Jones, 2009, p. 145). Objects can be traced to each other using arcs, which are functions like “is_a” and “part_of” that define the relationship between the two objects (Jones, 2009, p. 145). In this way knowledge of the cities and states that compose a country can be represented as knowledge, complete with relationships of what cities are in what states. Figure 9 shows an example of a semantic network.

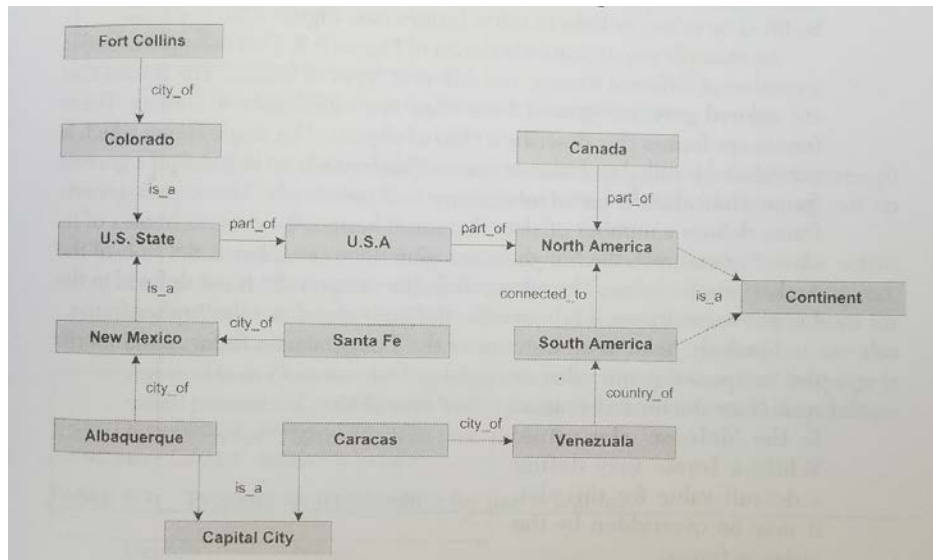


Figure 9. Example of a Semantic Network. Source: Jones (2009).

e. Semantic Network

First order logic, leverages atomic sentences and quantifiers to bring simple facts into knowledge inside of a computer (Jones, 2009, p. 153). An atomic sentence references the known objects in an environment, in a fashion that explains relationships or properties (Jones, 2009, p. 153). Say the computer knows about a person named Marc exists, bicycles are a form of transportation, and the word “rides” is the action of a person sitting on something. These are declarative forms of knowledge. An atomic sentence in a program would read as “Rides (marc, bicycle)” to the computer, but represent the knowledge of “Marc rides a bicycle” (Jones, 2009, p. 153). In this way the computer understands the relationships and actions of the people in the environment. By working these forms of knowledge into a knowledge base, a computer can reference more complex information while executing AI algorithms.

f. Machine Learning

Machine learning is a discipline revolving around AI algorithms that allow a computer to learn (Jones, 2009, p. 153). Jones (2009) writes, “given a set of data, machine learning [algorithms] can learn about the data and their relationships” (p. 171). The three main aspects of machine learning are supervised learning, unsupervised learning, and probabilistic learning (p. 171).

Supervised learning algorithms utilize a teacher to tell the computer whether it has made a correct decision, using predictors and target output (Jones, 2009, p. 171). Using perceptron, backpropagation, and decision tree tools the computer uses the predictor and target value to generate an output (Jones, 2009, p. 171–172). Using the feedback from the supervisor the algorithm adjusts its core function to better achieve the target (Jones, 2009, p. 172). Consider a person watching a video game where the autonomic player has the predictors of weapon (knife or gun), health (full or low), ammo (full or low) and output behaviors of fight or evade (Jones, 2009, p. 173). The various predictors are tracked and output decisions that lead to survival are tracked. The results are a series of predictor conditions that lead to successful or failed fights that should have been evasion tactics instead. When the supervisor provides the feedback that evade is correct in certain

predictor conditions, the computer can learn the predictor conditions that lead to successful predictor conditions to evade (Jones, 2009, p. 173). The computer can learn to behave more appropriately to fit the environment and increase the quality of decisions.

Unsupervised learning is different in that it doesn't use a target and grade, but rather the feedback is in the form of the output from the original function (Jones, 2009, p. 176). This feedback usually takes the form of a pattern derived from a set of data (Jones, 2009, p. 176). The best example for this is a smart home device. The device actively tracks the conditions that the house is set to, recognizes the settings are consistent at certain times [outputs], and adjusts at those times to the conditions it assumes the homeowner prefers (Jones, 2009, p. 177). In this case the output is a set of times and associated house conditions that the person prefers and can be incorporated into the basic settings. This is a form of Markov model, which has many applications in unsupervised learning.

These Markov models become probabilistic learning tools when they mix in probabilistic models with the observed feedback (Jones, 2009, p. 177). This serves to give the computer a sense of likelihood, when considering how to implement the feedback. Jones (2009) explains this approach can be used to, "generate reasonably syntactically correct words using known words for training" (p. 177). This means that the number of times a person typing follows the letters "th" with the letter "e" can be tracked and combined with known probabilities about the number of times people in general use the word "the" versus the word "theoretical" (Jones, 2009, p. 177). After using this as feedback, the computer can have a quantifiable expectation that after the letters "th" is typed, it will result in the word "the", rather than "theoretical" (Jones, 2009, p. 177). This training, analysis, and feedback of probability of next letter, allows machines to accurately perform word generation and offer a typist the word they are more likely to choose (Jones, 2009, p. 179). Word generation skills are relatively simple, but the effect can be used across a spectrum of skills including speech recognition, speech understanding, and music composition (Jones, 2009, p. 184).

g. Evolutionary Computation

Evolutionary computation is a technique that applies simulated evolution to algorithms, by evolving the population of potential solutions using natural selection (Jones, 2009, p. 195). A potential solution that is determined weaker is removed from the population and incrementally stronger solutions are entered into the decision space (Jones, 2009, p. 195). Inside of evolutionary computation are very natural algorithms, including genetic representations and biologically inspired algorithms that allow for swarm behavior.

Genetic algorithms work to eliminate weak solutions to problems and grow stronger solutions by treating the solutions themselves as chromosomes (Jones, 2009, p. 198). Possible solutions (chromosomes) are encoded as a sequence of potential actions (genes) and stored in the general population, then are tested as a solution to the problem. The portions of the code (genes) are graded for their strength and weaker portions are subjected to a recoding (Jones, 2009, p. 197). Figure 10 displays how the solutions are selected, mutated, and re-entered to the problem space. This process is repeated millions of times until the solution space consists of stronger and stronger solutions (Jones, 2009, p. 196). This process is slower than searching and analyzing solutions to the problems, but when considering complex problems with few defined correct solutions, it excels (Jones, 2009, p. 195). Modern computing power has further increased its use in multivariate optimization problems, or simply problems that are open ended and seeking to be optimized (Jones, 2009, p. 195).

Figure 10. Process for Mutation in Genetic Algorithms.
Source: Jones (2009).

A specific subset of evolutionary computation is biologically inspired algorithms that bring life to particle swarm optimization. Particle swarm optimization is a technique that uses a few simple rules, which give the individual particles freedom to define their relationships to each other to make a swarm function (Jones, 2009, p. 236). The solution space is all of the ways that particles can relate to the other particles in the swarm (Jones, 2009, p. 236). The first step is that a population of random vectors and velocities are generated and applied to each particle (Jones, 2009, p. 237). Of course, mayhem ensues. The fitness of each individual behavior is analyzed and good behaviors are stored as a best behavior for the particles that own them (Jones, 2009, p. 237). The poor behaviors are evolved using a mutating program and reinserted randomly to supply new velocities and vectors (Jones, 2009, p. 237). Incrementally, poor positioning and speed are removed from the swarm and particles store quality behaviors, causing an optimization of the swarm itself (Jones, 2009, p. 238).

h. Neural Networks

A *neural network* is a system of computers, arranged like a human brain, where each computer acts as a neuron in charge of conducting some portion of the computation (Jones, 2009, p. 249). Similar to a human nerve structure, the computers are networked through a series of axons and dendrites that route inputs and outputs around the network to the destinations they are needed at (Jones, 2009, p. 251). These types of networks are excellent at forming patterns and generalizations about groups of data that normal computers are unable to decipher (Jones, 2009, p. 249), as in the cases of credit risk assessment, adaptive signal processing, and defining functions inside of seemingly arbitrary data sets (Jones, 2009, p. 251)

Figure 11 displays how a simple neural network functions. Inputs come to the neuron (computer) through an axon, a function or algorithm performs a computation, and the output is routed using a transfer function through the network to other areas of the computers for processing (Jones, 2009, p. 252). If the input values are significant to the computer, the functions of that computer activate and perform their computation (Jones, 2009, p. 253). If the input is insignificant, the function lies dormant and routes the values to a computer it thinks can use the information. The brilliance is found by weighting the outputs. By weighting the outputs, values that leave one computer are transformed from significant to insignificant, or vice versa, based on the computer they are headed to (Jones, 2009, p. 253). Once the algorithm begins to sense a pattern in the information, it can adjust the weights to route significant data to the correct processing areas (Jones, 2009, p. 253). Not only is a pattern of significant inputs uncovered, but optimization and minimizations can be made from the pattern (Jones, 2009, p. 262).

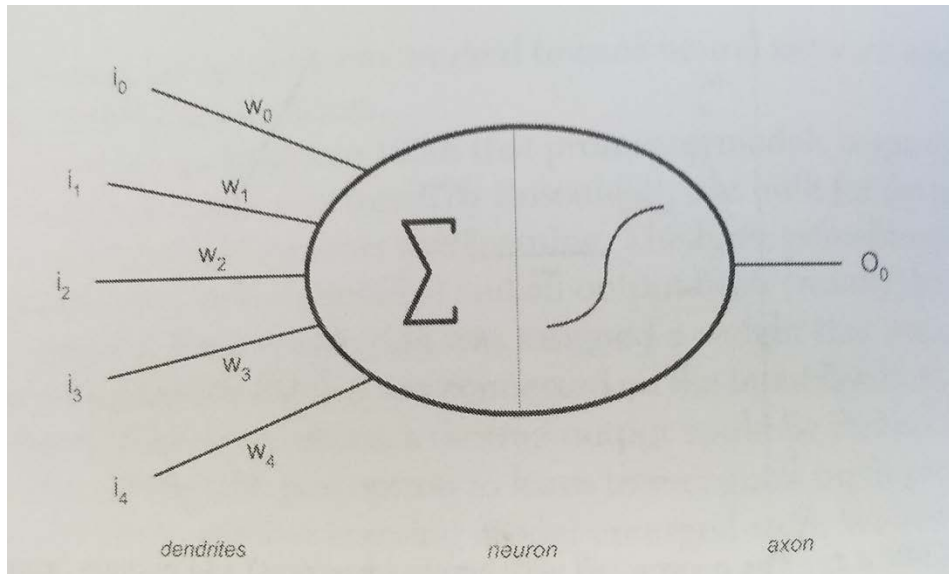


Figure 11. Singular Neuron in a Neural Network. Source: Jones (2009).

A common algorithm for processing inside of neural networks is called backpropagation. Jones (2009) defines backpropagation as,

for a test set, propagate one test through the MLP [a type of neural network] in order to calculate the output (or outputs). Compute the error which will be the difference of the expected outcome and the actual value. Finally, backpropagate this error through the network by adjusting all of the weights; starting from the weights closest to the output layer and ending at the weights to the input layer. (Jones, 2009, p. 266)

This means to decrease the error of outputs, compared to known and expected quantities, the algorithm runs the computation again using the same inputs, but adjusting the weights using the output error as a tuning fork. Similar to evolutionary computation, high performing weights and routing mechanisms can be retained, and the overall function of the network improves. This can be done as supervised or unsupervised learning, growing skill in the network until it can be used to seek out natural patterns in novel uses (Jones, 2009, p. 266).

The most useful skill enabled in neural networks is *generalization* (Jones, 2009, p. 268). This means that the network can make a general statement about inputs to the system, be they binary, Boolean, or graphical. This skill of generalization is very handy in pattern recognition. Figure 12 shows a set of grid based numbers that can be used to

train a neural network to seek out numbers in almost any type of image (Jones, 2009, p. 268). A specific number is input and the many connections of the network route the information from computer to computer, analyzing the grid squares that define the number. The weights illuminate certain computers to perform their calculations, and an output is made (Jones, 2009, p. 268). This could be correct or incorrect. The mean squared error is applied to the weights and through backpropagation and iteration the computers collectively get to the correct output, a matching number (Jones, 2009, p. 269). The next step would be to introduce a picture of the same number whose shape is not a perfect match to one defined with gridded squares. The computer would take the image, break it into pieces and apply the learned weights, across the spectrum of possible numbers, until it can minimize the errors (Jones, 2009, p. 272). It should result in a correct recognition and computerization of a numerical image. Noise is used in training to stave off overfitting, which would be the case of a computer being too picky about the exact shape of the numbers it is looking for (Jones, 2009, p. 269).

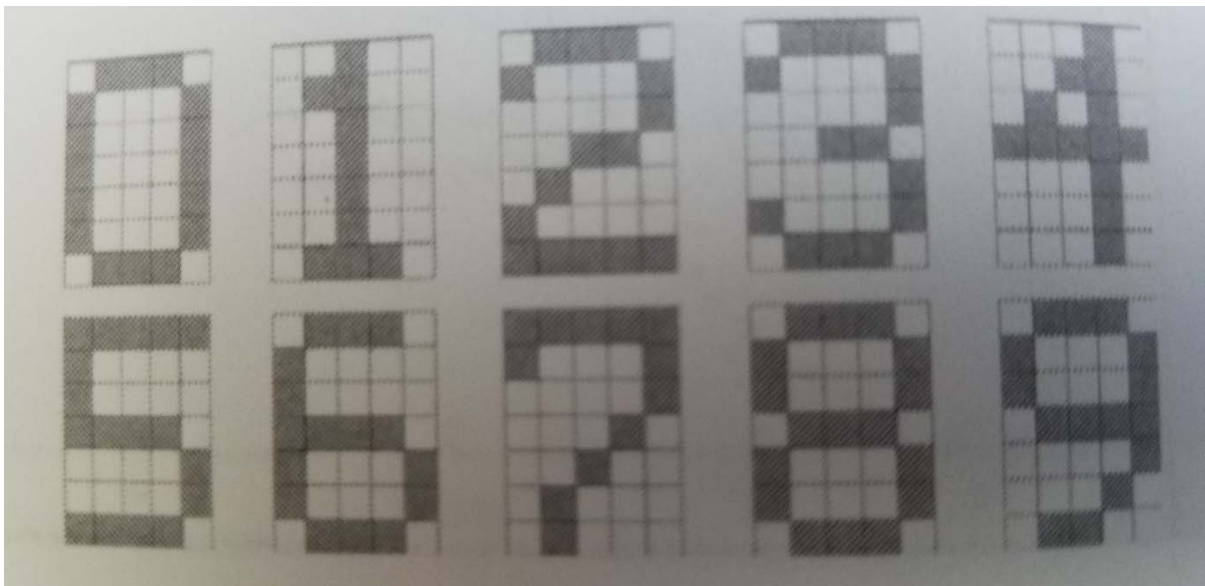


Figure 12. Possible Training Aid for Neural Networks. Source: Jones (2009).

These AI techniques are what bring to life the appearance of cognition in machine form. There are many specific nuances and subcategories that exist in these realms and

other publications may group and categorize differently. Often, the specific techniques could be grouped into several techniques domains. The key is to understand is that these machines are not magically smart based on a new type of microchip. They are an incremental build on hundreds of years of research in mathematics, computer processing speed, more efficient programming languages, linguistics and other disciplines that have led to an ever increasing level of usefulness. Additionally, the combination of these techniques with robust knowledge bases makes AI algorithms very different from a basic algorithm. Even after the research is done, the system must be trained to behave in the manor that is expected of it and to provide sufficiently accurate outputs for use. These AI algorithms are ready to be incorporated in AI systems, and fortunately there is growing commercial community ready to assist.

C. STATE OF THE AI COMMUNITY

The buzz around AI has been off and on since the 1950s, even hitting a period called AI's Winter in the 1970s (Jones, 2009, p. 8). Today the buzz is still low, but quickly gaining in research intensity and investment from the world's largest companies. Facebook, Google, and Amazon are all investing in AI technology to try and help them increase profitability and at the same time are endeavoring to further the research field along by openly publishing the results of their dedicated research arms. Visionaries are springing forward with opinions and predictions of all types. Companies dedicated to supplying the DOD with AI services and systems are emerging across the United States, positioning for the DOD's transition to those types systems. The state of this community is important for acquisition personnel to understand. It doesn't behave like the traditional markets that DOD has operated in and is full of more colorful characters than we have seen in a long time.

1. The Domain of AI Research

Apple Inc. spent \$10B on research and development (R&D) in 2016 (Apple Inc., 2017, p. 6). In an article at *Business Insider* online magazine, writer Kif Leswing analyzes the trends in Apple's recent R&D. He is capable of analyzing the amount spent and what products have come to market, but is only capable of speculating what apple

has been working on in its new or planned R&D centers in Japan, China, Indonesia, France, Sweden, and the United Kingdom (Leswing, 2017). Leswing's speculation is natural, given that protecting trade secrets to maintain a competitive advantage is relatively common across many industries. When asked about the research plans for Apple, CEO Tim Cook, skirts the question referencing growth in spending and insinuating that it is new ground for Apple (Leswing, 2017). Again it is natural for Cook to sidestep the question, but it is in stark contrast to the growing trend across the AI field of research. The primary difference is the amount of investment toward openly published AI knowledge from otherwise profit seeking companies. Four companies in particular Amazon, Google, Facebook, and IBM are setting the example for AI research in an open and results sharing world.

a. Amazon

Now the world's most valuable retailer, besting Wal-Mart in market value in July of 2015 and never looking back, Amazon is the king of selling products (Li, 2015). Part of what has made their climb so monumental is their approach of customer obsession, the notion where they will do anything to get a seller's product to the hands of customers as fast as possible (Amazon, 2017, p. 3). But they rest on a few other principles as well, one of which is passion for invention and competition (Amazon, 2017, p. 3).

In the world of online retailing they have plotted a course to compete with everyone from making television shows to online banking (Amazon, 2017, p. 4). Staying competitive in those realms, and the space in between, requires significant investment into new and existing products. Amazon has spent \$9.275B, \$12.540B, and \$16.085B on Technology and Content [R&D] in 2014, 2015, and 2016 respectively (Amazon, 2017, p. 37). A significant portion is of that R&D is used on, "employees involved in the application, maintenance, operation, and development of new and existing products" (Amazon, 2017, p. 28).

Amazon has already pioneered an in home AI device. In 2015 Amazon released a product called the Amazon Echo (Nuñez, 2015). This product is an immobot that offers a voice control of the machine and any Wi-Fi connected device, leveraging an open-source

developer's toolkit for third party appliance connection (Nuñez, 2015). It can be used to order products through Amazon's suite of services, or can be trained using the Alexa application to interact and control other smart products in your house (Nuñez, 2015). The Amazon Echo technology has been adapted since then to operate on screened devices, remote controls, and watches and is adding interfaces for streaming music connections and robotic vacuums (Amazon, n.d.). Its ability to standardize the code and facilitate interoperability across software platforms and architectures is one of a kind.

A key effort to their product development is in the effort of corporate support to OpenAI. OpenAI is a non-profit AI research company, discovering and enacting the path to safe artificial general intelligence" (OpenAI, n.d.). OpenAI has both private sponsors like Sam Altman, Elon Musk, and Peter Thiel, and corporate sponsors like Microsoft, Infosys, and Amazon Web Services; collectively they have pledged \$1B. From these investment OpenAI works to publish advanced AI research journals, focusing on topics like smart AI regulations and open-source software (OpenAI, n.d.). Two such open-source software kits are Gym and Universe, which help to develop learning algorithms and grade AI algorithm intelligence respectively (OpenAI, n.d.).

The DOD is not usually looking for a way to control robotic vacuums, but strictly from a facilities management standpoint it offers a real mechanism to exercise control over energy use and working conditions. Its ability to learn, predict, and implement generalized observations offer an advantage in terms of infrastructure optimization and a risk in the form of counterintelligence. Further, immobots like this have the capability of accepting inputs from a logistics system and reordering spare parts instantaneously, fully leveraging transparency across the logistics pipeline. Amazon is at the forefront of AI technology and research and will be a major driver in the direction of AI capabilities.

b. Google

Google is quickly becoming a pillar in the AI research and development field. Parent company, Alphabet, clearly articulates in its annual report that machine learning is a central focus of the company (Google, 2017, p. 3). They write, "it's what allows you to use your voice to search for information, to translate the web from one language to

another, to see better YouTube recommendations, and search for people and events that matter to you in Google Photos” (Google, 2017, p. 3). Further, they outline future products like learning thermostats and complete companies that they intend to establish in an effort to leverage AI technologies to make the world a better place (Google, 2017, p. 5). Alphabet spent \$9.832B, \$12.282B, and \$13.948B in R&D in 2014, 2015, and 2016 respectively (Google, 2017, p. 5). This is clearly a company that is dedicated to fleshing out future technologies like AI to help leverage their competitive advantage in the Internet industry.

Google is not only interested in AI products, but also swelling the scientific field of AI. Inserting a new normal for openly published research and development, Google has established a research arm, Research at Google. Its goal is to “tackle the most challenging problems in computer science and related fields” (Research at Google, n.d.). It has hundreds of research articles published in 21 research areas including, algorithms and theory, machine intelligence, machine perception, natural language processing, quantum AI, and a collection of articles focused on a Google technology called “Google Brain Residency” (Research at Google, n.d.). Google is absolutely interested in harvesting the results of this research, stating, “language, speech translation, visual processing, ranking and prediction relies on machine intelligence” (Research at Google, n.d.). Still, these articles are not illegitimate propaganda about Google products. These are scientific publications, peer reviewed, and published in accordance with the highest scientific standard. Many are co-authored and include researchers not permanently employed by the Google research team (Research at Google, n.d.). Further, they are openly published for review at Google and Cornell University Library (Research at Google, n.d.). The head of Research at Google is Peter Norvig, who helped this research define AI (Research at Google, n.d.)

A recent publication from Research at Google in the field of machine learning, *Adversarial Attacks on Neural Network Policies*, grapples with the information security aspect of training an AI machine to exhibit human behavior. It explains that similar to the overfitting in neural networks, machines are products of their training (Huang, Papernot, Goodfellow, Duan, & Abbeel, 2017, p. 1). The team explains miniature perturbations,

unobservable to human trainers, can cause the quality of learning to degrade in a network (Huang et al., 2017, p. 3). The journal concludes, machines are susceptible to picking up bad habits from adversarial attempts and, “it is indeed possible to significantly decrease the policy’s performance through introducing relatively small perturbations in the inputs” (Huang et al., 2017, p. 6). This research is directly applicable to the fielding of AI systems in the DOD. Any AI system entering Technology Maturation, will require training of the AI algorithm, to include protecting it from adversarial degradation.

c. Facebook

Another surprising leader in the dedicated AI research realm is Facebook Research. Facebook Research is divided into eleven separate programs: applied machine learning, connectivity, economics and computation, human computer interaction & UX, security & privacy, virtual reality, computer vision, data science, Facebook AI research, natural language processing & speech, and systems & networking (Facebook Research, n.d.). Each program has a dedicated lead, but products that are generated from each do not fit the usual mold of corporate production (Facebook Research, n.d.). Publications are group efforts generated by team members from various organization and while they do further Facebook’s corporate agenda, also further the generation and curation of AI (Facebook Research, n.d.). All publications are met with the rigor expected of scholarly work, then published as articles and posted in a blog style (Facebook Research, n.d.). This again redefines the normal roles associated with how technological research is conducted, leading to a corporate sponsored open source information.

Facebook AI Research is lead, since 2013, by Director of AI Research Yann LeChun. He states Facebook Research’s mission is to, “deeply engage with academia to drive progress” (Facebook Research, n.d.). This matches Facebook’s financial statements where they state they are funding, “long-term technology initiatives that we believe will further our mission to connect the world, such as virtual reality and artificial intelligence” (Facebook, 2017, p. 5). Facebook’s commitment is supported by research and development expenses of \$5.92 billion, \$4.82 billion, and \$2.67 billion in 2016, 2015, and 2014, respectively (Facebook, 2017, p. 6). Logically, these investments should

directly relate to Facebook's corporate goals, but the opposite is true. Many investments are concerned with general advancement of the field of AI.

A paper published under the Facebook Artificial Intelligent Research Group (FAIR) concerning the existence of observable footprints that reveal the "causal disposition" of the object categories in the collection of images, is a great example of an applicable AI effort (Facebook Research, n.d.). Causal direction stems from the idea that, "Correlation does not define causation, and very few instances allow for causal determination. Still aiming for correlation is possible" (Neale & Cardon, 2011). The process for finding a solution in a set of data is to, when correlations can be determined, measured, and "placed in order," creating a longitudinal analysis and space where a cause should reside (Neale & Cardon, 2011). The FAIR paper takes this idea and applies it to image interpretations. AI algorithms have the ability process image based data sets and be trained to understand the direction of causations imbedded in the picture (Lopez-Paz, Nishihara, Chintala, Scholkopf, & Bottou, 2017, p. 9).

They use the idea of a picture of a car to explain their research. Imagine a picture of a car on a bridge. A computer is capable of understanding the context of the image and identify the objects that exist in the picture (Lopez-Paz et al., 2017, p. 1). By using machine learning influences and training with scatter plots it is possible to understand several contextual relationships in the picture and the direction of causation for those relationships (Lopez-Paz et al., 2017, p. 2). If the car were removed nothing would change as it relates to the overall scene, but if the bridge were removed then certainly the other aspects of the image would change (Lopez-Paz et al., 2017, p. 1). Most likely the car would not exist in the location [picture and data] where it currently is. The algorithm can then conclude the direction of causation, the bridge is causing the car more than the car is causing the bridge to exists in its current location (Lopez-Paz et al., 2017, p. 2). While the bridge is not the only cause for a car to be on a bridge, the primary causes for the cars location can be found through the bridge to the other potential contributors for the cars location and behavior. If similar pictures of the car existed, the causal disposition could be extrapolated further to define other factors causing the car to behave in the way it is.

This may appear like a technology best used for showing a college student the best way to a local party, and in turn advertising dollars for Facebook, but at its core this is a publication about gleaning information from pictures. With increased use of advanced imaging technology in the intelligence community, carried out through UAVs, space systems, and in-computer cyber detection tools, there is a clear path for AI implementations like this in military systems.

d. IBM

International Business Machines Corporation (IBM) was once thought of as a hardware company, supplying computers and printers for other corporations that needed computing power. Today they are working to rebrand the company as a, “cognitive solutions and cloud platform company” (International Business Machines Corporation [IBM], 2017, p. 1). IBM defines cognitive solutions as, “the highest level of intelligence that exists in technology systems...ranging from answering client inquiries to helping physicians fight cancer” (IBM, 2017, p. 2). Averaging roughly \$5.5B in R&D spending over the last three years, aimed at bring their AI services approach to the global industry (IBM, 2017, p. 132).

Watson is a particularly interesting AI system that is working to generalize massive amounts of industry data, in order to supplement IBM customer support abilities (IBM, 2017, p. 2). Watson is a software agent that lives in the cloud owned and operated by IBM (IBM Think Academy, 2014). They claim that Watson escapes rigid decision tree processing by leveraging language processing skills that give it context into the literature of a specific field of study (IBM Think Academy, 2014).

First it defines its knowledge base by searching through as much literature as can be brought to it (IBM Think Academy, 2014). IBM uses industry experts to teach Watson what constitutes quality sources and what industry specific contexts are important (IBM Think Academy, 2014). From this, it can piece together patterns and generalizations from data and metadata that are graded for accuracy by the industry experts (IBM Think Academy, 2014). This could be medical information or data concerning professional athlete contracts. In either case the patterns are called hypothesis, and these hypotheses

are graded for accuracy (IBM Think Academy, 2014). If there are gaps in the knowledge base Watson can create questions for the experts to answer (IBM Think Academy, 2014). Watson even analyzes the inputs of various professionals and weights their inputs relative to the quality they provide (IBM Think Academy, 2014). This fashion of defining an industry, supervised machine learning, and grading of Watson's outputs, allows IBM to use the machine to define solutions to the hard to answer questions in an industry (IBM Think Academy, 2014). For example, the correct plan for treatment of a specific type of cancer or who is the correct player to draft for a sports franchise.

Watson is prime for analysis in the defense industry from manufacturing techniques and tolerance allowances to analysis of material choice and areas to reduce weight in systems. In the DOD, analysis on types of contract to use or minimum requirement of DAS acquisition thresholds to apply are in the realm of solvable. Anything that benefits from both quantitative and qualitative analysis, Watson can lend a hand given a dedicated team of capable trainers.

This is certainly not an exhaustive list of major companies interested in AI research. Many other American companies and foreign companies, like China's Baidu, are actively engaged in scientific publication and the advancement of the field of AI. Also, there are many traditional agencies, like DARPA and NASA, focusing their attention on AI development. Many of the research products sponsored and published under these companies, are assisted or coming directly from research centers at universities. Non-profits are being organized and federal funding is bolstering a large portion of AI research. The companies in this research are highlighted, however, because they represent a new playing field for defense procurement. paradigm is what is different about AI compared to traditional engineering disciplines. The companies that we have relied on for the last 100 years are not necessarily the companies with the focus and expertise in AI research or AI systems. Major technology companies, profiting in a digital world, are influencing the state of the art in AI, while many major defense contractors are operating under the decreased impact of DOD and federal spending. To fully understand and successfully procure AI will mean understanding industry effects caused by these and similar companies. In many ways, software intensive systems have

been inching this direction, but with AI there is a new research battle space that we must prepare for.

2. AI Industry Leaders

Throughout history many technological breakthroughs have occurred due to the guidance and curation of dreamers. The dreamers envision the future world and the various uses for each technology. AI is no different. It has a laundry list of colorful characters, at universities and private companies, who each expect to craft AI in a different way. Understanding these visionaries goes a long way toward understanding how AI will mature and can be applied to defense systems.

a. Jeff Bezos

In 2017 annual letter to shareholders, Jeff Bezos founder and CEO of Amazon, explained what AI means to Amazon in today's world. Bezos penned (Bezos):

much of what we do with machine learning happens beneath the surface. Machine learning drives our algorithms for demand forecasting, product search ranking, product and deals recommendations, merchandising placements, fraud detection, translations, and much more. Though less visible, much of the impact of machine learning will be of this type — quietly but meaningfully improving core operations (2017)

In writing this, Bezos outlines his vision for AI today, working in the background to create efficient and supportive systems. In the letter he also outlines the practical products and applications amazon is engaging in, from the already discussed Alexa software to the mind boggling autonomous Prime Air Delivery Drones (Bezos, 2017). From this we can learn his vision of a connected world, where he believes machines take on the actions once performed by humans and bring convenience to a mass market (Bezos, 2017).

Later in the letter he delves into the future. AI will be in widespread use, “from early disease detection to increasing crop yields” (Bezos, 2017). He explains the products like Amazon Polly and Amazon Rekognition, not yet in existence, that remove the heavy lifting from design and implementation of machine learning; no machine learning expertise required (Bezos, 2017). Bezos is less of a futurist though, than he is a

businessman. He is interested in shaping the AI in the near future and bringing it to world as products. Speaking at the 2007 Internet Associations Annual Gala in Washington Bezos spoke about AI, calling this, “a golden age” (Galeon & Gohd, 2017). His opinions are shaping the way AI will immediately shape the world around us.

b. Elon Musk

Elon Musk has a different focus. Born in South Africa in 1971, Musk is the CEO and founder of SpaceX, Tesla, Solar City, and has founded many other companies from the prominent PayPal to the lesser known Zip2 and X.com (Biography.com, 2017). As a child he was enamored with science fiction novels and computers, learning to program before the age of twelve (Biography.com, 2017). After a double bachelor’s degree in economics and physics at the University of Pennsylvania, Musk went to work (Biography.com, 2017). After selling his first company Zip2 to Compaq in 1999 for \$307 million in cash, Musk has been a power player in the technology industry and hell-bent on technology development (Biography.com, 2017). He started SpaceX in 2002 with the mission of commercial space travel and an ultimate goal of enabling human life on Mars (SpaceX, n.d.).

To go along with his technology savvy mind and entrepreneurial skill, Musk has crafted an image of a brash talker and futuristic thinker. At the National Governors Association Meeting 2017, Musk presented his opinion on AI. He spoke, “AI is a fundamental existential risk for human civilization, and I don’t think people fully appreciate that” (Domonoske, 2017). Musk has found celebrity in stark comments like that, and others made in a letter he penned in 2015 warning of an AI arms race (Domonoske, 2017). While these punchy lines play well in the news, he brings with them often unnoticed and salient arguments. At the Governors meeting he called for proactive government regulation to establish industry and ethicality standards surrounding AI (Domonoske, 2017). He followed with, “I keep sounding the alarm bell,” hammering home that a future with AI developers summoning demons they think they can control poses a risk to human life as we know it (Domonoske, 2017).

But this is not the complete story on Musk. He is one of the founding members of OpenAI (OpenAI, n.d.) and bolstering AI at both SpaceX and Tesla by poaching Andrej Karpathy from the ranks of OpenAI (Knight, 2017). Tesla is actively incorporating machine learning into auto-pilot functions, leveraging the millions of miles already driven by its cars to create detailed maps and behaviors (Knight, 2017). Even going further, Musk has presented the idea of *neural lacing*, a merger of a human brain with AI supported computer functions (Victorino, 2017). While his warnings ring loud, so to, should his efforts to push AI responsibly forward into the future.

c. Qi Lu

Once responsible for Microsoft's AI efforts, Qi Lu has made a move that is very telling about the future of AI (Bacchus, 2017). Qi Lu was raised by his grandparents in rural China (Helft, 2009). Still, he graduated from Fudan University in Shanghai and finished a Master's Degree in Computer Science (Helft, 2009). While working as a professor at Fudan University and making \$10US a day, he attended a talk given by Carnegie Mellon Professor Edmund M. Clarke (Helft, 2009). Being impressed by the question Lu asked, Clarke offered Lu a sponsored invitation to the doctoral program at Carnegie Melon (Helft, 2009). Lu leveraged this opportunity to become an international AI expert. He held a position at Yahoo! excelling at search algorithms until he was poached by Microsoft and became confidant to Steve Ballmer (Helft, 2009). Lu was instrumental in bringing Cortana AI software to life (Helft, 2009) and has taken the COO position at Baidu to usher in a new era of Chinese driven AI (Hempel, 2017).

Lu was a relatively quiet personality in computer science until the move to Baidu, but now has expressed opinions that are a must read for anyone interested in AI. "The company with the most data wins" is a quote he spoke during an interview in Silicon Valley with *WIRED* magazine's Jessi Hempel (2017). Lu believes the 731 million people online in China, nearly twice the population of the United States, offers a significant structural advantage (Hempel, 2017). Now in charge of R&D, sales, and marketing Lu seeks to expand Baidu's role on the AI grand stage (Hempel, 2017). Under Lu's control,

Baidu has created Baidu Brain, a first and one of a kind AI platform that is built of 60 different types of AI services (Hempel, 2017).

Part of what makes Lu's approach to the AI world unique is that he is embracing the arms-race dynamics Musk is concerned with. Lu spoke, "in this race to AI, it's actually about having the right application scenarios and the right ecosystems" adding, "you need an AI-first device to solidify an emerging base of ecosystems" (Hempel, 2017). Lu does explain some of his reservations in AI, citing protection of individual privacy as a chief concern (Hempel, 2017). It is clear from Lu that Baidu and China are as invested as their U.S. counterparts. Beyond establishing Baidu Research, Lu is working to define AI ecosystems for mass consumption and an Apollo self-driving technology (Hempel, 2017). Qi Lu is shaping the foreseeable future for international AI structures (Hempel, 2017) and his tone is generally optimistic stating, "I do feel there's opportunities for China and the United States to collectively drive the world forward" (Hempel, 2017).

3. Commercial AI Products

Research into the subject of AI appears to be surging ahead at a rapid pace, but the DOD is often more concerned with the application of technology. For use by a warfighter, there must be a safe and secure system. Fortunately, sprouting from the open source research and strategic visions are AI centric products aimed for DOD consumption. With varied uses from process management to autonomous transportation systems, these large and small business offerings are the beginning to develop mature commercial off the shelf AI systems.

a. NVIDIA

NVIDIA, was founded in 1993 as a generic computer company and in 1994 released their first product, a custom graphical user interface accelerator for SGS-Thompson (NVIDIA, n.d.-a). CEO and one of the founders, Jensen Huang, has grown the company through video game systems and computer enhancement projects to the point where they are a leader in microchip production (NVIDIA, n.d.-a). Now, the 2007 Forbes Company of the Year, is going all in on AI applications with the launch of their GPU

microchip series (NVIDIA, n.d.-a). These microchips are designed from the bottom up to support deep learning algorithms and be inserted as the brains of AI computers, robots, and autonomous cars (NVIDIA, n.d.-a). The GPU chip set has seen use in the movie industry, AI video game applications, and has been the brains of the ASIMO robot and Tesla automobiles.

Another key NVIDIA product are the JETSON AI supercomputer modules (NVIDIA, n.d.-d). These are like Arduino's consumer aimed processing units that are capable of efficiently running AI algorithms. They are designed to be the brains of any AI application including, robots, smart drones, smart cameras, and portable medical devices (NVIDIA, n.d.-d). More robust applications are seen in NVIDIA's complete data center products, like the DGX System. These can be purchased to include AI training already accomplished, fitting a client's AI needs directly out of the box (NVIDIA, n.d.-c).

Not only are these products purchasable at NVIDIA's website, there is a community for support and forums that can be used for free assistance; in addition to the consultative and customization work offered by NVIDIA (NVIDIA, n.d.-b). Similar to other computer system developers, they have taken to production of the hardware, giving away the software, and capitalizing on the opportunities for customization and consultation. NVIDIA's approach to AI hardware, offers opportunities for garage startups, small business, and defense contractors to tinker and innovate using existing commercial products.

b. SparkCognition

SparkCognition Inc is a company that was develop in 2014, by CEO Amir Husain, to try and bring AI systems to life (SparkCognition Inc, n.d.). It is based in Austin, Texas and has produced three products so far, DeepArmor, SparkPredict, and SparkSecure (SparkCognition Inc, n.d.).

DeepArmor is an AI program that provides end point computer security against malware, both detecting and preventing (SparkCognition Inc, n.d.). It falls under the machine learning category of AI, in that DeepArmor, "has trained on millions of malicious and benign files" (SparkCognition Inc, n.d.) SparkPredict uses cognitive AI to

enable predictive analysis of machine production variances (SparkCognition Inc, n.d.). This enables real time self-adjustments of production machines to keep variances within thresholds, which can adjust to maintenance schedules, machine load patterns, and environmental conditions. SparkSecure aims to replace human security analyst behaviors and support security analyst teams, by including a cognitive layer to software security systems (SparkCognition Inc, n.d.). A key to finding and validating zero-day types of attacks lies in signature [attack creator] identification, which is at present a very manual process including creative queries to the right sets of available data (SparkCognition Inc, n.d.). Using a self-improving system reduces cycle times from identification to validation, which can be the difference between successful address detection or complete loss of signature data (SparkCognition Inc, n.d.).

While there are many companies starting to develop these systems, SparkCognition is aiming its intelligent systems for government and military specific use from shipboard maintenance, to energy security, to enabling enhanced cyber warfare. They have even hired Major General (Ret.) Kenneth W. Wisain after 33 years in the Air Force and National Guard to be the Chief Architect for Defense (SparkCognition Inc, n.d.). This is a small scale example of AI companies listening to DOD requirements and attempting to respond with durable AI systems for military and government use.

c. SOARTECH

Soar Technology Inc, (SOARTECH) out of Ann Arbor Michigan, is another company that is trying to focus AI technology to military requirements. The CEO Dr. Michael Van Lent, who previously worked at the Navy Center for Applied Research in Artificial Intelligence, started SOARTECH as an effort to develop software that can understand and anticipate human behavior (Soar Technologies Inc., 2017). Their AI approach is to provide a software development service, which includes intelligent agents that can be used to simulate human behavior in warfare as close as possible (Soar Technologies Inc., 2017). This makes their AI algorithms valuable for decision support systems and effective training for soldiers in dynamically changing environments (Soar Technologies Inc., 2017).

SOARTECH has already entered into contracts with the DARPA, ONR, and AFRL focused on maximization of human performance and interaction with unmanned systems (Soar Technologies Inc., 2017). Perhaps, most interesting is their ability to leverage the video game industry technologies and re-apply them for military specific applications. Recently, SOARTECH and Epic Games have teamed up to provide licensing, training, and development services for a product called Unreal Engine 4, aimed at government and military applications (Cowley, 2015). The Unreal Engine, is an industry standard for creating high fidelity, real-time experiences for video game players, and SOARTECH has 17 years of experience as military contractor (Cowley, 2015). The resulting technology is a suite of development tools, including free code access and royalty free use, that can be tailored through consultation or work for hire services, generating 3-D visualization, training simulators, and movie capabilities (Cowley, 2015).

D. CHAPTER II SUMMARY

From this research we can see that AI is a viable technology and field of study, focused on developing machines that behave and act rationally. Even military environments can be significantly influenced and enhanced by a systems of sensors, AI algorithms, effectors, and practice. Whether they are taking the form of robots, software agents, or immobots, these systems are ready to leverage knowledge bases and AI techniques as they transition into Technology Maturation and the DAS. These AI essentials are the precepts that must be understood for DAS success. There is a tremendous effort happening in AI research and development. The world's biggest companies and many eccentric conceptualists are defining the AI landscape that the DOD will draw its systems from. Commercial applications are rapidly growing and many are ready for incorporation in defense systems. AI is ready for defense acquisition, but is acquisition ready for AI?

III. JCIDS

Now that there is a shared understanding for what AI is, the definition and its essentials, AI's path through the JCIDS can be analyzed. To support the necessary analysis this chapter will provide a brief summary of the JCIDS including inputs, inner workings, and key outputs. Not all of the JCIDS processes are analyzed in this research, merely the portions that will impact the fielding of AI systems. A systems engineering decision making methodology will be applied to key inputs to the JCIDS in order to determine whether or not they include the essentials of AI. This is important because, while the DAS fields the systems we fight with, the JCIDS is used to define the requirements that activate the DAS (JCIDS Manual, 2012, p. 1). Using strategic influences, warfighter inputs, and pointed technological exercises it informs DAS with capabilities the DOD actually needs and when. JCIDS outputs, validated requirements, have a direct impact on DAS performance and dictate the problems that can arise during DAS functions.

A. JCIDS PRIMER

To understand why inputs to JCIDS require AI essentials, an appreciation of how JCIDS relies on those inputs and what it does with them is necessary. The role of JCIDS is to, "facilitate the timely and cost effective development of capability solutions for the warfighter" (CJCS, 2012, p. 1). From this goal, it is apparent that JCIDS works as an integral cog in defense acquisition, by defining capabilities that warfighters require to meet mission demands. To understand the path for AI, this research considers the inputs to the JCIDS, the JCIDS steps that create requirements, and how the JCIDS outputs validated requirements. Figure 13 displays an overview of the JCIDS.

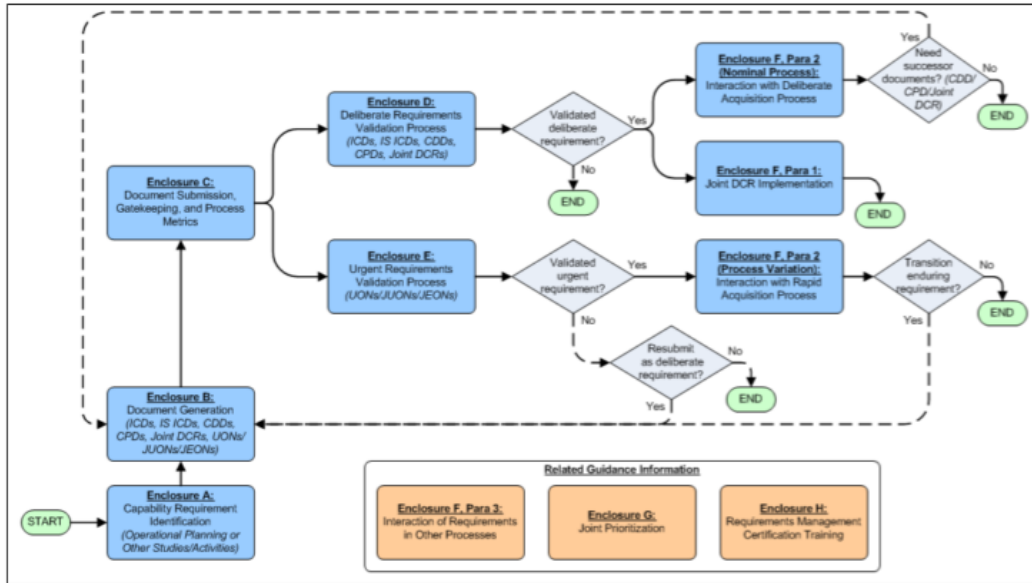


Figure 13. Overview of JCIDS. Source: CJCS (2012).

1. Inputs

Adding AI technology to warfighter capabilities, starts with identifying capability requirements (CJCS, 2012, p. 2). To execute the first step requires a mass of information, used to focus the JCIDS. This mass of information is known as the CONOPs (CJCS, 2012, p. A-B-2). Generally, the CONOPs gathers the military problem as well as other information and should consider the following concepts:

- the problem being addressed
- the mission
- the commander’s intent
- an operational overview
- the objectives to be achieved
- the roles and responsibilities of tasked organizations (CJCS, 2012, p. A-B-2).

While a CONOPs is very widely defined in the CJCS 3170.0I *Joint Capabilities Integration and Development System* and *JCIDS Manual*, it is important that CONOPs information answer the questions, “what operational outcomes they [capability requirements] provide, what effects they must produce, how they complement the integrated joint/multinational warfighting force, and what enabling capabilities are required

to achieve the desired operational outcomes” (CJCS, 2012, p. B-10). In the case of AI, nearly any document or information source, including intuition, could be used to build an effective CONOPs. Intuition of JCIDS personnel though, is difficult to analyze in a quantitative sense so it would make for a poor standard. For this research, the inputs analyzed are technologically focused strategic documents, joint capability technology demonstrations, and exercises aimed to provide a lessons learned for the JCIDS. The three areas are specifically referenced in the JCIDS sectioned that defines how identification of capability requirements should happen (CJCS, 2012, pp. A-1 – A-6). The CJCS JCIDS Manual defines these inputs as working in conjunction with the CBA to development high level program requirements, therefore they are suitable to be considered the make-up of a viable AI CONOPs (CJCS, 2012, p. A-4). These inputs are best positioned to inform the fit and function of potential AI solutions (CJCS, 2012, p. A-3).

2. Generating Requirements

To turn the CONOPs into significant and usable performance standards, JCIDS uses a mechanism called the capabilities based assessment (CBA) (CJCS, 2012, p. 2). The CBA is defined as, “an analytic basis to identify capability requirements and associated capability gaps (CJCS, 2012, p. A-4). This means, it takes the operational planning, technology demonstrations, and available exercises and assesses if the doctrine, organizational structure, training, materiel, leadership and education, personnel, facilities, and policy (DOTmLPF-P) support the future mission (CJCS, 2012, A-4). If capability gaps are identified the CBA works to solve the DOTmLPF-P problem with a joint Document Change Request (DCR), or a validated requirement document in the case of materiel gaps (CJCS, 2012, A-6).

3. Requirements Documents

The JCIDS manual explains the results of a CBA when a materiel capability gap is discovered, citing the, “results of a CBA or other study provide the source material for one or more Initial Capabilities Documents (ICD)” (CJCS, 2012, p. A-4). These requirements can take many forms, from the recommendation to transition systems from one branch to another or to field an entirely new system (CJCS, 2012, p. A-4). Two key

document outputs are the ICD and CDD (CJCS, 2012, p. A-4). These two requirements documents are the ones that activate processes inside the DAS (DOD, 2015, p. 6). All outputs from the JCIDS must be approved by the Joint Requirements Oversight Council (JROC) (CJCS, 2015, p. 1).

More than any other JCIDS output, the ICD is the most impacted by the quality of AI understanding used as an input to the JCIDS process and most critical to the DAS. This is because the publication of the ICD initiates the multiple processes inside of DAS, including funding requests and the expense of that money (DOD, 2015, p. 6). Without an ICD, a validated requirement, there would be no cause for a program to exist in the DAS. Additionally, the ICD and CDD carry forward to DAS the specifics of the CONOPs related to the system, that are must-have performance criteria (CJCS, 2012, p. B-10 – B-30). The ICD and CDD provide the contextual data that the entire acquisition program is framed on. The implications of low quality AI considerations in the CONOPs, are carried through the JCIDS process and to the DAS on ICDs and CDDs. Chapter IV analyzes the consequences of that logic.

B. METHOD

The method for analyzing a document for quality is a relatively subjective effort, but must be consistent and thorough. To achieve this, we are going to borrow from a decision management process outlined by the International Council on Systems Engineering (INCOSE) in their Systems Engineering Handbook (Haskins, Forsberg, Krueger, Walden, & Hamelin, 2011, p. 202). This technique, usually used to rack and stack design alternatives, can be equally good as a quantitative analysis mechanism for the quality of documents or CONOPs (Haskins et al., 2011, p. 207). The INCOSE Systems Engineering Handbook (Haskins et al., 2011) lists nine generic steps that can be used to work through a decision process:

- Frame the decision context, scope, constraints
- Establish communication with stakeholders
- Define evaluation criteria
- Define alternatives and select candidates for study
- Define measures of merit and evaluate selected candidates

- Analyze the results, including sensitivity analysis, and select best alternatives
- Investigate consequences of implementation
- Review the results with stakeholders and re-evaluate, if required
- Use scenario planning to verify assumptions about the future (p. 209).

Reviewing documents, technology demonstrations, and exercises for quality AI considerations will not require all of the steps, but they are still useful in generating the scoring criteria. This research will address the pertinent steps in order and that they apply.

a. Frame the Decision Context

Framing the decision means to clearly articulate the decision that needs to be made (Haskins, 2011, p. 209). This analysis aims to determine how well strategic documents, exercises, and technology demonstration are representing the essentials of AI presented in Chapter II of this project. The scope focuses on key inputs to the JCIDS process that originate across the world of DOD planning centers. Ideally, this research would also grade the outputs from the JCIDS process, but they are generally not for public release. That fact constrains the analysis to JCIDS inputs.

b. Establish Communication with Stakeholders

The primary stakeholders for this analysis are defense acquisition personnel. This includes future students at NPS, JCIDS personnel, DAS personnel, and those that output strategic documentation or design joint exercises. While they have not input their opinions for analysis, the idea is to consider them as customers requesting this type of analysis. Communication will be done by publishing the results of this research and analysis.

c. Define Evaluation Criteria

The INCOSE Systems Engineering Handbook (Haskins et al., 2011) states that selection criteria are the key desirable characteristics you want alternatives to have (p. 210). For this analysis, selection criteria are the AI essentials that strategic documents, technology demonstrations, and exercises should be considering during their execution. The AI essentials outlined by Chapter II will fill that role. For JCIDS to generate

requirements documents that lead to AI systems, the inputs should not only include the AI essentials, but also refer to them in the correct terms. These are the essentials extracted from Chapter II

AI Spectrum: Each technology proposed in the JCIDS inputs will fall on the spectrum of AI, from Automatic to Autonomic. Quality AI spectrum consideration goes beyond the use of the word drone or autonomous and recognizes the self-management properties contained in Autonomic Systems.

AI Definition: The definition of AI proposed in Chapter II captures what is meant by AI, when referring to its position in Defense Acquisition. A document could ignore the aspects of this definition, creating poor outputs in JCIDS, or articulate well the dynamics of AI, creating better outputs of the JCIDS process.

AI Components: AI systems consists of five major components (Sensors, AI Algorithms, Effectors, Environment, and Practice). Quality JCIDS inputs should be able to talk about these components inside of a CONOPs to correctly inform requirements generation.

AI Classes: The three classes of AI systems (Robots, Immobots, and Software Agents) each convey a different messages concerning physical requirements, distribution of sensors and actuators, and mobility. For the JCIDS process to correctly capture requirements it must articulate the difference between these classes.

AI Algorithms: An AI Algorithm consists of a special type of basic computer algorithm and a knowledge base. Quality JCIDS inputs will go further than referencing an algorithms and clarify that there is a need for select algorithms and a knowledge base.

AI Techniques: An AI technique is a specific use of an AI algorithm. Individual techniques excel in different ways, so it is important for the requirements generation process to convey the appropriate technique in each situation.

Each alternative, JCIDS input, is broken into elements. For this research an element is classified as *an instance, where AI technology is eluded to or should be included*. This could be a sentence, paragraph, figure or chart. It can be a video clip representing what happened at an exercise or an image that captures the complexity of any given document, vignette, or exercise. The document itself will define the propinquity of material, allowing for a person trained in the essentials of AI to understand the linkages and segregations natural to the artifacts that influence the JCIDS. It is important to break these JCIDS inputs into elements because some portions of them are not and should not be concerned with AI, they have different constructs, and different sections may have different levels of AI consideration. Admittedly, decisions surrounding elements will be the most subjective portion of analysis. For that reason, the generic term element is used to represent items elementary to the ability of an artifact to convey its content. It lets the document or exercise itself define how to be chunked based on its unique structure. The element is first described, then subjected to the evaluation criteria. There are different numbers of elements for each JCIDS input, so the evaluation criteria are averaged to give an overall sense of how well the document, exercise, or demonstration represents that characteristic of AI. Table 1 displays the construct for the evaluation criteria and elements.

Table 1. Sample Evaluation Chart Displaying Evaluation Criteria

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique	
1								
2								
...								
...								
n								Weighted Score
Average								

The final consideration for Evaluation Criteria, is that not each AI essential should be equally represented in these documents. Since, these are generally strategic inputs to the JCIDS process, they shouldn't have much specified in terms of specific types of algorithms, but should definitely point to the level autonomy desired. The evaluation

criteria are weighted to control their impact on the analysis, based what AI essentials are more strategic in nature. Table 2 displays the weights that each evaluation criteria hold.

Table 2. Criterion Weighting

AI Attribute	Weight
Spectrum	0.25
Definition	0.25
Components	0.25
Classes	0.1
Algorithm	0.1
Technique	0.05

d. Define Alternatives and Select Candidates for Study

The selection process for strategic documents, exercises, and demonstrations is done by searching for these documents and releases published at the requirements generation centers. These centers are the units and organizations that enact the JCIDS. Each service has a different structure and conducts this in a slightly different way, so it is somewhat subjective to define one strategic document or another as more influential. Generally, an object worth analyzing defines the future operations expected to be conducted by a service and spells out at least some of the expected performance standards that will be required. At least one strategic document from each service, with CONOPs published inside, is analyzed to allow for diversification across services. Many strategic documents are available for each service so those that act to inform future technology use are targeted. Published exercises that aim to incorporate lessons learned into JCIDS decisions, are sparser and diversification is not as easy. For that reason, more recent exercises are analyzed and technology demonstrations on the brink of execution are analyzed. This selection is not meant to be comprehensive, but rather a general inclination of whether or not the AI essentials are working their way in to the JCIDS and at what level of quality.

e. Define Measures of Merit and Evaluate Selected Candidates

The scoring for each element mostly copies the INCOSE Systems Engineering Handbook (Haskins et al., 2011) by rating each element on a 1–9 scale (p. 211). INCOSE actually recommends 1–10, but military technology readiness assessments use a 1–9 scale (AcqNotes, n.d.). Each score also has a qualitative attachment to clarify why a particular score is assessed. The qualitative portion is essential because inputs to the JCIDS normally have a narrative style. To assess the quality at which they address AI requires a less discrete formulation. This qualitative grade is the code used to apply a score. Additionally, each score is assigned a color, allowing for a reader to understand the level of quality at a glance. The standard red, yellow, and green colors are used. Table 3 outlines the scoring key.

Table 3. Scoring Legend

Score	Qualitative Grade	Color
1	Not addressing AI	Red
2	Poorly Addressing AI	Red
3	Under Addressing AI	Red
4	Marginally Addressing AI	Yellow
5	Addressing AI	Yellow
6	AI addressed	Yellow
7	AI Well Addressed	Green
8	AI Understood	Green
9	AI Well Understood	Green

f. Analyze the Results, Including Sensitivity Analysis, and Select Best Alternatives

The results of this analysis are discussed at the end of the chapter, without using sensitivity or making a selection. The result is several grades for JCIDS inputs that give the sense of how prepared each alternative analyzed is to inform about AI. Finally, a composite score for the overall AI CONOPs is developed by averaging the composite score for each individual JCIDS input. This offers a total sense of the AI essentials headed through the JCIDS process. Further, this analysis will also unveil the areas where

AI essentials are not addressed. This will highlight the essentials that are not well represented and subsequently offers an area where those preparing JCIDS inputs can improve.

g. Investigate Consequences of Implementation

The consequences of whether or not AI essentials are included in JCIDS inputs, is done in Chapter IV. Since JCIDS informs and activates the DAS, it is expected that the consequences will be realized during the DAS processes. The steps 8. Review the results with stakeholders and re-evaluate, if required and 9. Use scenario planning to verify assumptions about the future are not done. They are outside of the scope of this research.

C. ANALYSIS

With this methodology in hand an analysis of the JCIDS inputs can be conducted. The analysis is broken into two portions: AI in strategic documents and AI in exercises. AI in exercises includes both exercises that work to transfer lessons learned into JCIDS and technology demonstrations.

1. AI in Strategic Documents

The JCIDS Manual specifies several strategic inputs. It states the National Security Strategy, National Strategy for Homeland Security, National Defense Strategy or the most recent Quadrennial Defense Review Report, National Military Strategy, Defense Planning Guidance, Guidance for the Employment of Force, Chairman’s Risk Assessment, Joint Strategic Capabilities Plan are all important to provide a framework for JCIDS assessments (CJCS, 2012, p. A-1). Many of these documents are written in such a generic way, that the mention of “technology” could be eluding to AI Systems. For that reason, this research will key in to defense planning guidance that includes CONOPs and sets the stage for AI essential incorporation. The documents analyzed are the Unmanned Systems Integrated Roadmap, The U.S. Army: robotic and Autonomous Systems Strategy, Technology Horizons: A Vision for Air Force Science and Technology 2010–30, A Cooperative Strategy for 21st Century Seapower, 2014 Defense Intelligence Agency Innovation Strategic Plan, and the Marine Corps Operating Concept.

a. Unmanned Systems Integration Roadmap

The *Unmanned System Integration Roadmap* was published in 2011 and aims to direct Acquisition and JCIDS efforts running through the year 2036 (Chairman of the Joint Chiefs of Staff, 2011, p. i). It starts by describing the vision that the Under Secretary of Defense for Acquisition, Technology, and Logistics and Vice Chairman of the Joint Chiefs of Staff have for the integration of unmanned systems and technology into the Joint force structure (CJCS, 2011, p. 1). It not only points to the idea that this plan is expected to meet DOD affordability goals, but also where the DOD and Industry should be looking to make investments (CJCS, 2011, p. 1). It is worth noting that this is a combined effort between the owners of the JCIDS and DAS systems, and used many sources like surveys, previous and service specific roadmaps, and various service specific CONOPs (CJCS, 2011, p. 1). The result is a document that speaks of strategic direction for unmanned systems in three domains (Air, Land, Sea), including specific CONOPs that provide working examples of what that strategy would look like (CJCS, 2011, p. 2). It is the single unified source defining the problem sets facing unmanned systems and the ways the DOD believes we can maximize military utility from these systems (CJCS, 2011, p. 2).

Elements were broken down by the natural flow of the document, based on where allusion to AI were placed or should have been placed. There were elements taken from the vision statement, from whole chapters, and sub-chapters that each addressed AI to some degree. The analyzed elements can be seen in Table 4.

Table 4. Analysis of the Joint Unmanned Systems Integration Roadmap

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique	
1	Vision Statement	5	3	1	1	1	1	
2	DoD Vision	5	3	1	7	1	1	
3	Interoperability Across Domains Vignette, 2030	5	7	8	1	3	1	
4	African Maritime Coalition Vignette, 2030	5	7	6	1	3	1	
	Complex Unmanned Systems Test and Evaluation Scenario	5	6	7	8	8	6	
5	Current State: Requirements Definition and Acquisition system	5	1	1	1	1	1	
6	Current State: Logistics	5	1	1	1	1	1	
7	Current State: Autonomy	5	6	7	1	1	1	
8	Interoperability: The Way Ahead	5	5	5	6	5	1	
9	Autonomy	8	7	7	1	4	8	
10	Airspace Integration: Way Ahead	5	1	5	1	3	1	
11	Communication: Way Ahead	6	6	1	1	1	3	
12	Training: Way Ahead	5	1	1	1	1	4	
13	Propulsion and Power: Way Ahead	1	1	1	1	1	1	
14	Manned-Unmanned Teaming: Way Ahead	5	1	5	1	1	2	
15								
Average								Weighted Document Score
		5.0	3.7	3.8	2.2	2.3	2.2	3.7

The composite score is 3.7 indicating under addressing AI. It described most functions in CONOPs and technology development in terms of autonomous behavior rather than autonomic. The only exception is the autonomy section. Joint documents are usually composed of many peoples work, and this particular section covers very well the essentials of AI, that this research indicates as key to defense acquisition. The areas of AI that are poorly or not addressed are classes, algorithms, and techniques. Even the basic understanding of the differences between AI systems by definition and autonomous systems seems to be elusive. Sensors as an AI component, are often well addressed, but the idea that an algorithm produces an output that can be enacted by an effector mostly did not appear. Further, immobots and software agents were not addressed except in the element dedicated to Test and Evaluation. The *Unmanned System Integration Roadmap* must address AI essentials more completely, before JCIDS process can effectively transfer these concepts into capable AI systems.

b. The U.S. Army: Robotic and Autonomous Systems Strategy

In March 2017, the U.S. Army published a forward looking strategic support document titled *Robotic and Autonomous Systems Strategy* (Army Capabilities Integration Center [ARCIC], 2017). The document describes how the Army aims to provide overmatch in future wars using robotics, automation, and intelligent technologies (ARCIC, 2017, p. i). The forward is written by the Vice Chief of Staff, General Daniel B. Allyn, and sets the stage outlining that robotics and autonomous systems will both increase the effectiveness of Army forces and the means by which they acquire and

deliver support to soldiers (ARCIC, 2017, p. i). This publication by the Army’s lead requirements analysis center is positioned to be the foundation for future Army program initiation. It having the essentials of AI will be paramount for Army success in the AI domain.

Elements for this document take on a different form than most of the other documents. This is because the Army used essentials of AI in many small, but important areas. A notable example is the Title Page and Title of the document. It is very clear that the Army is perusing essentials of AI from the very first words. The forward is also included because General Allyn references his opinions about AI and its usefulness. The analysis includes three very valuable vignettes, dedicated to discrete time periods, and inclusive of operational view (OV-1) depictions of each. The scoring for the Army’s Robotic and Autonomous Systems Strategy is found in Table 5.

Table 5. Analysis of *The U.S. Army Robotics and Autonomous Systems Strategy*

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique	
1	Title Page	6	5	8	4	1	1	1
2	Forward by Gen Daniel B. Allyn	5	4	4	1	1	1	1
3	Section 1: Why the Army Pusues Robotic and Autonomous Systems	5	1	4	5	1	4	4
4	Section 2: Intro	7	9	6	7	7	6	6
5	Section 2: Near Term (2017-2020)	6	6	8	3	4	2	2
6	Vignette: Urban Operataions (2025)	5	4	6	5	1	1	1
7	Section 2: Mid Term (2021-2030)	7	7	6	3	4	7	7
8	Vignette: Setting the Theater in Future Crisis	4	1	1	1	4	1	1
9	Section 2: Mid Term (2031-2040)	5	8	4	3	4	3	3
10	Vignette: Reconnaissance and Security Operations	5	2	6	3	1	3	3
11	Section 4: RAS and Interim Solutions to Army Warfighting Challenges	5	5	8	3	2	3	3
Average		5.5	4.7	5.5	3.5	2.7	2.9	4.7

The composite score for *The U.S. Army Robotics and Autonomous Systems Strategy* is 4.7, indicating a document marginally addressing AI. While it does address many essentials of AI, in almost every element of the document it, scored poorly in terms of AI classes, AI algorithm dynamics, and AI techniques. Most often excluded, are the uses for software agents and the necessity to include a knowledge base to enable AI algorithms. Still, it scored well in terms of spectrum understanding and AI components. In these areas it is absolutely addressing AI, if only at a basic level of understanding.

Considering this is a strategic document, to have a strong representation of the understanding that systems will be self-prioritizing and will include sensors, algorithms, and effectors is a good sign. The key shortages in these two domains was the understanding that autonomic systems will have the ability to branch behaviors and practice is required to achieve maximum effect. Re-assessment of these essentials would greatly help this document and increase the usability of *The U.S. Army Robotics and Autonomous Systems Strategy* as a source document for successful AI requirements definition.

c. Technology Horizons: A Vision for Air Force Science and Technology 2010–30

Technology Horizons is a U.S. Air Force publication that aims to outline key science and technology focus areas that should be perused between the years 2010 and 2030 (Office of the US Air Force Chief Scientist, 2011). Generally, it is released every 15 years to reset Air Force vision for technology and clearly articulate overarching technology themes based on the strategic landscape and potential enduring realities (Office of the US Air Force Chief Scientist, 2011, p. ix). Together the Secretary of the Air Force and USAF Chief Scientist define the role and effort of the document and claim, “The future is ours to shape” (Office of the US Air Force Chief Scientist, 2011, p. ix). It is true, this document will shape future systems headed for the DAS and attempts to define them in terms of AI. Table 6 displays the elements and analysis for *Technology Horizons*.

Elements are based on the primary chapters and many sub-chapters. The chapters *Overarching Themes for Air Force S&T 2010–30*, *Technology-Enabled Capabilities for the Air Force 2010–30*, and *Grand Challenges for Air Force S&T 2010–30* comprise most of the many elements. The sub-chapters stood alone due to their size and scope. If the entirety of the chapters were analyzed, they would have appeared much more adept at addressing AI essentials than they actually are. Elements are chunked into complete ideas to get the best sense of how well AI is addressed for each independent concept and war fighting domain. The preface and executive summary are also analyzed due to the efforts of the authors to stress autonomy in those sections.

Table 6. Analysis of Air Force *Technology Horizons*

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique	
1	Forward by Secretary of the Air Force	5	1	1	1	1	1	1
2	Preface	5	1	1	1	1	1	1
3	Executive Summary	6	3	1	1	1	1	1
4	Strategic Context: Technology-Derived Challenges to Air Force Capabilities	5	1	1	1	1	1	4
5	Enduring Realities for the Air Force: Manpower Costs	5	1	1	1	1	1	1
6	Overarching Themes for Air Force S&T 2010-30: From Manned to Remotely Piloted	5	3	1	1	1	1	1
7	Overarching Themes for Air Force S&T 2010-30: From Fixed to Agile	1	1	1	1	1	1	1
8	Overarching Themes for Air Force S&T 2010-30: From Control to Autonomy	5	8	4	1	1	1	2
9	Overarching Themes for Air Force S&T 2010-30: From Integrated to Fractionayed	1	1	1	1	1	1	1
10	Overarching Themes for Air Force S&T 2010-30: From Preplanned to Composable	5	1	1	1	1	1	1
11	Overarching Themes for Air Force S&T 2010-30: From Sensor to Information	5	1	4	1	1	1	3
12	Overarching Themes for Air Force S&T 2010-30: From Cyber Defense to Cyber Resilience	1	2	1	1	1	1	1
13	Overarching Themes for Air Force S&T 2010-30: From Long system Life to Faster Refresh	1	1	1	1	1	1	1
14	Technology-Enabled Capabilities for the Air Force 2010-30: Technology Enabled Potential Capability Areas	5	7	5	1	1	1	3
15	Technology-Enabled Capabilities for the Air Force 2010-30: Automated Cyber Vulnerability Assessments and Reactions	6	1	1	1	1	1	3
16	Technology-Enabled Capabilities for the Air Force 2010-30: Decision Quality Prediction of Behavior	5	3	1	1	1	1	6
17	Technology-Enabled Capabilities for the Air Force 2010-30: Augmentation of Human Performance	2	1	3	4	2	2	2
18	Technology-Enabled Capabilities for the Air Force 2010-30: Advanced Constructive Discovery and Training Environments	1	1	1	1	3	4	4
19	Technology-Enabled Capabilities for the Air Force 2010-30: Trusted, Adaptive, Flexibly Autonomous Systems	5	5	1	1	6	1	1
20	Technology-Enabled Capabilities for the Air Force 2010-30: Processing-Enabled Intelligent ISR Sensors	1	1	1	1	1	1	2
21	Technology-Enabled Capabilities for the Air Force 2010-30: Embedded Diagnostic/Prognostic Subsystems	5	1	4	1	1	1	2
22	Technology-Enabled Capabilities for the Air Force 2010-30: Improved Orbital Conjunction Predictions	1	1	4	1	1	1	1
23	Key Technology Areas Supporting Potential Capability Areas	7	7	6	1	4	5	5
24	Grand Challenges for Air Force S&T 2010-30: Challenge 1, Inherently Intrusion-Resilient Cyber Networks	5	2	2	1	1	1	1
25	Grand Challenges for Air Force S&T 2010-30: Challenge 2, Trusted, Highly Autonomous Decision Making Systems	7	5	4	1	3	8	8
26	Grand Challenges for Air Force S&T 2010-30: Challenge 3, Fractionated, Composable, Survivable, Autonomous Systems	6	1	4	1	1	1	1
27	Grand Challenges for Air Force S&T 2010-30: Challenge 4, Hyperprecision Aerial Delivery in Difficult Environments	5	1	1	1	1	3	3
Average		4.1	2.3	2.1	1.1	1.5	2.3	Weighted Document Score 2.5

The composite score for *Technology Horizons* is a 2.5, indicating a document that poorly addresses AI. Only once are AI classes addressed. This points to a critical misunderstanding on how AI systems can behave. Robots, immobots, and software

agents each have a different perspective and mobility capabilities. Algorithms scores a 1.5, which is essentially not addressed. To enable any form of beyond autonomous behavior specific algorithms must be used and infrastructure investments must be made toward multi-system knowledge base access. It may seem immaterial for a strategic document to address a knowledge base in specific terms, but it is absolutely imperative to address information databases and accessible architectures, or these capabilities can be forgotten during requirements formulation. Understanding that the AI spectrum beyond automated behavior includes self-aware and potentially branching systems, was also missing from *Technology Horizons*. The most well developed element is number 23. This include descriptions specific to AI, abstractly three of the four AI components, and honed in on specific AI techniques that could be useful (Office of the US Air Force Chief Scientist, 2011, p.118). It is worth noting that Air Force Technology Horizons had the most elements analyzed. This generally decreases the scoring, in that smaller sections allow for less aggregation of concepts. Extension of the terms and understandings used in element 23, concerning trusted highly autonomous decision systems, would benefit the entirety of the document.

d. A Cooperative Strategy for 21st Century Seapower

This strategic document produced by the Office of the Secretary of the Navy in 2015 is designed to inform how the Navy intends to design, organize, and employ sea services to support the United States (Office of the Secretary of the Navy, 2015, p. iii). Generally, this document would be considered a higher level document, than one that directly informs the JCIDS. The introduction, however, sends the message that it proposes to not only inform the JCIDS, but also in a technical sense. It claims, “[A *Cooperative Strategy for 21st Century Seapower*] describes how naval forces will enhance their effectiveness, employ new warfighting concepts, and promote innovation” (Office of the Secretary of the Navy, 2015, p. 2). This may not specify the technologies we should use, as most strategic documents don’t, but how the processes will engage and what technologies they aim to engage with. For that reason, and its nautical persuasion, makes it an ideal candidate for scoring against the AI essentials outlined in this research.

Elements for this document are relatively few, mostly due to its high level nature. Still, very specifically it defines technology and capabilities expected of the Navy in its future state. These elements, along with the introduction and the general expected future use of seapower are analyzed. The results of the analysis are in Table 7.

Table 7. Analysis of *A Cooperative Strategy for 21st Century Seapower*

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique	
1	Introduction	1	1	1	1	1	1	
2	Seapower in Support of National Security	1	1	1	1	1	1	
3	Force Design: Building the Future Force: Flexible, Agile, and Ready Forces	3	1	3	1	1	1	
4	Force Design: Building the Future Force: Concepts	1	1	1	1	1	1	
5	Force Design: Building the Future Force: Capabilities	5	1	4	1	1	1	
Average		2.2	1.0	2.0	1.0	1.0	1.0	1.6

The composite score for *A Cooperative Strategy for 21st Century Seapower* is a 1.6. Qualitatively this is not addressing AI. Very few instances speak in terms of AI, and very often utilization of terms representing the current standard of autonomous efforts are forgone. The highest scoring section is the AI spectrum. A score of 2.2 indicates the document is considering systems that perform automatic functions. This area of the AI spectrum would relate to machines that replace simple processing functions formerly done by humans. The scoring strategy is relatively subjective, based on the knowledge of AI gained from earlier portions of this research. Still a rule or two surfaced. The word “autonomous” when used to indicate a system capable of carrying out a series of tasks and prioritizing actions to achieve it’s given goal, is scored as a 5. Most of the documents analyzed had little trouble scoring well in the AI spectrum category because envisioning a future of autonomous vehicles that support operations is not a significant stretch from current technology capabilities. *A Cooperative Strategy for 21st Century Seapower* refrains from defining future capabilities as autonomous, with one notable exception, and has little to no other introductions of other AI terms or essentials. This document would not perform the job of informing JCIDS personnel to consider AI essentials.

e. **2014 DIA Innovation Strategic Plan**

To try and sample strategic documents that live outside of the typical Army, Navy, and Air Force realm the Defense Intelligence Agency (DIA) *2014 Innovation Strategic Plan* is analyzed. The introduction by Lieutenant General Michael T. Flynn describes DIA innovation as, “a driver for productivity gains, the ability to do more with less” (Defense Intelligence Agency [DIA], 2014, p. ii). General Flynn also emphasizes that, “DIA must build agility into the core of all activities enabling us to address tomorrow’s challenges more efficiently and effectively” (DIA, 2014, p. ii). Over the course of 15 pages it lays out a vision for the future, several processes to support innovation, and defines an entire section to supporting agility of the organization.

There are two elements where General Flynn allowed for AI essentials to be addressed, his forward and the agility section (DIA, 2014, p. ii). Despite the General’s narrow focus, every section could have used the essentials of AI to support the future directions for DIA. Rather than score each section concerning processes about leveraging international partnerships and allowing innovation to bubble up from the ranks poorly, they were omitted. From an analysis standpoint this makes a poor example, but to test the tertiary efforts of the DOD a DIA document is ideal considering intelligence gathering runs a gamut from human collection to the various forms of sensor enabled collection. This research also reviewed the 2016 *Defense Intelligence Agency Strategy* and the *Defense Intelligence Agency Strategic Vision 2012–2017* with strikingly similar results. Elements including the essentials for AI were spurious, infrequent, and poorly addressed on all accounts. The natural instinct would be to exclude the DIA from analysis, but perhaps the fact that so few inclusions of AI are found throughout the DIA is equally telling of their plans for AI. The analysis for the *2014 DIA Innovation Strategic Plan* is found in Table 8.

Table 8. Analysis of the 2014 DIA *Innovation Strategic Plan*

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique	
1	Innovation Imperative	2	2	2	1	1	1	
2	Innovation Goal 2: Agility	5	1	1	1	1	1	
Average		3.5	1.5	1.5	1.0	1.0	1.0	1.9

This document achieved a cumulative score of 1.9 and qualitatively is not addressing AI. Techniques, classes, and algorithms are missing. Notable allusions to AI technology are found in Innovation Goal 2: Agility, addressing that autonomous functions will be required to achieve innovation (DIA, 2014, p. 5). Much of this score has to do with how the document is written, focusing on the efforts of DIA employees and processes to achieve future gains. This research does not suggest efforts to promote innovation through personnel empowerment are unimportant, but rather simply scoring on how well strategic documents will support the JCIDS. The *2014 DIA Innovation Strategic Plan* is not addressing AI.

f. Marine Corps Operating Concept

The *Marine Corps Operating Concept* explains how the Marine Corps expects to operate as an expeditionary force in the 21st Century (Marine Corps, 2016). It is a forward looking document that outlines the future security environment, what is expected to drive change, and what shape the future Marine Corps should look like (Marine Corps, 2016, p. ii). This includes the identification of a central problem facing Marines and the identification of five critical tasks that must be achieved to solve that problem (Marine Corps, 2016, p. ii). It uses an interesting question and answer CONOPs to start the discussion. The landscape is set to 2026 and a moderator is interviewing marines after an operation; their responses reveal what the landscape of the future looks like from a Marine's eyes (Marine Corps, 2016, p. 1).

Elements were derived from the section headings and pertinent sub-headings. As with most strategic documents the chapters work naturally as element partitions. One chapter in particular, concerning the five critical tasks required for future Marine success was broken into five elements, as each task was a discrete thought (Marine Corps, 2016, p. 10). The forward from General Robert B. Neller was also included because it articulates the ideals Marines leadership are aiming for (Marine Corps, 2016, p. i). The result is 11 elements. The results of the analysis can be found in Table 9.

Table 9. Analysis of *Marine Corps Operating Concept*

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique
1	Foreword (Gen Robert B. Neller)	1	3	1	1	1	1
2	Looking Ahead - How we will operate and Fight in 2025	5	1	4	5	1	1
3	The Future Security Environment: Key Drivers of Change	5	1	2	4	3	2
4	Statement of the Central Problem	1	1	2	1	1	1
5	Our Operating Concept: Maneuver Warfare in Every Dimension; Combined Arms in All Domains	5	1	2	1	2	1
6	Creating the Future Force: Critical Tasks and Issue Areas: Intro	1	2	1	1	1	1
7	Creating the Future Force: Critical Tasks and Issue Areas: Critical Task: Integrate the Naval Force to Fight At and From the Sea	3	1	2	1	1	1
8	Creating the Future Force: Critical Tasks and Issue Areas: Critical Task: Evolve the MAGTF	5	3	1	4	2	2
9	Creating the Future Force: Critical Tasks and Issue Areas: Critical Task: Operate with Resilience in a Contested-Network Environment	3	1	5	3	3	4
10	Creating the Future Force: Critical Tasks and Issue Areas: Critical Task: Enhance Our Ability to Maneuver	6	2	2	1	1	1
11	Creating the Future Force: Critical Tasks and Issue Areas: Critical Task: Exploit the Competence of the Individual Marine	3	1	2	1	1	4
Average		3.5	1.5	2.2	2.1	1.5	1.7

The *Marine Corps Operating Concept* composite score is 2.2. While the entirety of the document included AI essentials throughout, the vocabulary was often over generalized or not sufficient. The opening CONOPs was both interesting and inclusive of AI essentials. It dealt with automatous machine behavior, including an understanding that sensors will be engaging with environments and robots can perform, perceive, and operate in a mobile fashion. Element 10 also reiterated the need for autonomous systems and AI support. The key factor missing from the document is the AI lexicon. In many elements it addresses AI systems and support, but does so without a solid understanding or the terms or functions inherent in AI. For this reason, it poorly addresses AI.

2. AI in Exercises

Joint exercises and technology demonstrations work together to support CBA analysis (CJCS, 2012, p. A-A-5). They deliver a very tangible and recognizable input to the overall CONOPs, which provides depth to requirements and grounds them in practical scenarios. The source for these exercises and demonstrations are similar to strategic documents, originating at the requirements generation centers. For exercises, this research focuses on the exercises hosted by requirements generation agencies that directly tie to the efforts of supporting future materiel acquisition. There are many types

of exercises across the military that can influence JCIDS, through accidental discovery or planned doctrinal testing, but the portion that pertains to AI is the portion concerned with helping JCIDS define future technological system requirements. Similarly, demonstrations come in many shapes and forms. Trade shows and defense acquisition competitions are examples. This research will focus on two examples of demonstrations that are actively working to encourage development of future warfighter technologies. The exercises analyzed are the Army's Unified Quest 2016, the 2016 Marine MAGTF Exercise, Air Force Red Flag Exercise, and TechWarrior 17 Exercise. The demonstrations analyzed are ThunderDrone and DARPA's Cyber Grand Challenge.

a. UQ 16 Future Force Design II Final Report

The *UQ 16 Future Force Design II Final Report* is a publication that outlines the findings of an ARCIC hosted exercise called UQ16. While, ARCIC conducts many exercises and demonstrations to support future materiel decisions, the Unified Quest series is a comprehensive look at the generic assessment of future conditions and best representation of a complete Army approach. Unified Quest 2016 Future Force Design Exercise has the purpose of examining how armies, corps, and divisions will operate in 2030 and seeks to inform development of operational and organizations concepts (Army Capabilities Integration Center Future Warfare Division [ARCIC Future Warfare Division], 2016, p. 1). This includes conducting a broad review of political and technological trends that can be incorporated to defeat the four scenario extracts included in the document (ARCIC Future Warfare Division, 2016, p. 1). The end result for Unified Quest Exercises are capability gaps, strategic responses, and lessons learned that influence the overall CONOPs for Army acquisition (ARCIC Future Warfare Division, 2016, p. 1).

Elemental breakdown for this exercise is similar to the strategic documents previously analyzed. Most of the pertinent information is contained in a single after action report, complete with a table of contents, chapter headings, and sub-headings. Since, all of the available classified information is in one place the headings and sub-headings are used as elements. Table 10 displays the results of the analysis.

Table 10. Analysis of *UQ 16 Future Force Design II Final Report*

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique	
1	Introduction	1	1	2	1	1	1	1
2	Method	1	1	2	1	1	1	1
3	Results: O&O Concept Insights	6	1	2	5	1	1	1
4	Results: Army Warfighting Challenge Insights	1	2	1	1	1	1	1
5	Results: Operational Environment	1	1	1	1	1	1	1
6	Results: Implications for an Army Conducting Expeditionary Maneuver	1	1	1	1	1	1	1
7	Results: Emerging Insights into Future Requirements	1	1	1	1	1	1	1
Average		1.7	1.1	1.4	1.6	1.0	1.0	1.4

The composite score for the *UQ 16 Future Force Design II Final Report* is a 1.4, indicating it is not addressing AI. While the document does have an instance of highlighting the need for robotic assistance and autonomous support (element 3), it generally focuses on the domains that a future army is expected to fight in. Outside of highlighting the expected future environment, AI is not considered. Perhaps UQ16 was initially missing the directive of highlighting AI as an aspect of future wars. This should not be the case, given the Army’s 20 defined warfighting challenges on the ARCIC webpage. Warfighting Challenge #4 has eight learning demands centered on creating an adaptable and innovative force (Army Capabilities Integration Center [ARCIC], n.d.). At a minimum, AI essentials should be included in element seven dealing with insights for future requirements. In terms of AI, this is a missed opportunity.

b. Marine MAGTF Integrated Experiment 2016

On 26 July 2016, Company K, 3rd Battalion, 5th Marine Regiment, Company Landing Team began a five-day experiment where they were reconfigured, re-equipped, and re-trained in order to test their abilities to enter an enemy controlled area (Schultz, 2016). The Marine Warfighting Laboratory fitted the Marines with 40 possible technologies to test surrogate techniques and develop Tactics Techniques & Procedures (TTPs), attempting to make the force smarter, faster and more lethal (Le, 2016). The Marine Warfighting Laboratory highlights many of these technologies on their website (United States Marine Corps, n.d.). Additionally, the Marine Corps hosts on its website a video of the action. Marines can be seen utilizing various Unmanned Ariel Systems (UAS) and unmanned ground vehicle (UGV) technologies to carry out their missions (Marine Corps Warfighting Laboratory, 2016). The whole experiment has an emphasis

on creative solutions in a condition of limited resources, while remaining combat effective (Le, 2016).

Unlike the UQ16 exercise conducted by the Army, the Marine Corps did not compose the results of the experiment into a singular and easily digestible document. Resources outlining the events of the exercise are mostly found in news articles, complete with interviews and pictures of the Marines executing the various missions. Elements are each individual picture and a pair of news articles. More articles would increase the validity of the events analysis, but after review of many artifacts it appeared mostly duplicative. Since the tools used by Marines were supplied by the Marine Warfighting Laboratory, the highlighted technologies on the Marine Warfighting Lab website are also analyzed. The results of the Marine MAGTF Integrated Experiment 2016 (MIX16) analysis are seen in Table 11.

Table 11. MIX16 Analysis

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique	
1	Coung Article	4	1	4	4	1	1	
2	Schultz Article	1	1	2	1	1	1	
3	The Warfighting Lab	6	8	7	4	1	2	
4	MIX16 Exercise Picture 1	1	1	1	1	1	1	
5	MIX16 Exercise Picture 2	1	1	1	1	1	1	
6	MIX16 Exercise Picture 3	1	1	1	1	1	1	
7	MIX16 Exercise Picture 4	1	1	1	1	1	1	
8	MIX16 Exercise Picture 5	1	1	1	1	1	1	
9	MIX16 Exercise Picture 6	3	4	6	4	1	1	
10	MIX16 Exercise Picture 7	1	1	1	1	1	1	
11	MIX16 Marine Corps Video	6	6	7	4	1	1	
Average		2.4	2.4	2.9	2.1	1.0	1.1	2.3

The composite score for the MIX16 is 2.3, indicating it is poorly addressing AI. One of the articles covers a considerable amount of the equipment used by the Marines in the experiment, including robot systems and unmanned aerial vehicles. It alludes to information gathering and the ability to effect the enemy force, which leads to the component score of 4. A higher score would have occurred with inclusion of sensors, effectors, algorithms, environment and practicing applications. Element 9 is a picture displaying a tracked robot and soldier teamed together. The robot has visible sensors and a weapon mounted, which elevates its component score. An interesting element is the Marine video depicting the scenario highlights. This is the first video evaluated and required a somewhat different sense to evaluate. The scoring was conducted from two

angles: the words spoken in and video images of systems operating. At one point a manned ground system complete with sensors and effectors, was employed using a practicing behavior. This is displayed in the component score of 6. Another, scene showed an unpiloted ground vehicle that followed Marines with its suite of sensors, causing it to increase the Spectrum grade from simply unmanned score of 4 to a fully autonomous score of 6. Algorithms and techniques are mostly missing from what is easily learnable about MIX16. The Warfighting Lab scored considerably stronger than other elements, as its role is to bring autonomous systems to Marines. Their inventory of technologies specifically highlights AI essentials and autonomous systems that follow goal oriented behavior (United States Marine Corps, n.d.). The efforts of the Marine Corps Warfighting Lab and the production of the MIX16 video that offer the best indicators of what really happened in MIX16. It is likely in this case, the score, built from a wide range of sources analyzed, does not match the true consideration of AI essentials during MIX16.

c. Air Force Red Flag 2016

The 414th Combat Training Squadron hosts an annual exercise called “Red Flag,” consisting of allied air combat operation, maintenance, and recently including intelligence, cyber, and electronic warfare functions (414th Combat Training Squadron, 2012). While the main driver of the event is preparation for air combat using the existing systems available in the air force inventory, new iterations of the exercise have included the testing of new capabilities (Bultman, 2017). A video debrief posted on Youtube.com by the 414th Combat Training Squadron, using a question and answer session, highlights the partner nations involved and the effects the F-35 had on the battlefield (Sinbad, 2017). A key aspect of the 2016 Red Flag exercise was the incorporation, for the first time, the F-35 Lightning II Aircraft. F-35A variant aircraft were on full display and its capability of sensor fusion was put to the test (Sinbad, 2017). Australia and Great Britain aided in pitting our most advanced systems in a multi-domain exercise against a simulated “Red Team,” outputting air, land, sea, space and cyberspace lessons learned into the greater CONOPs (Bultman, 2017). This process of testing current and potential future equipment is a great input to the JCIDS.

Elements of the Air Force Red Flag Exercise are similar to the Marines MIX16. An article published by Joint Base San Antonio is used as one element as it highlights the space, cyberspace, and electronic warfare elements used at the exercise. Another element is taken in the form of a video out brief, featuring the key leaders for various exercise participant groups. This video uses a question and answer format. Several answers by respondents', reference sensor fusion from the F-35 platform, so inclusion of that technology seems important. To understand sensor fusion, a white paper by builder Lockheed Martin outlining the nuances and capabilities of the sensor suite is entered as an element (Lockheed Martin, n.d.). The analysis for Red Flag 2016 can be seen in Table 12.

Table 12. Analysis of Red Flag Exercise 2016

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique	
1	Red Flag Debrief (Video)	4	6	3	2	1	1	
2	Bultman Article	6	4	5	2	2	1	
3	Lockheed Martin F-35 White Paper	7	8	7	3	4	4	
Average		5.7	6.0	5.0	2.3	2.3	2.0	4.7

The Red Flag Exercise 2016 has a component score of 4.7, indicating marginally addressing AI. Element 1, the debrief video, highlights generically many of the capabilities of the F-35 and its involvement in the exercise. The exercise personnel reference information sharing rather than autonomous information sharing which limits the scoring in the AI spectrum and AI definition categories. One participant references the Link 16 sharing ability of the F-35, includes partially autonomous and algorithmic behaviors inherent to the aircraft. Element 3 supports the assumption made by element 1 highlighting how sensor fusion works (Lockheed Martin, n.d.). It denotes sensors specifically, autonomous manipulation of the data sensed, and transmittal of the information across the battlespace without pilot influence (Lockheed Martin, n.d.). This requires an algorithm, but the white paper doesn't address specific types or the use of a knowledge base, which deflates the algorithm score (Lockheed Martin, n.d.). AI classes address the perspective and type of AI system and none of the elements address classes well. Much of what is present in the Red Flag artifacts include essentials of AI, but do not reference them in the terms that are specified in this research.

d. Tech Warrior

AFRL Tech Warrior is an annual event that attempts to take the technicians from the lab and put them into combat situations, but also lets them bring their toys with them (Tech Warrior, 2017). In Fairborn, Ohio AFRL scientists and engineers attend an 11-day emersion to experience field, mobility, and combat training, and often include a particular focus like emergency medical readiness or rescue operations (Tech Warrior, 2017). Some of the technologies included in this years’ event were novel to battlefield situational awareness systems, human performance monitoring and sensors, and augmented battlefield solutions (Tech Warrior, 2017). The results of a test like this are captured and worked into the CONOPs that influences and informs the JCIDS.

The first element is an article published by the 88th Air Base Wing that describes the event in 2017 and some of the technologies incorporated (Tech Warrior, 2017). Element two is based on a video depicting the events at Tech Warrior 2016 (Air Force Research Laboratory, 2016). These two elements compose a quality assessment of the AI essentials that can be transitioned into requirements generation. The analysis for Tech Warrior can be found in Table 13.

Table 13. Analysis of Tech Warrior

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique	
1	Tech Warrior 17 Article	4	1	4	1	1	1	
2	Tech Warrior 16 Video	6	3	8	4	3	4	
Average		5.0	2.0	6.0	2.5	2.0	2.5	3.8

The composite score for AFRL Tech Warrior is a 3.8, indicating the exercise in under addressing AI. It does make perfect sense that an article the generalizes the event and mostly informs what took place would be lacking in AI. The video however, was a great forum to display and share the specific technologies that were used to support field operations. The scores for AI spectrum and AI components describe an event that is using a lot of AI systems to support their effort. Still they were not described well and once again it is the missing AI vocabulary that stops the scores from improving rapidly. Neither the article or video, scored well in the AI algorithm category, due to the lack of

identifying a knowledge base. A bright spot is found in element 2, in the AI component category. The video depicts systems with sensors, processors, and effectors which are veterans of several Tech Warrior campaigns. This indicates the AFRL already understands the necessity of practice for AI systems.

e. ThunderDrone

Sofwerx just finished ThunderDrone, an open competition held to demonstrate small drones and swarming applications (Carey, 2017). The competition includes a series of events starting with design submission and culminated with a Rodeo. The Rodeo was conducted 3 November, 2017, and the results are still being processed (Carey, 2017). While the results are not yet published and no lessons learned have been derived, the competition framework and design are widely available. The competition was run by Sofwerx, a non-profit organization initiated by the Doolittle Institute, and USSOCOM, it is open to private companies, military services, and research organizations (Carey, 2017). ThunderDrone took place in 7,000 square foot indoor test range and included four classes for competition: 3 drones, 33 drones, 333 drones, and 3003 drones (Sofwerx, 2017). Many more events of this type can be expected as the U.S. Air Force is hosting a drone racing championship series to test the arena of first person view (FPV) drone technology (Carey, 2017).

Elements for analysis are from the press coverage of the event and the Sofwerx website. The website can be broken into two elements, the information brochure (Sofwerx, 2017) and the web page dedicated to keeping interested parties informed (Sofwerx, n.d.). To be sure the webpage has a consistent content it was extracted from the Internet on 20 October 2017. These elements work to paint the picture for the use of AI essentials in the ThunderDrone competition. Analysis for ThunderDrone can be seen in Table 14

Table 14. Analysis of ThunderDrone

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique	
1	ThunderDrone News Article	5	6	5	4	1	4	
2	ThunderDrone Brochure	7	8	7	5	6	5	
3	Sofwerx Website: Event Page	6	6	8	3	1	1	
Average		6.0	6.7	6.7	4.0	2.7	3.3	5.7

ThunderDrone is an excellent event for addressing quality AI essentials. In many of the categories, it alludes to all of the characteristics for a several AI essentials. The only downside to the ThunderDrone event is that it isn't expressly referencing AI essentials by standardized nomenclature. While this may sound like a nitpicky, to receive a score of well addressing AI or better, perfect understanding of AI and standardization of the field must be considered. The biggest contributor to quality from Sofwerx is the event brochure, specifically the section that outlines the different challenges that drones will be subject to. Over the course of 13 challenges drones and swarms will be subjected to autonomous mapping, GPS denied environments, cloud sensing conditions, self-healing demands, and different scenarios requiring sensing and effecting (Sofwerx, 2017). The brochure also outlines the events as a means to develop, test, asses and share results indicating practice and training of the machines will be a part of the challenge. For these reasons, ThunderDrone has a composite score of 5.7. ThunderDrone is addressing AI, if slightly mislabeling its effort.

f. DARPA Cyber Grand Challenge

DARPA has launched an annual competition called the Cyber Grand Challenge (CGC) to help define the need for automated, scalable, machine-speed vulnerability detection and patching (Fraze, n.d.). This is the world first all-machine cyber hacking and defense tournament and took place on 4 August 2016 in Las Vegas, Nevada (Fraze, n.d.). The game played by the machines was Capture-the-Flag, a classic hacking event that as of 2013, had zero machines on the planet capable of navigating the rules (Fraze, n.d.) (Defense Advanced Research and Projects Agency [DARPA], 2016). The top three scoring teams received \$2 million, \$1 million, and \$750 thousand dollars as a prize (Fraze, n.d.). An interesting twist to this competition is that it unfurled live, using a two-and-a-half-hour studio recorded show style (DARPA, 2016).

The two elements used to analyze DARPA's CGC are press coverage on the DARPA website and the aired show. All of these elements combine to demonstrate the level of quality AI essentials received in the planning execution, and lessons learned for the demonstration. This analysis can be seen in Table 15.

Table 15. Analysis of DARPA’s Cyber Grand Challenge

Element	Description	Spectrum	Definition	Components	Classes	Algorithm	Technique	
1	DARPA Website Coverage	7	8	5	2	4	6	
2	Cyber Grand Challenge Show	7	8	6	4	5	7	
Average		7.0	8.0	5.5	3.0	4.5	6.5	6.2

DARPA’s CGC has a composite score of 6.2, indicating a grade of AI addressed. Much of this score is derived from the nature and description of the event. The machines go through the challenge live on their own network and are required to engage the game without human assistance (DARPA, 2016). This highlights fully autonomous behavior, encroaching on autonomic behavior without specifically saying it. Further, the systems must engage a knowledge base and utilize many AI techniques to score well. The coverage of the game actually reveals many specific AI techniques that are not covered by this research like fuzzing and symbolic emulation (DARPA, 2016). CGC does not score higher in AI techniques however, because of its specificity. For the information to be perfectly used in JCIDS it must be understood well throughout the JCIDS by business level administrators and engineers alike. One area where CGC does not score well is AI classes. This is partially because the environment only allows for software agents and partially because it displays software agents, but never directly explains this domain. All in all, CGC addresses AI with quality throughout, and could be expected to transfer many quality AI lessons learned into the JCIDS process.

D. CHAPTER SUMMARY

These artifacts and many others combine to build the CONOPs, impacting the information transitioned through the JCIDS and on to the DAS. The CONOPs quality directly impacts the quality of requirements that move on to define defense systems and many other aspects of DAS like design considerations, life-cycle management principles, and the means to verify and validate the resulting systems. To gain a general sense of how well the CONOPs is addressing AI essentials a final calculation is necessary. By averaging the composite scores, an understanding of the average level of AI essentials consideration across the DOD is understood. The resulting score is in Table 16.

Table 16. CONOPs Composite Score

CONOPs Inputs	Score
Joint Unmanned Systems Integration Roadmap	3.7
The U.S. Army Robotic and Autonomous Systems Strategy	4.7
Air Force Technology Horizons	2.5
A Cooperative Strategy for 21st Century Seapower	1.6
2014 DIA Innovation Strategic Plan	1.9
Marine Corps Operating Concept	2.2
UQ 16 future Force Design II Final Report	1.4
Marine MAGTF Integrated Experiment 2016	2.3
Air Force Red Flag 2016	4.7
Tech Warrior	3.8
ThunderDrone	5.7
DARPA Cyber Grand Challenge	6.2
CONOPs Composite Score	3.4

The result is a score of 3.4, indicating the CONOPs is under addressing AI. Out of the 12 artifacts analyzed, four are considering AI essentials to a level that can inform JCIDS and improve the output of AI requirements. Even the strategic documents, exercises, and demonstrations that do address AI, are generally not using the correct vocabulary or terminology. Fortunately, the information in this research can work to inform the strategic documents, exercises, and demonstrations and the analysis portion can again grade how well AI essentials are incorporated. At present, some adjustment to better prepare JCIDS inputs with AI essentials, is necessary for quality AI requirements.

Inputs to the JCIDS process are insufficiently representing the field of science and technology of AI. Logically the outputs from the JCIDS would also insufficiently address AI essentials. This means, the ICDs headed for DAS will not include the AI essentials that represent quality AI systems. This does not mean that systems will not be fielded and that all of the work analyzed is without merits. Still, if new systems are developed and the tenets of AI are missing, this could cause significant problems in the processes that follow, specifically inside of DAS.

IV. DAS

A. INTRODUCTION

Once the essentials of AI traverse the JCIDS process, and a validated need is defined, they are then subjected to the Defense Acquisition System (DAS). The DAS is the formal process for creating material solutions that match the requirements defined in the JCIDS (Department of Defense [DOD], 2007). This section of research starts with a basic primer of the actions inside of the DAS. Next, this chapter outlines where in the DAS failures are likely to occur, based on the assumption evolving from Chapter III: Low quality JCIDS inputs will result in poor validated requirements. The areas of DAS analyzed are systems requirements generation, technology transition, test and evaluation, and life cycle management. In all cases, this research seeks to determine if cost, schedule, or performance will suffer due to insufficient AI essentials entering the JCIDS.

B. DAS PRIMER

In order to understand the landscape AI systems will work through, a basic understanding of the Defense Acquisition System is needed. The DAS is the management process that provides effective, affordable, and timely systems for DoD users (Department of Defense, 2007, p. 4). It is this defined process that allows for acquisition programs to exist, receive appropriated funding, and progress in a responsible and accountable manor (DOD, 2007, p. 4). The Department of Defense Directive 5000.01, which outlines policy and applicable instructions, is used to enact the process (DOD, 2007, p. 1). DoDD 5000.01 lays out the goals and important process interdependencies that DAS is aiming for and subject to. Fundamental goals for the DAS are to enhance flexibility, deliver responsiveness, support innovation, secure business discipline, and streamline an effective management process (DOD, 2007, p. 3). Key interdependencies outlined are the federal financial systems of Defense Finance and Accounting System and PPBE, as well as, research and development, test agencies and policy, information security, JCIDS, legal, and logistics efforts (DOD, 2007, p. 6-7).

The document that best explains the acquisition processes steps is the Department of Defense Instruction 5000.02, which defines the methods and authorities that defense acquisition personnel have to enact the acquisition process (DOD, 2017, p. 1-2). This document starts by outlining the specific roles in the DAS for various defense acquisition personnel and the authorities and statues that apply to each position (DOD, 2017, p. 2). It further refines the thoughts in DoDD 5000.01, pointing out the different types of programs, program structures, and the ways in which supporting processes engage the DAS (DOD, 2017, p. 3-5). One of the most important aspects of the DAS is the step by step process by which a program is run and where foundational interactions occur. The basic outline of the DAS is displayed in Figure 14, as well as, how the outputs of the JCIDS engage it.

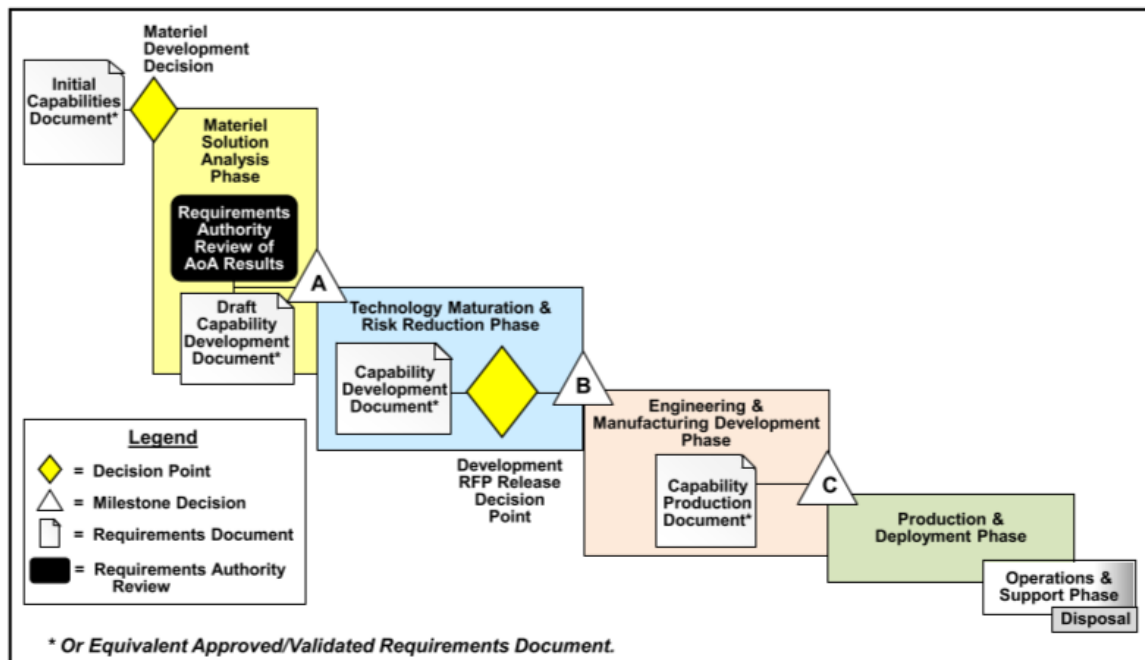


Figure 14. Illustration of a Basic DAS Process, Complete with JCIDS Connections. Source: DOD (2017).

Important aspects of this process follow a systems engineering waterfall modeling technique. Inputs from the JCIDS are white document symbols (DOD, 2017, p. 6). These inputs are the instruments that take the work done in the JCIDS and apply them to the

various stages of system development (DOD, 2017, p. 6). In yellow diamonds are key decision nodes, where responsible DAS personnel determine the progress to date and future steps necessary for developing systems (DOD, 2017, p. 6). The five key phases of the DAS are also displayed. These are the material solutions analysis (MSA) phase, technology maturation and risk reduction (TMRR) phase, engineering and manufacturing development (EMD) phase, production and deployment (PD) phase, and operations and support (O&S) phase (DOD, 2017, p. 6).

The MSA phase is used to conduct the analysis and other activities necessary to choose solutions to capability gaps and refine requirements into system-specific aids (DOD, 2017, p. 18). Often, this effort is concurrent with JCIDS efforts and includes research into applied sciences, development of CONOPs, and defining funding lines necessary to analyze potential material solutions (DOD, 2017, p. 19). The MSA phase is initiated by an ICD output from JCIDS as well as a material development decision (MDD) (DOD, 2017, p. 18). A key effort that must be accomplished during the MSA phase is the construction of an acquisition strategy (DOD, 2017, p. 19). This strategy must contain at a minimum justification for the preferred materiel solution, affordability and feasibility analysis, the scope of the effort, an understanding of technical risks, cost risks, schedule risks, a plan for managing intellectual property, and threats to the program (DOD, 2017, p.20). Exiting the MSA phase requires a successful Milestone A review decision, which grades the program management team on its acquisition strategy and the maturity of technology that will be included (DOD, 2017, p. 20).

The TMRR phase is used to reduce technology, engineering, integration, and lifecycle risks to the point where a decision concerning manufacturing development can be made (DOD, 2017, p. 21). The crux of this phase is to make preliminary design trades and technological improvements until an affordable and reliably scheduled program is outlined (DOD, 2017, p. 21). Throughout this phase technology levels are continuously tracked to determine if each technology supporting the design is achievable. Another key aspect of this phase comes from the JCIDS, in the form of the official requirements document the CDD (DOD, 2017, p. 22). The CDD focuses the technology maturation effort using a refined and well-articulated user need, including materiel amounts needed,

operating conditions, mission profiles, and other CONOPs considerations that speak to the nature of use for the materiel solution (DOD, 2017, p. 24). While systems engineering and developmental test efforts are underway, lifecycle management principles and basic production knowledge is accumulating (DOD, 2017, p. 22). A Milestone B review assesses the technology readiness, updated acquisition strategy, and other required engineering documents, to authorize movement from TMRR into an official program of record.

The EMD phase is used to develop, build, and test products designed earlier in the DAS (DOD, 2017, p. 27). At this point detailed engineering designs are formalized, open risks are attacked, prototypes are built and tested, and a product baseline is paved (DOD, 2017, p. 27). Additionally, manufacturing support is designed, fabrication methods are tested, and early test articles are produced (DOD, 2017, p. 28). EMD is considered complete when the design is stable, the system meets the requirements based on developmental and early operational testing, manufacturing processes are under control, software lifecycle considerations are under control, and industrial production capabilities are available (DOD, 2017, p. 28). A Milestone C review analyzes those developments and others to determine if the system is ready to proceed to production and operational testing.

The PD phase is used to produce and deliver requirements-compliant products to operational units (DOD, 2017, p. 30). It starts with limited production and operational test and evaluation, and leads to warfighters achieving initial operating capability (DOD, 2017, p. 30-31). It is very important that critical deficiencies identified during OT&E are resolved before the program progresses past low-rate initial production or limited deployment (DOD, 2017, p. 31). Once all major activities of the PD phase are complete the decision authority will determine if the system can continue to full-rate production and full deployment (DOD, 2017, p. 32). From this point on, the system is considered a part of the DOD's arsenal, ready for operational use, and requires sustainment until it is retired.

The O&S phase is used to execute the support strategy, maintain or improve system readiness, and even dispose of the system legally at the end of its lifecycle (DOD,

2017, p. 32). Sustainment is a term the DOD uses to describe all efforts to keep the system online, improve its performance, and decrease operating costs by leveraging systems engineering, contractor support, and equipment management (DOD, 2017, p. 32).

These five processes take a considerable amount of time, resources, and individual efforts to take an idea from basic science, to system development, to the end of its useful life. It performs in a very linear fashion, but many efforts dedicated to one specific system can be iterative and concurrent. Figure 15 displays a notional software intensive system that expects steady improvement in performance to be delivered throughout the systems lifecycle (DOD, 2017, p. 13).

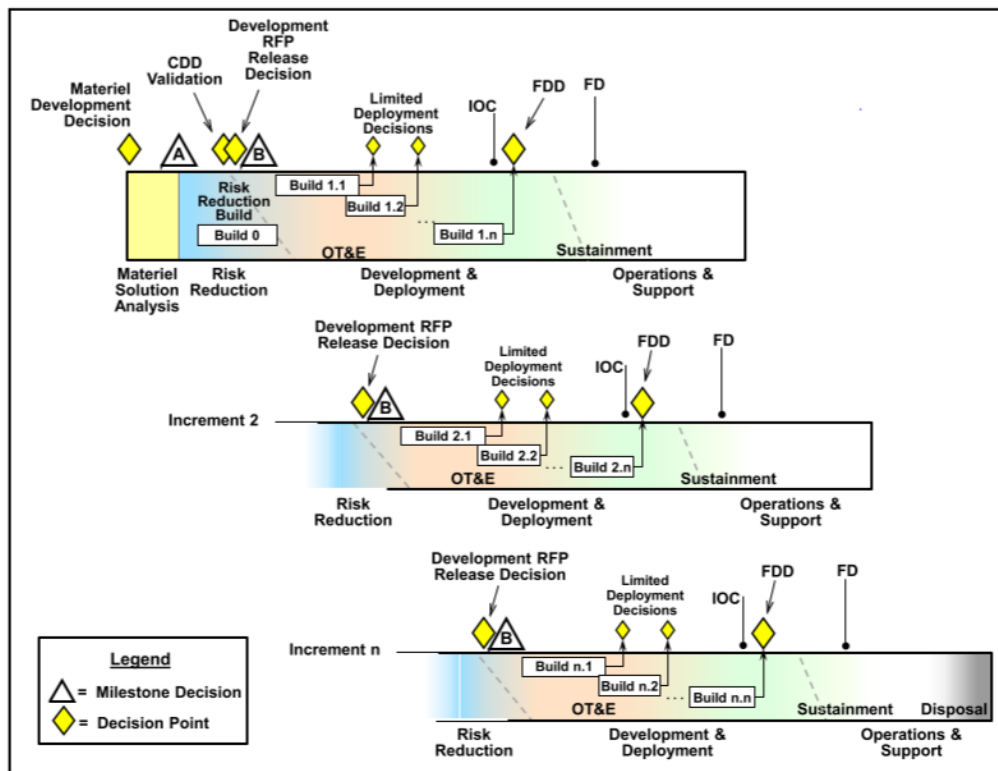


Figure 15. Incrementally Deployed Software Intensive Program.
Source: DOD (2017).

From Figure 15, it is clear that there are many independent efforts to manage. Each of the individual efforts is defined as an increment, and any given increment can be in a different DAS phase. The steps are still the same, but management of each baseline starts with addressing the users need and timeframe for that need, basic technologies available to support, maturation of the technology, integration and development, production, and deployment. Test events must be scheduled throughout to reduce risk, verify and validate performance, and prove to users the system will conform to their mission needs. At many points cross functional teams will be necessary to define the scope of the program, grade the business ability of the acquisition strategy, anticipate support requirements years in the future, and establish supply chains.

The core of DAS is a disciplined systems engineering effort that has been modified to incorporate the diverse processes of defense acquisition (DOD, 2017, p. 87). It incorporates warfighters needs, acquisition monitoring and management, and an industry base that ultimately grows technology and performs the gross amounts of engineering work (DOD, 2017, p. 87). Success is defined as delivering to the warfighter a system that meets its requirements, when it is asked for and at the program price point specified in the program baseline (DOD, 2007, p. 3). Overruns in cost and schedule are disruptive and can cause a program to be cancelled. Additionally, systems that fail to meet performance objectives risk perishing or leave warfighters vulnerable on the battle field. The outcomes of DAS efforts have a direct impact on readiness and capability for warfighters, which is the driving force for determining the expected problems AI technologies will encounter.

Systems that incorporate AI technology are destined to travel through the DAS and its various management disciplines. A validated requirement will activate the MSA phase, where potential solutions will be explored (DOD, 2017, p. 18). AI has many uses and will undoubtedly surface as a solution to warfighter needs. AI technology will be matured and the application will be enhanced during TMRR (DOD, 2017, p. 21). A stable design must be made to work through EMD, to the point that production processes prove the system can actually be built (DOD, 2017, p. 27). During the PD phase, AI production units will be delivered and tested to see if they do in fact match the validated requirement

that started the fielding process (DOD, 2017, p. 30). Finally, procurement, repair, improvement, and eventually retirement will await these systems (DOD, 2017, p. 32). For AI technology to last throughout these processes, it must allow for program schedules to remain intact, system performance to meet warfighter expectations, and not outpace program funding.

C. DAS FAILURES

The DAS is a process that has both proven to be reliable and yet leaves many involved with it wanting more. True, it can be tailored to fit almost any type of acquisition, from commercial purchases to transition of stealth technologies into deep strike combat capabilities (DOD, 2017, p. 9). But as notable as its successes can be, its failures are often awesome. This research is not suggesting that the DAS is inadequate or must be changed, but rather that a confluence of events can cause the demise of very promising systems. The current standard for technologies that have struggled throughout the DAS, are software intensive systems. Management guides of every type are available to support software intensive systems as they work through the DAS.

Guidelines for Successful Acquisition and Management of Software-Intensive Systems is a guide dedicated to fielding software intensive systems (United States Air Force, 2000, p. 1). It includes the role for software in today's forces, a vision of software's future, the effects of software in reducing budgetary demand, education on how software intensive systems behave and are built, traditional problems with software acquisition, a navigation guide to statutory requirements, and specifically outlines every step a program team needs to take a program from initial development to deployment (United States Air Force, 2000, pp. 1-945). Despite this guide, sufficiently outlined DAS processes, and talented people, the DoD still has failures in software intensive systems. A 2004 General Accounting Office (GAO) report titled *Defense acquisitions: stronger management practices are needed to improve DODs software-intensive weapon acquisitions*, explains the mixed results DOD software intensive systems have had. They state, "The F/A-18 C/D, a fighter and attack aircraft, and the Tactical Tomahawk missile had fewer additional cost and schedule delays. For these programs, developers used an

evolutionary approach, disciplined processes, and meaningful metrics” (Schinasi, 2004, p. 2). This shows that managing requirements and systems engineering well can usher through programs that meet performance, cost, and schedule requirements.

The GAO also exposes the other side, the programs that couldn’t perform well inside the DAS. They write, “the following programs, which did not follow these management strategies, experienced schedule delays and cost growth: F/A-22, an air dominance aircraft; Space-Based Infrared System, a missile-detection satellite system; and Comanche, a multimission helicopter” (Schinasi, 2004, p. 2). AI based systems, if not properly managed can cause the type of consequences that end up as low-lights in GAO reports, or worse, cause the cancelation of systems that would greatly benefit warfighters.

The GAO outlines technology development control, requirements growth, and lack of sufficient metrics to track progress as the major failures that caused cost and schedule growth (Schinasi, 2004, p. 3). Since software-intensive systems live on the AI spectrum near the autonomous system level, these failures serve as areas worth investigating using an AI lens. The first research effort will be to define the impact requirements have on a program and what struggles AI may have in defining system requirements. Since, system requirements act to inform many of the processes inside of the DAS, it behaves as the first domino in a chain of failures. The next domino outlined by the GAO, is technology transition and this research tries to identify problems that can arise in technology transition due to poor system requirements. Next, test and evaluation relies on the metrics identified by the GAO, and this research seeks out the consequences of faulty requirements. Finally, since life cycle management is used to control system progress after fielding, it is analyzed to determine what significance poor requirements will have for AI.

1. System Requirements Generation

A requirement is a shared understanding between the warfighting community and acquisition community, concerning the minimum performance of a materiel solution (Department of Defense [DOD], 2012, p. A-1). These requirements that are an output of

the JCIDS process, in the form of key performance parameters (KPPs) and key system attributes (KSAs), work as the thread between business functions and expected mission applications (DOD, 2015, p. 18). They enter the DAS after a MDD, triggering R&D processes, and again after a successful Milestone A decision, via a CDD (DOD, 2017, p. 24). This triggers funding profiles and directs basic science efforts leading to the technology transition arena analyzed in the next section (DOD, 2017, p. 24). After Milestone A, the validated requirement triggers a new round of funding, the building of an acquisition strategy, and a request for proposal (RFP) release decision in order to prepare for Milestone B (DOD, 2017, pp. 19-25). While KPPs and KSAs are what normally come to mind when the discussion of requirements comes up in acquisition, they are not the focus of this section.

KPPs and KSAs are known as stakeholder requirements, as defined in the *INCOSE Systems Engineering Handbook* (Haskins et al., 2011, p. 56). Stakeholder requirements do not do the job of completely representing a system, in that they are void of considerable analysis that leads to system requirements (Haskins et al., 2011, p. 72). System requirements are the set of requirements needed to meet project and design constraints (Haskins et al., 2011, p. 74). The process of transforming KPPs, KSAs, and other supporting data into system requirements is called requirements analysis. INCOSE outlines the process of requirements analysis as having five major components (Haskins et al., 2011, p. 73). These are inputs, controls, activities, enablers, and outputs. Figure 16 displays how these five factors work together to produce a complete set of system requirements.

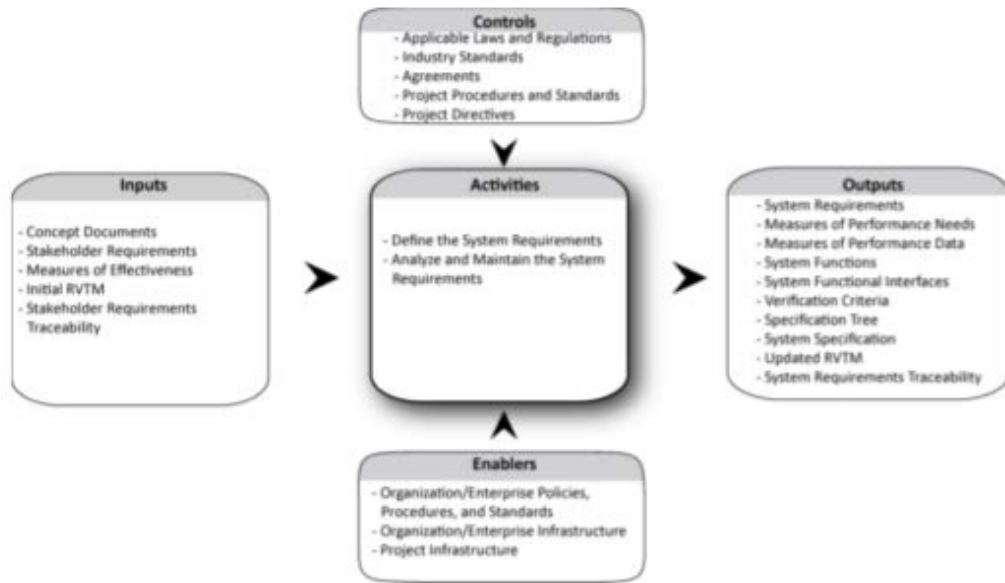


Figure 16. System Requirements Generation Process.
Source: Haskins et al. (2011).

The inputs component has terminology this research has discussed previously. Notably the CONOPs and stakeholder requirements (Haskins et al., 2011, p. 73). It also has a few terms that will not be fully satisfied by the CDD or direct engagement with warfighters. Measures of effectiveness may be based on a KPP, but more often than not, are derivatives of KPPs and KSAs that must be developed by a program team through integrated project teams (IPTs) (Haskins et al., 2011, p. 73). Organizational dynamics and other enablers constrain and confine the solution space, often pitting KPPs and KSAs against each other necessitating performance trades (Haskins et al., 2011, p. 73). Program procedural requirements can extend the warfighter’s timeframe, and other program controls work to shape the realm of what is and isn’t possible. At the end of this, a new set of system requirements emerges that enable PMs to understand and design a system that will meet warfighter needs and usher the system through the DAS.

System requirements, not JCIDS requirements, are used to define the necessary DAS management processes. They define the standards, system boundaries, and interfaces which underpin system design (Haskins et al., 2011, p. 75). They lead to performance measures which are the crux of a testable system and they facilitate efficient

and cost effective lifecycle dynamics (Haskins et al., 2011, pp. 73-76). It may seem like the program is distancing itself from KPPs and KSAs, but a key characteristic for good requirements is that they are traceable; every system requirement is accurately supporting the warfighter need (Haskins et al, 2011, p. 79).

A 2015 GAO report titled *Military Service Chiefs’ Concerns Reflect Need to Better Define Requirements Before Programs Start*, highlights the ballooning nature of program management costs based on this process of requirements analysis (Sullivan, 2015, p. 1). The report cites “creep – or growth – in the high-level requirements [KPPs and KSAs] is rare” (Sullivan, 2015). This indicates that requirements analysis must do a considerable amount of work to actually develop a system the warfighter can use. Figure 17 displays how requirements typically grow over the life of a system, expressed chronologically with a generic DAS graphic (Sullivan, 2015).

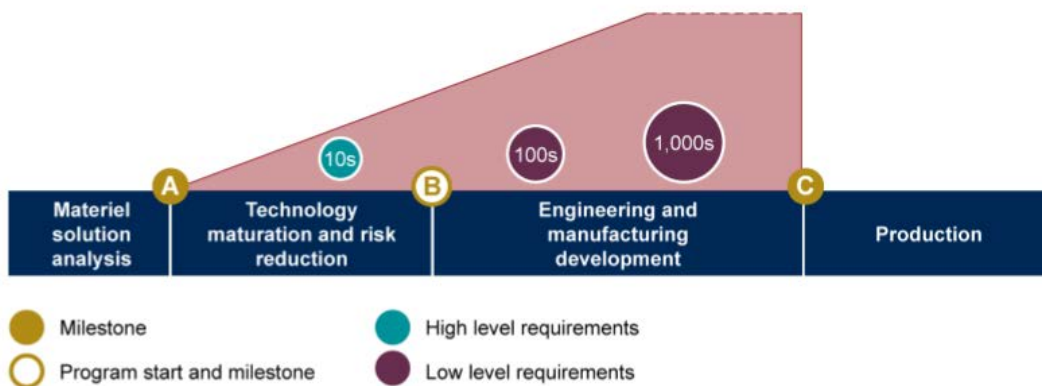


Figure 17. GAO Chart Displaying Requirements Growth. Source: Sullivan (2015).

The GAO also identified that, “cost growth and other problems are more directly related to developing lower-level requirements after the program has started (Sullivan, 2015, p. 11). The cost growth is well explained by considering program trades. GAO writes, “requirements are insufficiently defined at program start; when their full consequences are realized, trade-offs are harder to make—cost increases and schedule delays become the preferred solutions” (Sullivan, 2015, p. 12). Given the low level of AI

essentials quality influencing the JCIDS, it is unlikely that the KPPs and KSAs will be sufficient once the programs are initiated. This will cause program teams to build a considerable amount of the requirements needed to field an AI system, after program initiation. The follow on processes that depend on system requirements for proper management will start just as slowly. The trade space that these management processes depend on will evaporate, and the consequences indicated by the GAO will occur.

2. Technology Transition

Technology transition is the process of taking emerging technologies, packaging them into relevant systems, and finding a mainstream market for it (Moore, 1991, p. 14). *Crossing the Chasm*, by Geoffrey A. Moore, is a book dedicated to examining the growth of technology through a business space, specifically focusing on how a technology is expected to be adopted throughout its lifecycle. It starts by explaining the High-Tech Marketing Model, as seen in Figure 18., which explains the complete lifecycle of any notional technology (Moore, 1991, p. 17). The lifecycle has five major stages of technology adoption from innovators, early adopters, early majority, late majority, and laggards which are labelled as such for the types and number of people that congregate in each stage (Moore, 1991, p. 21). Moore (1991) also explains that there are cracks between each region, where products that will not get accepted by the next stage can leak out of the market and become confined to an existence of limited use (p. 22). The crack that holds the most significance and caused the title of the book, is, “the deep dividing chasm that separates the early adopters from the early majority” (Moore, 1991, p. 25). It is crossing this chasm that is most problematic for commercial and defense systems. In the DAS, this chasm would equate to milestone A. It is at milestone A reviews, that acquisition leadership must be convinced a concept that satisfied the ICD requirements is viable for a defense system (DOD, 2017, p. 20). This is the point here emerging technologies are either included and head toward early majority acceptance, or fall through the chasm and become relegated to zombie status (Moore, 1991, p. 30)

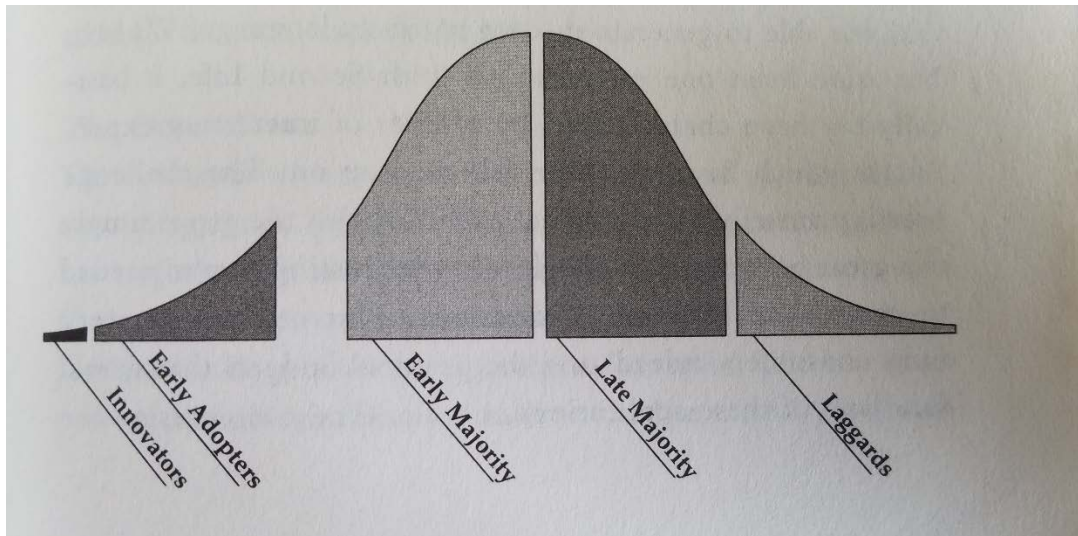


Figure 18. High-Tech Marketing Model. Source: Moore (2014).

In typical military fashion, guides to support technology transition have been developed to support acquisition personnel and PMs. The two that provide the best insight for this research are the *Defense Acquisition University: Manager's Guide to Technology Transition in an Evolutionary Acquisition Environment* and the *Air Force Technology Development and Strategy Guidebook (TDTS)*.

The Manager's Guide to Technology Transition in an Evolutionary Acquisition Environment is designed to help defense acquisition personnel deliver, “the latest technology into the hands of the warfighter in the quickest, most cost effective manner possible” (Defense Acquisition University [DAU], 2005, p. iii). DAU initiated this effort with the Under Secretary of Defense for Advanced Systems and Concepts (DAU, 2005, p. iii). The stated use for this document is to support program efforts in, “taking full advantage of science and technology,” by promoting effective collaboration and informing program teams of the appropriate resources at the right time (DAU, 2005, p. xv). Important to this research, are the challenges the guide lists in the arena of technology transition. It highlights that technology is changing rapidly, commercial items with considerable technology are increasing, and our adversaries have access to our defense technology (DAU, 2005, p. xvi). To achieve successful technology transition, the guide offers mechanisms for defining the environment of technology transition,

technology transition planning and tools, programs that facilitate technology transition, and challenges that face a typical technology transition effort (DAU, 2005, p. xviii).

The environment for technology transition, echoes Moore's definition of technology transition. It defines it as, "the use of technology in military systems to create effective weapons and support systems – in the quantity and quality needed by the warfighter to carry out assigned mission" (DAU, 2005, p. 1-1). This definition highlights the first problem that AI will encounter inside of technology transition. If the warfighter has ambiguity in the mission they intend to pursue, given the complexity of AI, then the goals for technology transition cannot be achieved. The ability to transition even commercially mature technologies, with their reduced requirements for contracting and testing, can fall subject to inaccuracy. If the warfighter cannot state the need correctly, then the resulting technology will be a poor match to missions.

The Manager's Guide to Technology Transition in an Evolutionary Acquisition Environment also outlines the planning efforts necessary for successful technology transition. It defines the first steps as originating in the science and technology community (DAU, 2005, p. 2-1). This is where the management of a technology begins, from its inception in the research and development laboratories. From a timeline approach this is not, however where the technology transition begins. While the military science and technology community efforts are sometimes conducted independent of capability, the majority of efforts are initiated based on expected future warfighter capability needs (DOD, 2017, p. 171). The process of identifying a technology and supporting it with RDT&E funding requires funding requests 18-24 months in advance (DAU, 2005, p. 2-1). This puts the activating burden for technology transition on the needs authorized by the JROC. Since, inputs to JCIDS are under addressing AI essentials, future warfighter needs are not likely to properly inform the S&T community. This will cause two possible outcomes. R&D efforts will be delayed however long it takes to build effective warfighter needs plus 18–24 months. This means that all subsequent processes in the DAS will be delayed, displacing the timeframe that a system can actually be delivered back to warfighters. The second possibility is that technology development will be a poor match for warfighter needs (DAU, 2005, p. 4-1). This occurrence would cause

delays later in DAS, while the program team struggles to fix deficiencies in system design late in the acquisition process.

The Manager's Guide to Technology Transition in an Evolutionary Acquisition Environment also lists a series of challenges for technology transition. These challenges cover many activities, but this research is focused on the challenges found in the *inserting and enabling technology* area (DAU, 2005, p. 4-1). Challenges in this area are further partitioned into the communities that they are expected to impact (DAU, 2005, pp. 4-1 – 4-30). The challenges that have direct application to this research are from the capability needs community (DAU, 2005, p. 4-2). The challenges this research will focus on are:

- Do your capability need documents describe the essential warfighting capabilities?
- Do your capability needs documents employ an incremental approach?
- Do your capability needs documents support technologies that reduce life cycle costs?
- Are you involved in the S&T planning?

Are the capability needs documents available for supporting transition? (DAU, 2005, p. 4-1).

It is the first challenge that should give the most pause, even though all are good considerations for AI when transitioning technology. This research has shown that inputs to the JCIDS are having difficulty describing AI. Specifically, AI characteristics of required mobility, domain, branching capability, level of reasoning, technique required, and knowledge base source. If these AI essentials are not correctly incorporated to capability need documents, then accurate technology transition will falter. The tools in *The Manager's Guide to Technology Transition in an Evolutionary Acquisition Environment* that increase technology transition effectiveness, like utilizing OTAs to support prototyping, IPPD planning, and small business innovation and research programs will all suffer the inaccuracies found in capability need documents (DAU, 2005, pp. 3-2 – 3-11).

The TDTS, published by the U.S. Air Force, takes a wholly different approach to managing technology transition. They articulate a process of stages, with gates at the end

that control an end to end process of technology development (United States Air Force, 2010, p. 4). It claims that the process of moving technology from the beginning stages of DAS is generally well known, but maturing of that technology has been using an inefficient process (United States Air Force, 2010, p. 4). To achieve a disciplined process, the document emphasizes creating a TDTS strategy, with a dedicated IPT, and tracking the technology through each technology gate (United States Air Force, 2010, p. 5). To meet the exit criteria at each gate, means to ensure the readiness of the technology at each stage of the DAS. Figure 19 displays the process by which technology is expected to transition alongside the DAS.

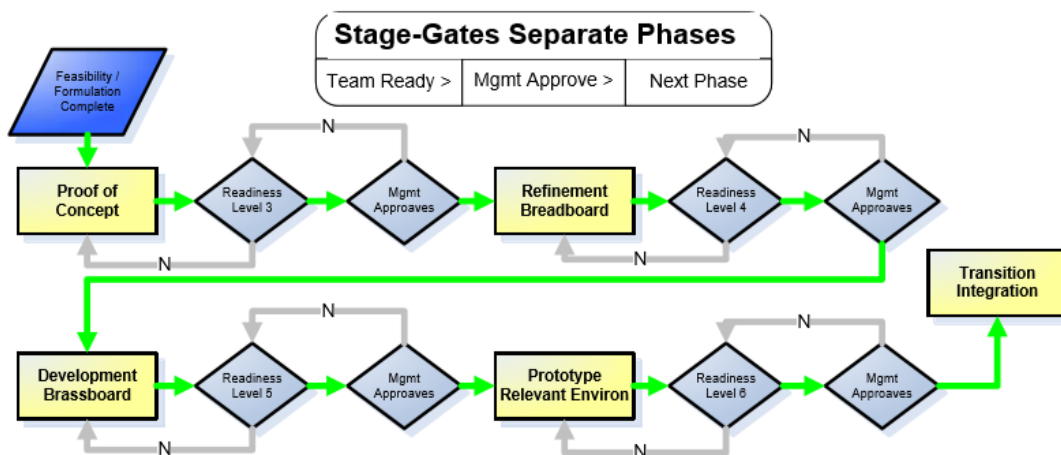


Figure 19. TDTTS Stage Gates Process. Source: Kropas-Hughes, Rutledge, & Sarmiento (2008).

The first problem that a PM and technology manager (TM) will run into at their first IPT, is the quality of the concept of exploration refinement (CER). The CER, is the process of identifying relevant issues, recommending a transition strategy, and identifying support issues and ultimately leads to the effective capture of technology maturation objectives (United States Air Force, 2010, p. 5). It establishes the technology baseline for the TDTS strategy and enables the IPT to initiate the development of the life cycle management plan (LCMP). A key driver of the CER are the outputs of the JCIDS, including any weapon systems roadmaps, user significance for having the technology,

and warfighting capability that this technology concept is addressing (United States Air Force, 2010, p. 13). A key output of the CER are the identification of exit criteria for the several future stages, expected deviations or waivers, and a well-defined capability need (United States Air Force, 2010, p. 25). Since this effort of TDTS planning using the CER occurs in the proof of concept phase, it is subject to the exit criteria of Stage-Gate 3 (United States Air Force, 2010, p. 14). The exit criteria are:

- Technology readiness level (TRL) 3 and manufacturing readiness level (MRL) per MRL plan.
- Technology concept has been proven sufficient to meet the user need in a laboratory environment, and a proof-of-concept has been documented.
- The TDTS document has been approved (see Section 2.6), and an acquisition agency has shown a level of interest.
- A Breadboard Laboratory Validation Plan for the refinement stage has been developed, with purpose, objectives, and scope adequately described. (United States Air Force, 2010, p. 14).

While three of these criteria can be accomplished without a significant input from the JCIDS or CER, determining if the proof of concept is sufficient to meet warfighter needs is explicitly reliant on the validated requirements. Missing essentials from AI in requirements will corrupt this exit criteria resulting in two possible consequences. The first, is that exit criteria will not be satisfied. Too much ambiguity would alert an adept PM that the concepts and its sufficient technology baseline is not possible given the AI inputs. This would delay the process until the warfighter can effectively articulate the branching ability or machine learning requirements for example. The second consequence, and much more likely given the can-do nature of PMs, is that the team would piece together an inadequate concept. This concept, with unforeseen risk, would continue forward and carry on the inaccuracies late into the DAS where changes and delays get costlier. All of the program control that stage-gates offer, supporting technology transition process account for, and technology transition guides organize, will erode at the first stage-gate without strong consideration of AI at the beginning of the JCIDS.

3. Test and Evaluation

Early in the acquisition lifecycle, often concurrently with requirements analysis, the PM establishes the test and evaluation working IPT in order to formulate a test and evaluation master plan (TEMP) (DAU, 2013, p. 713). The purpose of testing and evaluation is to verify and validate the performance capabilities which are documented as requirements, track and reduce technical risk, and determine if the resulting system is operationally effective (DOD, 2012, p. 23). Early in the life of a program T&E can be used to help determine feasibility of concepts and support trades in the design space (DOD, 2012, p. 23). Later, T&E answers the question of operational effectiveness and suitability of the resulting system (DOD, 2012, p. 23). All of the T&E efforts throughout the acquisition lifecycle, are predicated by the requirements delivered from JCIDS and the resulting requirements analysis process (Haskins et al., 2011, p. 82).

The requirements analysis process ultimately defines the T&E strategy, by identifying functionally every action a system must be able to perform and a complete set of non-functional requirements (Haskins et al., 2011, p. 74). From this effort the functions can be broken down into key operational issues and then measurable characteristics, which are testable and measurable whereas a requirement often is not (Haskins et al., 2011, p. 74). From these quantifiable measures of suitability, performance tests can be designed to determine the success of the design and a means for tracking improvement. (DOD, 2012, p. 77). It is only through this functional decomposition that data requirements, test metrics, test plans, and evaluation criteria can be managed (DOD, 2012, p. 77). Figure 20 displays the path from validated need to testing and evaluation requirements.

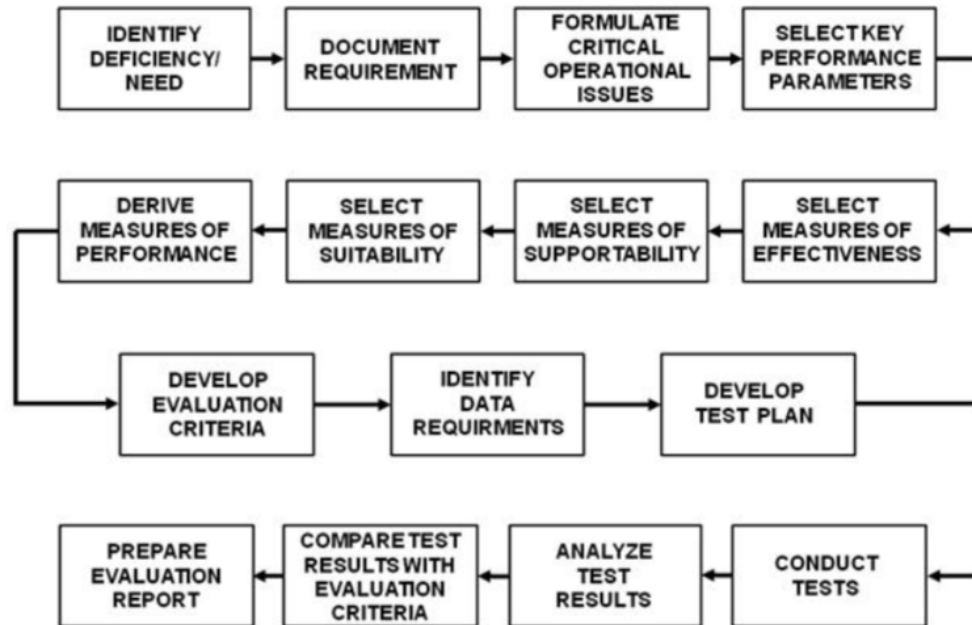


Figure 20. Flow Diagram from Requirement to Test Report.
Source: DOD (2012).

Only once all of the details pertaining to test specifics are understood, can a comprehensive TEMP be built and executed (DAU, 2013, p. 745). DAU defines the TEMP as, “the overarching document for managing a T&E program (DAU, 2013, p. 745). This means it is the document by which a PM can effectively coordinate the wide array of required actions from reducing technical risk, estimating the many testing and evaluation costs, and ultimately certifying the system (DAU, 2013, p. 746). Program success cannot be had without an accurate and measurable plan.

The GAO reported in 1993 that the DOD has been slow to improve in the testing of software intensive systems (Cooper, 1993). It cites four main reasons why the DOD has failed to improve in this effort:

- Failed to address the critical nature of software inside of systems
- Failed to develop standardized processes and tools for cost, performance, and schedule decision making
- Failed to develop test and evaluation policy that provides software maturity guidance

- Failed to define and manage requirements for complex software (Cooper, 1993, p. 15).

Admittedly, 1993 is a long time ago. To the DOD's credit, it has addressed many items, including how critical software is to system performance and developed T&E policy to support software acquisition. Other effective efforts are the insertion of software specific guidance in the 5000 series (DOD, 2017, p. 32), T&E Management Guide (DOD, 2012, p. 165), and the inclusion of software acquisition training into the career development for acquisition personnel.

Still, in 2010 the National Defense Industrial Association released a report titled, *Top Software Engineering Issues within Defense Department and Defense Industry* that found DOD efforts had not yet tackled the problems with requirements management (National Defense Industrial Association [NDIA], 2010, p. 1). The report tracked problems with software acquisition using a base year in 2006 and checking again in 2010 for improvement (NDIA, 2010, p. 1). The number one issues in 2006 was that, "the impact of requirements upon software is not consistently quantified and managed in development or sustainment" (NDIA, 2010, p. 1). The status as of 2010 reads, "some progress, but still inconsistent software requirements definition in planning and sustainment" and they added, "issues persist with JCIDS documents and resulting definition, management, and flow down of software requirements" (NDIA, 2010, p. 1). In 25 years, the DOD has failed to properly manage system requirements throughout the effort of T&E for highly technical systems. This will continue with AI, if JCIDS outputs do not expertly articulate the need for AI in future systems.

4. Life Cycle Sustainment

Life cycle sustainment includes a range of planning efforts, implementation actions, and execution steps that work to support defense systems after they are procured (DAU, 2013, p. 1). The goal of life cycle sustainment management is to maximize readiness throughout the life of a system and incorporate upgrades and modifications as requirements change (DAU, 2013, p. 1). The bulk of program execution happens well after the system is fielded, but adept life cycle sustainment can only be achieved with

considerable planning early in the DAS process (DAU, 2013, p. 1). The Defense Acquisition University’s *Defense Acquisition Guidebook* dedicates an entire chapter to the proper planning and execution of life cycle management (Defense Acquisition University, 2013). Figure 21 displays holistic view of life cycle management and where in the program specific efforts should be considered.

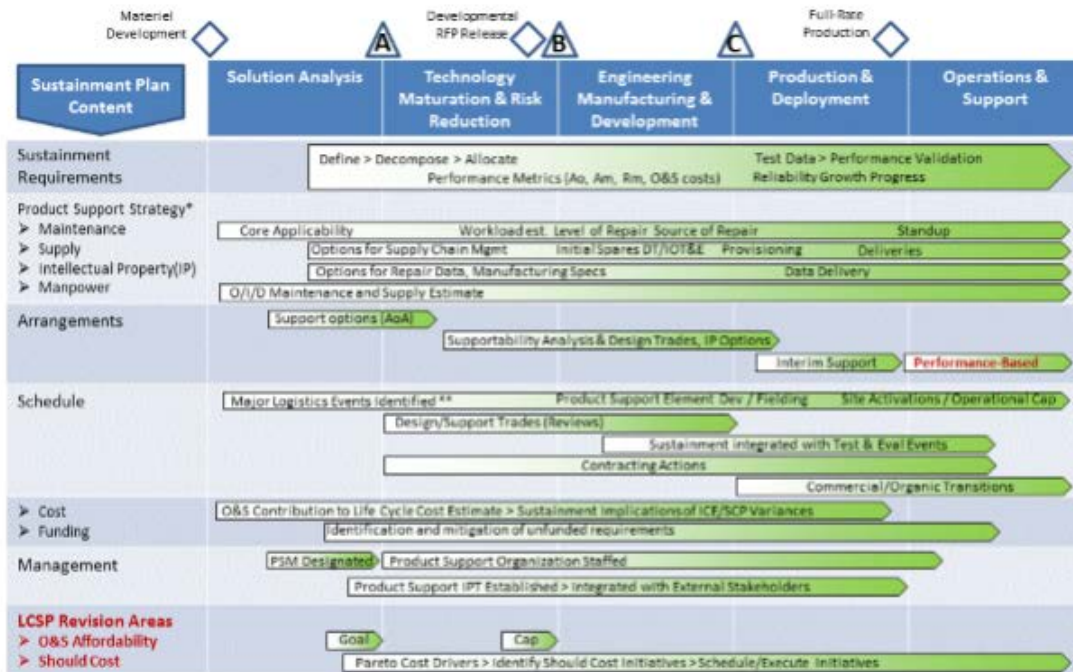


Figure 21. Life Cycle Management throughout the DAS.
Source: DAU (2013).

Figure 21 displays many actions as arrows, with the left most part of the arrow used to indicate where a lifecycle effort should begin. Not only does the sustainment plan need to be penned before the program can begin at milestone B, a considerable number of sustainment efforts begin at the earliest stages of DAS (DAU, 2013, p. 2). The *Defense Acquisition Guidebook* explains well what Figure 21 is suggesting, writing “successful post-fielding sustainment performance depends on thoughtful consideration during requirements development and solution analysis.” (DAU, 2013, p. 2) Like technology transition and T&E, life cycle sustainment planning is highly dependent on the

requirements received from ICDs and CDDs through the requirements analysis process. Like technology transition and T&E, lifecycle management will suffer the same inaccuracies when considering AI systems.

In 2014, the GAO released a report titled *F-35 SUSTAINMENT: Need for Affordable Strategy, Greater Attention to Risks, and Improved Cost Estimates*, which catalogs the mounting life cycle risk for the F-35 (Russell, 2014). The F-35 is a multinational acquisition, intended to produce a jointly used and affordable fifth generation fighter aircraft (Russell, 2014, p. 5). Not only is this an advanced aircraft, but it is also a software intensive system. Its primary sustainment tool is an autonomic logistic information system (ALIS), which predicts maintenance and supply issues, automates support processes, and is designed to optimize its own lifecycle management decisions (Russell, 2014, p. 9).

Unfortunately, the GAO reports that largely due to immaturity in sustainment ability ALIS is seven years behind schedule and the DOD's cost estimate of \$1 trillion over a 56-year lifecycle is likely inaccurate (Russell, 2014, p. 14). They cite that, "weaknesses exist with respect to a few of the assumptions" and that, "estimates did not use reasonable assumptions about part replacement rates and depot maintenance." (Russell, 2014). Currently, the F-35 is going through low rate production, but has not yet finalized a sustainment strategy (Russell, 2014, p.10). Much of this is caused by miss understanding the AI training requirements for the highly autonomous ALIS system (Russell, 2014, p. 14). GAO (Russell, 2014) explains the process for improving ALIS functionality as:

[ALIS Maintenance Team] had to use multiple approaches to identify the best maintenance solution. Once identified, the maintainer can submit this solution as an update for AFRS [ALIS' Anomaly and Failure Reporting System]. However, this update must first be reviewed by field support on site and then sent—in the form of an action request—to the contractor for approval before it is integrated into AFRS. Maintenance officials told us that they have submitted several thousand action requests to date and have thereby created a backlog, leaving maintainers to wait multiple days for an approval (p. 14).

The warfighter, resource sponsors, PEO, PM, and lifecycle management effort did not properly understand the requirement for AI practice or the techniques for machine learning. This and other F-35 requirements like it have delayed the F-35 fielding timeline and are causing uncertainty throughout lifecycle management. Complete understanding of AI essential early in the JCIDS and DAS process could have yielded a requirement that was supportive of ALIS training demands; saving funds later in the lifecycle.

D. CHAPTER SUMMARY

DAS and its associated processes present a disciplined process for taking basic sciences and converting them into usable defense systems. This starts with the warfighter's validated need and applies requirements analysis to fully understand the system requirements, technology transition demand, T&E claim, and lifecycle management stresses. Quality requirements lead to quality DAS functions and poor outputs from the JCIDS process cause significant heartburn throughout DAS and system use. Effort should be taken to ensure requirements are as mature as possible, inclusive of AI essentials, before the activation of the DAS. If not AI program cost, performance, and schedule will continue to meet the fate of software intensive systems; over budget, over schedule, and short on performance.

THIS PAGE INTENTIONALLY LEFT BLANK

V. CONCLUSION

A. INTRODUCTION

This research aimed to analyze the issues that AI will have navigating defense acquisition, by answering the primary and secondary research questions. The key areas of defense acquisition, as defined by the project scope, are inside of the JCIDS and DAS structures. To properly answer these questions, the research gathered the opinions of experts in the AI field, studied education textbooks for the essentials of AI, and conducted a basic market research surrounding the AI field of study and the technology as it exists today. Further, the processes used by the DOD to field systems were laid out using published DOD documents. The areas that deal with managing technology and bringing them into system design were highlighted to understand what inputs are important to the JCIDS and DAS processes. Then, key inputs to JCIDS were taken from across the various services and the expected quality of JCIDS output were assumed. This assumption of JCIDS outputs, specifically low poor appreciation for AI essentials in requirements, was used to predict the expected outcome of the DAS. This analysis does not serve to belittle the processes or planning efforts, but works to highlight the areas that require improvement to achieve success in AI system acquisition.

B. SECONDARY RESEARCH QUESTIONS

The answers to the secondary research questions can be found by analyzing Chapters II, III, and IV. These answers build and synchronize to answer the primary research questions. Truly, each can function independently to inform PMs and other defense acquisition personnel, but to answer the primary research question each secondary question must balance with each other.

1. What Does AI Mean for DOD Acquisition and Industry Today?

To the DOD, artificial intelligence is one of many technologies that will be incorporated into defense systems. It is both a scientific field of study that requires investment and cultivation, and it is the process of harnessing human functions inside of

machines. This research outlines a workable definition that can be used by DOD personnel to help them better understand the field and technology.

Artificial Intelligence for Defense Systems: *a multidisciplinary scientific field that aims to study human intelligence, and the attempt to bring rational and human behaviors into a system that can think, act, and self-manage.*

Key considerations for these AI systems are mobility, composition, and types of algorithms which are enlightened by the analysis of AI components, AI classes, and AI techniques. A complete spectrum from automatic to autonomic must be understood in order to properly plan for system independence and functionality. Not only is it important to understand these facets of AI, but also to communicate these facets when describing future solutions to capability gaps.

To close these capability gaps, DOD acquisition personnel will ultimately engage industry. This quickly growing industry has many powerful and influential corporations, creative figures that are driving its direction, and small businesses that are aiming to leverage the new technology for defense use. This industry is generally driven by massive investment from the commercial sector, so defense and federal government actors will have to engage this new community to ensure DOD demands are considered and sufficient regulations are in place across the field.

2. How Well Does the Joint Concept of Operations Account for AI Technologies and How Does That Impact The JCIDS?

The DOD is in a consistent state of planning for the future. As the world changes through externalities and technology development, so should the plans. Strategic documents built from the services' requirements management branches and derived from national security priorities, explain the operational application of expected future technologies. These documents, assisted by technological research, joint technology demonstrations, and technological based exercises, all work to inform the joint CONOPs. The CONOPs outlines a potential future battle space and the systems we plan to navigate that battle space with. The CONOPS is a primary input to the CBA, which is the process we use to outline capability gaps and validated requirements.

The documents, demonstrations, and exercises that inform our concepts are under addressing the essentials of AI. The teams and personnel that should be incorporating smart technology use into our planning effort, do not know the correct vocabulary, understand the spectrum of AI autonomy, nor have a sense to the training required to make machines behave humanly. Insertion of AI essentials, specifically, understanding of knowledge base, the concept of internal versus external system sensing, mobility, and AI components would greatly increase the effectiveness of the CONOPs.

Since the CONOPs is not correctly considering AI essentials, it is expected that outputs from the JCIDS will not be sufficient. That is to say, requirements validated by the JROC will not properly respect AI system functionality and behavior. This research was not able to analyze the actual outputs of the JCIDS in the form of ICDs or CDDs. It does not prove that outputs from JCIDS are corrupted. It does prove that inputs are corrupted, and given the processes inherent to JCIDS, there are expected negative consequences in processes and outputs. Likely, AI requirements will surface as DOD processes trudge forward through great effort by dedicated people, but the accuracy of those AI requirements will suffer.

3. If Poor Requirements Are Transitioned from JCIDS to DAS, What Problems Will AI Encounter?

Once a warfighter has decided what future battle field requirements are, the DAS starts several processes to inform, build, and procure these systems. Basic science research is conducted to develop the technology necessary to achieve the future plans. Industry is engaged in order to understand the systems that are immediately available and what the contractor pool is capable of developing. Funds are requested and contracts are signed to achieve the goals stated by the warfighter. The requirements that emerge from JCIDS are a primary catalyst for DAS action.

The most basic systems engineering effort in DAS is defining system requirements. These represent the validated requirements for JCIDS, the realm of possible functionality, and the acquisition landscapes various rules and regulations. System requirements are the complete foundation for the system being built. If the outputs from JCIDS are poor, system requirements will be inherently corrupted.

If the system requirements are poor in definition or accuracy, then the wrong investments are made in basic sciences. The wrong applications of available technology are pursued and conceptual systems are analyzed for the wrong performance. If poor AI requirements induce this behavior in DAS, the wrong technology applications will transition from the S&T community to defense programs.

Testing and Evaluation works to determine whether or not the system purchased by the government meets what was defined in the contract and meets what the warfighter needs. It achieves this by decomposing system requirements into quantifiable and testable measures, that when achieved prove the system can perform in the ways requested by the warfighter. Inaccuracies in system requirements will carry forward through the testing process. The consequence is failure to achieve the stated system performance, or the achieved system performance will not match what the warfighter actually needs. If the system requirements contain a quality perception of AI essentials, then it is much more likely that testing and evaluation will benefit the system development.

Quality AI systems will only be useful if they are available to the warfighter in the quantities, capacities, and durations required. To deliver this system availability requires systems life cycle planning, based on quality system requirements. Modifications are required to maintain system viability throughout its lifecycle, and confusion in system requirements handicaps that ability. Without consistent inclusion of AI essentials in validated and system requirements, life cycle effort will cause cost, schedule, and performance, problems for AI systems that generally will not be uncovered until after the system is fielded.

If validated requirements are low in AI quality, the problems that plague software acquisition will plague AI acquisition. The end result will be delays to schedule, poor performance, and inflated system costs.

C. PRIMARY RESEARCH QUESTION

The primary research question is:

What Problems AI Based Systems Expecting to Encounter as They Transition from Basic Science to Executable Program?

To answer the primary research question requires a synthesis of the answers to the secondary research questions. By considering them totally, a true sense of acquisition viability for AI systems can be uncovered. This does not suppose that a complete understanding of the potential problems for AI system acquisition are uncovered, but rather the problems innate to JCIDS and DAS given the highlighted AI essentials, are addressed. Influences from organizational structure, world events, and changes in political structure could all have an impact, but are not addressed in this research.

The first problem facing AI based systems, is that DOD personnel will be unable to recognize and define AI. AI is both a scientific field and the assembly of AI components into a humanly behaving machine. It is related to computer science, which is generally not well understood throughout defense acquisition. At the same time, there are independent assumptions to AI that, unlike computers, deals with mobility and the directionality of sensing. Currently, these facets are not well understood by the personnel engaged in planning for future battle spaces, nor are they understood well by the regulators, or those that have the duty of defining industry safety standards. The consequence of poorly understanding AI essentials is defense personnel will poorly define throughout the DOD. Planning efforts will fail to articulate AI's use or expected behavior. Warfighters will be unable to picture the uses or abilities that AI systems offer, and they will shy away from defining what is actually most helpful to them.

The second problem is that JCIDS will output poor validated requirements. This research proves that inputs to JCIDS are flawed. The CONOPs does not have a good sense of the technology or its uses. Planning agents are not up to the task of articulating AI essentials, specifically in terms of the spectrum, definition, and components required to make the JCIDS process function properly. The logical conclusion is that CBA process would churn, outputting insufficient requirements. Garbage into the JCIDS will cause garbage out. This research cannot prove the JCIDS outputs will be flawed. Additionally, this research does not prove how software intensive requirements became misaligned with hopeful system performance. Still, the research does identify that software intensive systems suffered from low quality requirements. These can only be built using the JCIDS process and growth in systems requirements points to poorly understood technology

throughout the JCIDS. The conclusion then, is validated requirements will be corrupted if AI is not understood well. The flawed validated requirements will push forward into the acquisition world and no sense of their flaws will be available.

The third problem is cost, schedule, and performance goals will continue in the same fashion as they have for software intensive systems. While this research has not proven the AI systems will suffer this fate. It does prove that software intensive systems succumbed to these problems, largely due to poorly defined requirements. There is no reason to believe AI system will be any different than software intensive systems. This conclusion, is based on the assumption that poorly defined validated requirements lead to poorly defined system requirements. Further, that poorly defined system requirements have a waterfall effect causing downstream management processes to become unmanageable or even impossible. If requirements analysis cannot account for the misstep in the JCIDS process or fails itself to correctly understand AI essentials, the confluence of events that cause management problems is triggered. Technology transition, test and evaluation, and life cycle management cease to be management techniques and are merely bureaucratic hurdles. Every day that schedules increase, increases in program costs by the burn rate occur. Just like many software intensive systems, an AI system would be delivered late and over budget, that doesn't meet the needs of the warfighter.

D. RECOMMENDATIONS

Fortunately, AI is still a young science and AI technology is not yet widespread. The DOD has time and resources that can better prepare personnel for operating inside of the JCIDS and DAS. This research has developed four recommendations for the DOD based the results of the literature review, JCIDS analysis, and DAS analysis.

The first is to share this project. It isn't a perfect representation or all-encompassing of the problems facing AI, but it does announce there is a pending problem. This research did not uncover any other analysis that presents AI as a potential problem for defense acquisition. By understanding that each technology has its peculiarities, software acquisition has had its share of problems, and that planning efforts

are under addressing AI defense acquisition personnel can at least be forewarned, if not already trained.

The second is that defense acquisition processes must label and include the AI component of practice. This notion of practice, is identified in this research as an AI component and important for several reasons. First, it is a component that is non-physical. It is attached to a program like a testing strategy or life cycle plan, but is also embedded in the brain of the processor. The personnel that prepare mission plans and exercises will need to be able to define the level of AI performance required to satisfy a mission. This will drive the validated requirements and system requirements that emerge. Additionally, it is the first time a weapon system will require the same amount, or more training than the user. PMs will have to incorporate practice as another management process inside of DAS.

In 2015, Task and Purpose Magazine published an article outlining the fact that the United States Air Force has a pilot shortage. In the article, they cite government calculations that estimate the cost of training a pilot at \$6 million and that F-22 pilots require over three years of specialized training before they are ready to perform their jobs (Gjertsen, 2016). Understanding AI essentials, specifically practice, means the machine demands a similar investment to training. The plane and the pilot will each need to be trained, which requires organizing resources, budgeting time, and achieved levels of performance before certification. Generally, the other AI essentials, when managed through the DAS with discipline, will not pose this significant of a threat to the program. They are more incremental in nature than punctuated. Practice, however, can change the schedule, demand for testing, demand for operator support, and costs of training significantly. Understanding that practice is a component of an AI system is imperative.

The third recommendation is to establish an AI acquisition class or curriculum. Preparing personnel for the challenges of the future is commonly accepted as inherent for the DOD and its members. At present defense personnel who are engaged in planning efforts and requirements definition are lacking the vocabulary and conceptual understanding of AI essentials. DAS personnel are prepared for systems engineering and software acquisition, but few are ready for systems that branch behavior and function

humanly. To improve the products that enter and exit the JCIDS and DAS, a curriculum could be implemented focusing on the AI essentials outlined in this research. Similar software acquisition courses have been developed for use at DAU and NPS, given the struggles acquisition has had with software intensive systems (Skertic, n.d.). Matching AI essential training to program technology demands is a boon for programs overall.

The fourth recommendation is to institute AI into process control, before attempting to institute AI into mobile systems. Yes, robots are flashy and depict the essentials of AI well, but AI is equally good at making generalizations concerning massive amounts of data. Additionally, this research has revealed that system acquisition using DAS is challenging at best. The GAO reports studied outlined many programs and even the successful attempts contained cost and schedule over runs. It is possible to use AI to analyze the CONOPs, determine if it is sufficient, insert missing information, and output quality requirements for nearly any technology base. JCIDS and DAS experts could train the process control system, similar to IBM's Watson, concerning the context for defense acquisition. Why struggle through years of failed AI system programs, when the very technology can eliminate the human inefficiencies inherent to the DAS?

The final recommendation is to implement an AI system designator into the DOD 5000 series. The DODD 5000.01 already includes the designation of *Software Intensive Systems*, which garners it respective instructions throughout the DODI 5000.02 (DOD, 2007, p. 9). Identifying a class of technology in the regulation, declares to acquisition personnel that this technology behaves differently, requires specific understandings to manage, and additional training could be required. This research highlighted that AI is susceptible to the same problems as software intensive acquisition and takes the complexity of technology to another higher level. Perhaps the label of software intensive would suffice? Still, AI behaves fundamentally differently than software given that it must consider direction of sensing, practice, and the potential for branching of tasks. This requires knowledge points that are discrete and independent from software. Informing the JCIDS and DAS, that a system is classified as AI system would alert personnel that the essentials outlined in this research apply.

These recommendations are not a complete list of the steps necessary to achieve success in AI system acquisition. They are a small list that seeks to solve the immediate problems identified in this research. The problems are the inability to recognize and define AI, poorly defined requirements, and failure to achieve cost, schedule, and performance goals. Developing an AI essential curriculum could address all three problems by educating defense personnel to understand AI and prepare for the failures that can occur in the JCIDS and DAS. Incorporation of practice is acquisition programs, as a funded objective and mandated process, directly addresses the risks that AI practice can place on a system. Just as training must be considered, considering practice, can alleviate costly schedule delays and performance failures. Implementation of AI terminology into defense regulations attempts to address all problems by informing defense personnel that this is a technology they must contend with. If there isn't a class for AI, then acquisition personnel can at least get smart on their own and be ready to develop requirements, system requirements, and manage technology transition, testing, and life cycle efforts. By

E. RECOMMENDATIONS FOR FUTURE RESEARCH

This section lists prime opportunities for future research that were uncovered based on this research. These subjects often revealed themselves during writing as something that could enhance a particular section. This research found that much research has been done into the basic science that is AI (algorithms and techniques), there was a very limited research into the possible uses. Using the structure of this research as a guide, recommendations for future research are organized into three parts based on technological dynamics, operational uses, and DAS solutions.

In the realm of AI technology, specific applications could be analyzed for usefulness to meet the demands outlined in today's CONOPs. This could include research into COTS solutions available to support today's operations or compiling commercial technologies that are mature enough to be incorporated into today's programs. Current CONOPs could even be reworked, including the AI essentials to enhance their ability to translate information into JCIDS. Additionally, research could be

done into the algorithms that would specifically benefit the leading edge of defense systems.

In an operational context, key areas of future research would include how to fight with AI systems. Development of new AI CONOPs for specific battle spaces would go a long way to getting the DOD started with understanding how we expect to fight in future wars with this technology. Special focus should be taken to consider how that research informs the joint CONOPs and revealing the tactical benefits that could be achieved given the nature of AI. Another area for research, in the realm of operational use, would be to strategize how to AI could support the JCIDS process. Software agents could eliminate human processes, automatically reduce risks inherent in requirement generation, and increase the pace at which requirements can be developed. A roadmap concerning how to implement AI to achieve those benefits would be useful.

Finally, research into management practices that will enhance DAS performance when fielding AI systems would be helpful. Development of an AI specific acquisition strategy, including practice and the expectation for branching behavior, would greatly service PMs. Additionally, the application of AI support to assist PMs through DAS processes would be helpful. Defining the system requirements that an AI process management system, capable of supporting DAS, would be the first step to harnessing the process support power of AI. Remember, software agents are equally as good at supporting decisions and processes as robots are for assisting warfighters in the field.

LIST OF REFERENCES

- 414th Combat Training Squadron “Red Flag.” (2012, July 06). Retrieved October 17, 2017, from <http://www.nellis.af.mil/About/Fact-Sheets/Display/Article/284176/414th-combat-training-squadron-red-flag/>
- AcqNotes (n.d.) Technology Readiness Level (TRL). Retrieved October 04, 2017, from <http://acqnotes.com/acqnote/tasks/technology-readiness-level>
- Air Force Research Laboratory. (2016, October 03). AFRL Tech Warrior 2016. Retrieved October 20, 2017, from <https://www.youtube.com/watch?v=cb2yHzGqFO4>
- Amazon (2017). 2016 Annual Report Form 10-K. Retrieved September 26, 2017, from <https://www.sec.gov/Archives/edgar/data/1018724/000101872417000011/amzn-20161231x10k.htm>
- Amazon. (n.d.). Alexa. Retrieved September 26, 2017, from https://www.amazon.com/b/?ie=UTF8&node=9818047011&tag=mh0b-20&hvadid=77721756043382&hvqmt=e&hvbmt=be&hvdev=c&ref=pd_sl_iwlt1gvek_e
- Apple Inc. (2017, February 02). 2016 Annual Report Form 10-K. Retrieved September 28, 2017, from <https://www.sec.gov/Archives/edgar/data/320193/000162828016020309/a201610-k9242016.htm>
- Army Capabilities Integration Center [ARCIC] (n.d.). Army warfighting challenges. Retrieved October 17, 2017, from <http://www.arcic.army.mil/Initiatives/ArmyWarfightingChallenges>
- Army Capabilities Integration Center Future Warfare Division. (2016, January 29). *UQ 16 Future Force Design II Final Report*. Fort Eustis, VA.
- Army Capabilities Integration Center. (2017). *The U.S. Army Robotic and Autonomous Systems Strategy*. Fort Eustis, VA.
- Artificial Intelligence. (n.d.). In *Merriam-Webster’s* online dictionary. Retrieved September 19, 2017, from <https://www.merriam-webster.com/dictionary/artificial%20intelligence>
- Bacchus, A. (2017, August 14) Former Microsoft veteran Qi Lu opens up about AI, how Amazon beat Microsoft. *On Msft*. Retrieved September 29, 2017, from <https://www.onmsft.com/news/former-microsoft-veteran-qi-lu-opens-up-about-ai-how-amazon-beat-microsoft>

- Bezos, J. P. (2017). "What does day 2 look like?" [Letter written 2017 to Amazon Shareholders]. Retrieved September 29, 2017, from <https://www.geekwire.com/2017/full-text-annual-letter-amazon-ceo-jeff-bezos-explains-avoid-becoming-day-2-company/>
- Biography.com. (2017, April 11). Elon Musk. Retrieved September 29, 2017, from <https://www.biography.com/people/elon-musk-20837159#!>
- Bultman, L. A. (2017, January 25). Red flag evolves as ISR, cyber presence increases. *25th Air Force Press*. Retrieved from <http://www.nellis.af.mil/News/Article/1061770/red-flag-evolves-as-isr-cyber-presence-increases/>
- Carey, B. (2017, August 15). USSOCOM plans 'ThunderDrone' technology demonstration. *AIN online*. Retrieved from <https://www.ainonline.com/aviation-news/defense/2017-08-15/ussocom-plans-thunderdrone-technology-demonstration>
- Chairman of the Joint Chiefs of Staff. (2011). *Unmanned systems integration roadmap FY2011-2036* (11-S-3613). Retrieved October 17, 2017, from <https://www.unols.org/sites/default/files/usroadmap2011.pdf>
- Chairman of the Joint Chiefs of Staff. (2012, January 19). *Manual for the operation of the joint capabilities integration and development system*. Washington, D.C.: Author
- Chairman of the Joint Chiefs of Staff. (2015, January 23). *Joint capabilities integration and development system (JCIDS)* (CJCSI 3170.0I). Washington, DC. Retrieved 18 November, 2017 from <http://acqnotes.com/wp-content/uploads/2014/09/CJCS-Instruction-3170-01I-Joint-Capabilities-Integration-and-Development-System-23-Jan-15.pdf>
- Clark, J., Welinder, P., McGrew, B., Schneider, J., Duan, R., Tobin, J., Fong, R., Ray, A., Wolski, F., Kumar, V., Ho, J., Andrychowicz, M., Stadie, B., Handa, A., Plappert, M., Reinhardt, E., Abbeel, P., Brockman, G., Sutskever, I., & Wojciech, Z. (2017, July 26). Robots that learn. *OpenAI*. Retrieved September 26, 2017, from <https://blog.openai.com/robots-that-learn/>
- Cooper, D. E. (1993) *Test and evaluation: DOD has been slow in improving testing of software-intensive systems United States* (GAO/NSAID-93-198). Washington, DC: General Accounting Office.
- Cowley, D. (2015, December 01). SoarTech and Epic Games announce official partnership for government and military applications. Retrieved September 28, 2017, from Unreal Engine website: <https://www.unrealengine.com/en-US/blog/soartech-partnership-government-military-applications>

- Crevier, D. (1995). *AI: the tumultuous history of the search for artificial intelligence*. New York, NY: Basic Books.
- Defense Acquisition University. (2005). *Manager's guide to technology transition in an evolutionary acquisition environment*. Fort Belvoir, VA: Defense Acquisition University Press.
- Defense Acquisition University. (2013, September 16). *Defense acquisition guidebook*. Retrieved October 31, 2017, from <https://www.dau.mil/tools/dag>
- Defense Advanced Research and Projects Agency. (2016, August 08). *DARPA's Cyber Grand Challenge: Final Event Program*. Retrieved October 20, 2017, from https://www.youtube.com/watch?time_continue=8452&v=n0kn4mDXY6I
- Defense Intelligence Agency (2014). *2014 DIA innovation strategic plan*. Washington, DC: Author.
<http://www.dia.mil/Portals/27/Documents/Business/Innovation/DIA%20Innovation%20Strategy.pdf>
- Department of Defense. (2007, May 12). *The defense acquisition system (DOD Directive 5000.01)*. Washington, DC: USD (AT&L). Retrieved September 29, 2017 from <http://www.acqnotes.com/Attachments/DoD%20Directive%205000.01.pdf>
- Department of Defense. (2012). *Test and evaluation management guide (6th ed.)*. Fort Belvoir, VA: The Defense Acquisition University.
- Department of Defense. (2017, August 10). *Operation of the defense acquisition system (DOD Instruction 5000.02)*. Washington, DC: USD (AT&L). Retrieved September 29, 2017 from <http://acqnotes.com/wp-content/uploads/2014/09/DoD-Instruction-5000.02-The-Defense-Acquisition-System-10-Aug-17-Change-3.pdf>
- Department of the Air Force. (2003, February). *Guidelines for successful acquisition and management of software-intensive systems (Condense Version)*. Software Technology Support Center. Retrieved November 16, 2017 from <http://www.toadland.net/works/GSAM%20V4.pdf>
- Domonoske, C. (2017, July 17). Elon Musk warns governors: artificial intelligence poses 'existential risk'. *NPR*. Retrieved September 29, 2017, from <http://www.npr.org/sections/thetwo-way/2017/07/17/537686649/elon-musk-warns-governors-artificial-intelligence-poses-existential-risk>
- Facebook Research. (n.d.). Research areas. Retrieved September 26, 2017, from <https://research.fb.com/>

- Facebook. (2017, February 2). 2016 Annual Report. Form 10-K. Retrieved September 1, 2017, from https://s21.q4cdn.com/399680738/files/doc_financials/annual_reports/FB_AR_2016_FINAL.pdf
- Fraze, D. (n.d.). Cyber Grand Challenge (CGC). Program information. Defense Advanced Research Projects Agency. Retrieved October 20, 2017, from <https://www.darpa.mil/program/cyber-grand-challenge>
- Frenzel, L. E. (1987). *Crash course in artificial intelligence and expert systems*. Indianapolis, IN: Howard W. Sams & Co
- Galeon, D. Gohd, C. (2017, May 09). Amazon's CEO says we're living in the golden age of AI. *Futurism*. Retrieved September 29, 2017, from <https://futurism.com/amazons-ceo-says-were-living-in-the-golden-age-of-ai/>
- Gjertsen, M. (2016, February 19). Why the Air Force is paying big bonuses to some pilots and forcing others out. *Task and Purpose*. Retrieved November 14, 2017, from <http://taskandpurpose.com/air-force-paying-big-bonuses-pilots-forcing-others/>
- Google (2017, February 02). 2016 Annual Report. Form 10-K. Retrieved September 01, 2017, from https://abc.xyz/investor/pdf/20161231_alphabet_10K.pdf
- Haskins, C. Forsberg, K., Krueger, M. Walden, D., Hamelin, R.D., (2011 October). *Systems engineering handbook: A guide for systems life cycle processes and activities*. San Diego, CA: International Council on Systems Engineering.
- Helft, M. (2009, August 30). A hired gun for Microsoft, in pursuit of Google. *New York Times*. Retrieved September 29, 2017, from <http://www.nytimes.com/2009/08/31/technology/internet/31search.html>
- Hempel, J. (2017, August 11). How Baidu will win China's AI race-and, maybe, the world's. *Wired*. Retrieved September 29, 2017, from <https://www.wired.com/story/how-baidu-will-win-chinas-ai-raceand-maybe-the-worlds/>
- Huang, S., Papernot, N., Goodfellow, I., Duan, Y., & Abbeel, P. (2017, February 8). *Adversarial attacks on neural network policies*. Retrieved September 1, 2017, from Cornell University Library: <https://arxiv.org/pdf/1702.02284.pdf>
- IBM Think Academy. (2014, November 21). *How it works: IBM Watson* [video file]. Retrieved September 26, 2017, from https://www.youtube.com/watch?time_continue=467&v=AtdJ1DGJjXA
- IBM. (2001, February 23). How deep blue works. IBM research. Retrieved September 19, 2017, from <https://www.research.ibm.com/deepblue/meet/html/d.3.2.html>

- International Business Machines Corporation. (2017, February 02). 2016 Annual Report. Form 10-K. Retrieved September 26, 2017, from <https://www.sec.gov/Archives/edgar/data/51143/000104746917001061/a2230222z10-k.htm>
- Jones, M. T. (2009). *Artificial intelligence: A systems approach*. Sudbury, MA: Jones and Bartlett.
- Knight, W. (2017, June 22). Tesla's new AI guru will help its cars learn for themselves. MIT technology review. Retrieved September 29, 2017, from <https://www.technologyreview.com/s/608155/teslas-new-ai-guru-could-help-its-cars-teach-themselves/>
- Kropas-Hughes, C., Rutledge, L., & Sarmiento, G. (2008, September 26). *High confidence technology transition planning through the use of stage-gates* (TD-13) [Presentation]. Retrieved October 26, 2017, from http://www.afadaytonwright.com/Downloads/TD-13_Kropas-Hughes_Claudia_080926-Final.pdf
- Le, C., Sgt. (2016, August 05). Darkhorse' Marines assault California during MAGTF integrated experiment 2016. *U.S. Marine Corps*. Retrieved October 17, 2017, from <http://www.marines.mil/News/News-Display/Article/905784/darkhorse-marines-assault-california-during-magtf-integrated-experiment-2016/Print>
- Leswing, K. (2017, February 01). Apple is spending billions on secret R&D projects - and it keeps spending more. *Business Insider*. Retrieved September 28, 2017, from <http://www.businessinsider.com/apple-rd-spend-charts-2017-2>
- Li, S. (2015, July 24). Amazon overtakes Wal-Mart as biggest retailer. *Los Angeles Times*. Retrieved September 26, 2017, from <http://www.latimes.com/business/la-fi-amazon-walmart-20150724-story.html>
- Littlefield, S. (n.d.). Anti-Submarine Warfare (ASW) Continuous Trail Unmanned Vessel [ACTUV]. Program information. Defense Advanced Research Projects Agency. Retrieved September 11, 2017, from <https://www.darpa.mil/program/anti-submarine-warfare-continuous-trail-unmanned-vessel>
- Lockheed Martin. (n.d.). *The F-35 and advanced sensor fusion*. Retrieved October 17, 2017, from <http://www.sldinfo.com/whitepapers/the-f-35-and-advanced-sensor-fusion/>
- Lopez-Paz, D., Nishihara, R., Chintala, S., Scholkopf, B., & Bottou, L. (2017, July 2017). *Discovering causal signals in images*. Facebook Research. Retrieved September 5, 2017, from <https://research.fb.com/publications/discovering-causal-signals-in-images/>

- Marine Corps Warfighting Laboratory. (2016, August 05). *Darkhorse Attacks MIX-16*. Retrieved October 17, 2017, from <http://www.marines.mil/News/Marines-TV/videoid/478372/dvpTag/MCWL/>
- Marine Corps. (September 2016). *Marine Corps operating concept*. Retrieved October 17, 2017, from <http://www.mccdc.marines.mil/Portals/172/Docs/MCCDC/young/MCCDC-YH/document/final/Marine%20Corps%20Operating%20Concept%20Sept%202016.pdf?ver=2016-09-28-083439-483>
- NASA (n.d.). Curiosity overview. Mars curiosity. NASA. Retrieved September 19, 2017, from https://www.nasa.gov/mission_pages/msl/overview/index.html
- National Defense Industrial Association. (2010, September). *Top software engineering issues within the Department of Defense and defense industry* (Report No. 5a). Arlington, VA: Author.
- Neale, M. C. & Cardon, L. R. (2011). *Methodology for genetic studies of twins and families*. Dordrecht, ND: Kluwer Academic Publishers B.V. Retrieved September 5, 2017, from <http://ibgwww.colorado.edu/workshop2006/cdrom/HTML/book2004a.pdf>
- Nuñez, M. (2015, June 25). Amazon echo is the first artificial intelligence you'll want at home. *Popular Science*. Retrieved September 26, 2017, from <http://www.popsci.com/amazon-echo-first-artificial-intelligence-youll-want-home>
- NVIDIA. (n.d.-a). History: A timeline of innovation. Retrieved September 28, 2017, from http://www.nvidia.com/page/corporate_timeline.html
- NVIDIA. (n.d.-b). DRIVE PX2. Autonomous Car Development Platform. Retrieved September 28, 2017, from <http://www.nvidia.com/object/drive-px.html>
- NVIDIA. (n.d.-c). Deep learning and AI. Products. Retrieved September 28, 2017, from <https://www.nvidia.com/en-us/deep-learning-ai/solutions/>
- NVIDIA. (n.d.-d). NVIDIA Jetson. Embedded systems. Products. Retrieved September 28, 2017, from <http://www.nvidia.com/object/embedded-systems-dev-kits-modules.html>
- Office of the Secretary of the Navy. (2015, March). *A cooperative strategy for 21st century seapower*. http://www.au.af.mil/au/awc/awcgate/maritime/maritime_strat_oct07.pdf
- Office of the U.S. Air Force Chief Scientist (2011, September) *Technology horizons: A vision for Air Force science and technology 2010–30* (AF/ST-TR-10-01-PR). Maxwell Air Force Base, Ala.: Air University Press, Air Force Research Institute.

- OpenAI. (n.d.). About OpenAI. Retrieved September 26, 2017, from <https://openai.com/about/>
- Pearl, J. (1985). *Heuristics: Intelligent search strategies for computer problem solving*. Taipei: Ju Lin.
- Russell, C. (2014, September). *F-35 sustainment: Need for affordable strategy, greater attention to risks, and improved cost estimates* (GAO-14-778). Washington, DC: Government Accountability Office.
- Russell, S. J., & Norvig, P. (2010). *Artificial intelligence: a modern approach* (3rd). Upper Saddle River, NJ: Pearson.
- Schinasi, K. V. (2004, March) *Defense acquisitions: stronger management practices are needed to improve DODs software-intensive weapon acquisitions* (GAO-04-393). Washington, DC: General Accounting Office.
- Schultz, L., Cpl. (2016, August 08). Combat center hosts Marine Corps Warfighting Lab. *U.S. Marine Corps*. Retrieved October 17, 2017, from <http://www.29palms.marines.mil/News/News-Article-Display/Article/907961/combat-center-hosts-marine-corps-warfighting-labs-magtf-integrated-experiment-2/>
- Simon, H. A. (1985, Spring). Artificial intelligence: Current status and future potential. *The Charles H. Davis lecture series* (Ser. 10, pp. 5–23). (1985). Washington, D.C.: National Academy Press.
- Sinbad, S. (2017, February 3) *Red Flag debrief Jan-Feb 2017 USAF, RAAF & RAF participants* [video file]. Retrieved October 17, 2017, from <https://www.youtube.com/watch?v=DkVMaN1i3tE>
- Skertic, R. P. (2015, February). Training the DOD software acquisition professional. *Defense Acquisition University*. Retrieved November 14, 2017, from <http://www.dtic.mil/docs/citations/ADA615595>
- Sofwerx. (2017). ThunderDrone rapid prototyping event [presentation].
- Sofwerx. (n.d.). ThunderDrone. Retrieved October 20, 2017, from <https://www.sofwerx.org/pdf-wrappers/>
- SpaceX. (n.d.). About. Retrieved September 29, 2017, from <http://www.spacex.com/about>
- SparkCognition Inc. (n.d.). DeepArmor. Retrieved September 05, 2017, from <https://sparkcognition.com/deeparmor-2/>

- Sullivan, M. J. (2015, June) *Defense acquisition process: Military service chiefs' concerns reflect need to better define requirements before programs start* (GAO-15-469). Washington, D.C.: Government Accountability Office. Retrieved October 27, 2017, from <http://www.gao.gov/assets/680/670761.pdf>
- Sullivan, M. J. (2017, April) *F-35 joint strike fighter: DOD needs to complete developmental testing before making significant new investments* (GAO-17-351). Washington, DC: Government Accountability Office.
- Tech Warrior: AFRL scientists gain insights during field exercise. (2017, September 22). *88th Air Base Wing Office of Public Affairs*. Retrieved October 17, 2017, from <http://www.daytondailynews.com/news/tech-warrior-afrl-scientists-gain-insights-during-field-exercise/jaY0fVsXIYIsDOOrAYLYHJ/>
- Truskowski, W., Hallock, H.L., Rouff, C., Karlin, J., Rash, J., Hinchey, M., & Sterrit, R. (2009). *Autonomous and autonomic systems: with applications to NASA intelligent spacecraft operations and exploration systems*. London: Springer.
- United States Air Force. (2010, July). *Air Force technology development and transition strategy guidebook* (Version 2). Wright-Patterson AFB, OH: HQ AFMC/A5S. Retrieved November 15, 2017 from www.acqnotes.com/Attachments/TDTS_Guidebook-v2.doc
- United States Marine Corps (n.d.). The warfighting lab. Retrieved October 17, 2017, from <https://www.marines.com/what-we-do/adapt-and-overcome/warfighting-lab.html>
- Victorino, C. (2017, March 27). SpaceX, Tesla CEO Elon Musks [sic] wants to implant AI devices to human brain. *IBT*. Retrieved September 29, 2017, from <http://www.ibtimes.com/spacex-tesla-ceo-elon-musks-wants-implant-ai-devices-human-brain-2515910>

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California