



Calhoun: The NPS Institutional Archive
DSpace Repository

Theses and Dissertations

1. Thesis and Dissertation Collection, all items

2022-06

EXPLORING THE POTENTIAL OF A MACHINE TEAMMATE

Barton, Andrew C.; Chuprevich, Joel M.

Monterey, CA; Naval Postgraduate School

<https://hdl.handle.net/10945/70629>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>



**NAVAL
POSTGRADUATE
SCHOOL**

MONTEREY, CALIFORNIA

THESIS

**EXPLORING THE POTENTIAL OF A MACHINE
TEAMMATE**

by

Andrew C. Barton and Joel M. Chuprevich

June 2022

Thesis Advisor:

Anthony Canan

Co-Advisor:

Steven J. Iatrou

Approved for public release. Distribution is unlimited.

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>
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 June 2022	3. REPORT TYPE AND DATES COVERED Master's thesis	
4. TITLE AND SUBTITLE EXPLORING THE POTENTIAL OF A MACHINE TEAMMATE			5. FUNDING NUMBERS
6. AUTHOR(S) Andrew C. Barton and Joel M. Chuprevich			
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.			
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.			12b. DISTRIBUTION CODE A
13. ABSTRACT (maximum 200 words) Artificial intelligence has been in use for decades. It is already deployed in manned formations and will continue to be fielded to military units over the next several years. Current strategies and operational concepts call for increased use of artificial-intelligence capabilities across the defense enterprise—from senior leaders to the tactical edge. Unfortunately, artificial intelligence and the warriors that they support will not be compatible “out of the box.” Simply bolting an artificial intelligence into teams of humans will not ensure success. The Department of Defense must pay careful attention to how it is deploying artificial intelligences alongside humans. This is especially true in teams where the structure of the team and the behaviors of its members can make or break performance. Because humans and machines work differently, teams should be designed to leverage the strengths of each partner. Team designs should account for the inherent strengths of the machine partner and use them to shore up human weaknesses. This study contributes to the body of knowledge by submitting novel conceptual models that capture the desired team behaviors of humans and machines when operating in human-machine teaming constructs. These models may inform the design of human-machine teams in ways that improve team performance and agility.			
14. SUBJECT TERMS human-machine teams, HMT, artificial intelligence, AI, team performance, interactive team cognition, ITC, agility, communications, coordination, interaction			15. NUMBER OF PAGES 213
			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.

EXPLORING THE POTENTIAL OF A MACHINE TEAMMATE

Andrew C. Barton
Major, United States Marine Corps
BA, University of Arizona, 2010
MBA, University of Phoenix, 2016

Joel M. Chuprevich
Major, United States Marine Corps
BS, Montreat College, 2006
MA, American Military University, 2019

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN INFORMATION TECHNOLOGY MANAGEMENT

from the

**NAVAL POSTGRADUATE SCHOOL
June 2022**

Approved by: Anthony Canan
Advisor

Steven J. Iatrou
Co-Advisor

Alex Bordetsky
Chair, Department of Information Sciences

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Artificial intelligence has been in use for decades. It is already deployed in manned formations and will continue to be fielded to military units over the next several years. Current strategies and operational concepts call for increased use of artificial-intelligence capabilities across the defense enterprise—from senior leaders to the tactical edge. Unfortunately, artificial intelligence and the warriors that they support will not be compatible “out of the box.” Simply bolting an artificial intelligence into teams of humans will not ensure success. The Department of Defense must pay careful attention to how it is deploying artificial intelligences alongside humans. This is especially true in teams where the structure of the team and the behaviors of its members can make or break performance. Because humans and machines work differently, teams should be designed to leverage the strengths of each partner. Team designs should account for the inherent strengths of the machine partner and use them to shore up human weaknesses. This study contributes to the body of knowledge by submitting novel conceptual models that capture the desired team behaviors of humans and machines when operating in human-machine teaming constructs. These models may inform the design of human-machine teams in ways that improve team performance and agility.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	BACKGROUND	1
	1. The Rise of the Machine.....	1
	2. Don't Just "Sprinkle Some AI on It"	2
	3. Different Flavors of Artificial Intelligence	5
	4. Teams Solve Problems.....	6
	5. Centaurs as a Model	7
	6. We Need Models Grounded in Team Performance.....	8
B.	RESEARCH QUESTIONS, METHODOLOGY, AND PATH	9
	1. Research Questions.....	9
	2. Methodology	10
C.	PURPOSE OF RESEARCH	10
II.	ANALYSIS OF THE DEPARTMENT OF THE NAVY ARTIFICIAL INTELLIGENCE STRATEGY.....	13
A.	DON'S PURPOSE	13
B.	DON'S PROBLEM SPACE.....	14
C.	KEY OPERATIONAL VARIABLES AND A TECHNICAL FRAMEWORK.....	15
D.	TECHNOLOGY ADOPTION.....	17
E.	WHY IT IS IMPORTANT? WHAT IS THE IMPACT OF NOT DOING IT?.....	18
F.	AI AND GREAT POWER COMPETITORS	19
G.	IMPLEMENTATION AND STRATEGY GAPS.....	19
H.	IMPLEMENTATION PLAN	21
I.	ANTICIPATED BENEFITS.....	22
J.	ORDERED EFFECTS.....	23
K.	CHAPTER SUMMARY.....	24
III.	A REVIEW OF TEAM THEORY	25
A.	SHARED MENTAL MODELS.....	25
B.	INTERACTIVE TEAM COGNITION.....	26
	1. Team Member Schemas	27
	2. Teamwork and Taskwork	29
C.	TEAM SITUATIONAL AWARENESS.....	30
D.	COMMUNICATIONS, COORDINATION, AND INTERACTION DYNAMICS IN TEAM CONSTRUCTS.....	31

E.	TEAM PERFORMANCE AND AGILITY	37
F.	SUMMARY	39
IV.	A REVIEW OF COMMAND AND CONTROL CONCEPTS AND DOCTRINE	41
A.	THE NATURE OF COMMAND AND CONTROL.....	41
B.	COMMAND AND CONTROL SYSTEMS.....	42
C.	SEMANTIC EFFECTS ON DECISION MAKING	44
1.	The Nature of Meaning and the Semantic Problem	44
2.	The Cognitive Hierarchy	45
3.	The Decision and Execution Cycle	46
D.	ORGANIZATIONAL STRUCTURES.....	47
1.	Simple Structure	48
2.	Machine Bureaucracy.....	49
3.	Professional Bureaucracy.....	49
4.	Divisionalized Form	50
5.	Adhocracy.....	51
E.	BOYD'S OODA LOOP	51
F.	SUMMARY	53
V.	A REVIEW OF ARTIFICIAL INTELLIGENCE THEORY, CAPABILITIES AND LIMITATIONS.....	55
A.	WHY AI?	55
B.	THE DIFFERENCE BETWEEN AUTOMATION AND AI	56
C.	AI IS ALGORITHM BASED	58
D.	THE LIMITATIONS OF AI.....	62
E.	THE BENEFITS AND CAPABILITIES OF AI	66
F.	CHAPTER SUMMARY.....	69
VI.	HUMAN-MACHINE TEAMING	71
A.	WHAT DOES IT MEAN FOR AI TO BE A TEAMMATE?.....	71
B.	WHY HUMAN MACHINE TEAMS?	72
C.	HUMAN MACHINE TEAMING METHODS AND MODELS	73
1.	Processes and Characteristics of Effective Human Machine Teaming.....	73
2.	Structures, Roles, and Tasks.....	75
3.	Human Machine Team Interactions	78
4.	Design Principles Specific to Human Machine Teaming	81
D.	EFFECT OF AI ON TEAM PERFORMANCE	87
E.	CHAPTER SUMMARY.....	87

VII.	METHODOLOGY	89
A.	VIGNETTE DEVELOPMENT	89
1.	Operation Provide Comfort: AWACS, Black Hawks, and F-15 Crews	89
2.	USS Vincennes Shootdown of Iran Air Flight 655.....	89
B.	APPROACH.....	90
1.	Conceptual Systems Dynamics Models.....	90
2.	System Dynamic Approach Using Causal Loop Diagrams.....	93
3.	System Archetypes	94
C.	MODEL DEVELOPMENT	96
D.	CHAPTER SUMMARY.....	98
VIII.	ANALYSIS	99
A.	DESIRED TYPES OF HUMAN-MACHINE INTERACTION.....	99
B.	VIGNETTE 1: OPERATION PROVIDE COMFORT.....	99
1.	Background	99
2.	Overview	105
3.	Causal Loop Diagrams	111
C.	VIGNETTE 2: USS VINCENNES CASE STUDY	126
1.	Background and Overview.....	126
2.	Causal Loop Diagrams	130
D.	CHAPTER SUMMARY.....	170
IX.	SUMMARY AND FUTURE WORK.....	171
A.	CONCLUSIONS	171
1.	Research Question 1: What does it mean to be a teammate?.....	172
2.	Research Question 2: How can machines be partners with humans in teams?.....	172
3.	Research Question 3: How can the integration of machine partners into previously human-human teams improve the team’s agility?	174
4.	Research Question 4: Can machine partners be better understood with shared team mental models or interactive team cognition approaches?.....	174
B.	ADDITIONAL FINDINGS AND LIMITATIONS.....	175
1.	Additional Findings	175
2.	Limitations.....	176
C.	RECOMMENDATIONS FOR FUTURE WORK.....	176

LIST OF REFERENCES.....179

INITIAL DISTRIBUTION LIST195

LIST OF FIGURES

Figure 1.	Intelligent Autonomous System Technical Framework. Source: DOD (2021a).	15
Figure 2.	Artificial Intelligence Applications and Operational Contexts. Source: Lewis and Vavrichek (2019).	16
Figure 3.	United States Artificial Intelligence-Related Strategies. Source: NIST (2019).	20
Figure 4.	Colonel John Boyd’s OODA Loop. Source: Osinga (2007).....	52
Figure 5.	Simple Neural Network Design	60
Figure 6.	Failure to Communicate.....	112
Figure 7.	AI to Integrate Team Communications.....	115
Figure 8.	Misidentification of Black Hawks	117
Figure 9.	AI as Authenticator	120
Figure 10.	Non-integration of Eagle Flight	122
Figure 11.	AI to Regulate Team Integration	124
Figure 12.	Avoiding the Fate of the USS Stark.....	131
Figure 13.	AI Regulates Team Structure.....	133
Figure 14.	Time Pressure Loops.....	136
Figure 15.	AI Back Up Behavior Completes Tasks and Improves Taskwork Capacity	139
Figure 16.	Ineffective Team Structure and Behaviors	142
Figure 17.	AI Regulates Team Interactions.....	149
Figure 18.	Poor Team Behaviors Erode Communications Patterns and Situational Awareness.....	154
Figure 19.	AI Regulates Team Communications	159

Figure 20.	Interactions Between the USS Vincennes and Iran Air Flight 655	162
Figure 21.	AI Enables Higher Levels of Situational Awareness.....	168

LIST OF ACRONYMS AND ABBREVIATIONS

ACE	airborne command element
AI	artificial intelligence
AIM	advanced intercept missile
AMRAAM	advanced medium range air to air missile
ASO	air surveillance officer
ATO	air tasking order
AWACS	airborne warning and control system
C2	command and control
CEO	chief executive officer
CFAC	combined forces air component
CTF	combined task force
DOD	Department of Defense
DON	Department of the Navy
GPHIN	Global Public Health Intelligence Network
HMS	Her Majesty's Ship
HMT	human machine team
IFF	identify friend or foe
ISR	intelligence, surveillance, and reconnaissance
JADC2	Joint All Domain Command and Control
JAIC	Joint Artificial Intelligence Center
JSOC	joint special operations component
JTF	joint task force
LoE	line of effort
MCC	military coordination center
NATO	North Atlantic Treaty Organization
NIST	National Institute for Standards and Technology
OODA	orient, observe, decide, act
RMA	revolution in military affairs
SD	senior director
TAO	tactical action officer

TAOR	tactical area of responsibility
TD	target designator
TSA	team situational awareness
TTP	tactics, techniques, and procedures
USS	United States Ship
VID	visual identification

ACKNOWLEDGMENTS

This research and thesis is the result of interactions with numerous staff and faculty at the Naval Postgraduate School. We are indebted to them all.

In particular, we would like to thank Dr. Mustafa Canan and Professor Steve Iatrou. Both are world-class instructors and researchers who provided us with many hours of inspiring debate. Their challenging ideas helped to propel us forward in reaching the finish line. Hopefully, the concepts and models we present in this study will be used in the future defense of our country.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

Nuclear missile launches were detected. It was October 5, 1960, and NATO was at the highest level of alert. With 99.9 percent accuracy, incoming Soviet ballistic missiles were detected by early warning systems in Greenland. Thankfully, NATO's retaliation was halted, and operators figured out that the "smart" system was tracking the rising moon (Singer, 2009). Naturally, this was not the only time the world almost perished in an artificial intelligence (AI) induced nuclear exchange. On September 26, 1983, Lieutenant Colonel Petrov of the Soviet Union found himself as the duty officer inside the Serpukhov-15 bunker near Moscow. Orbiting in space, the Soviet Oko early warning satellite system reported with complete certainty that multiple missiles were on their way to Moscow. The problem was the Oko system misidentified sunlight reflecting from the cloud tops as a series of American missile launches (Scharre, 2018). Deciphering the limitations of their systems, and placing events in context, operators were able to prevent catastrophe. Surely these extreme cases are outliers, and we are not as reliant on human judgement for today's artificial intelligence, right? Unfortunately, not quite. Humans and artificial intelligence (AI) work differently, so organizations such as the Department of Defense (DOD) need to be very deliberate when they insert AI-enabled systems into operational teams.

A. BACKGROUND

1. The Rise of the Machine

AI has been here for decades. It is already deployed with human formations and will continue to be fielded to more units over the next several years. People have been reliant on machines and variants of automation for decades, yet humans still have not learned how to make them context sensitive and adaptive. Most are intimately familiar with the unintended effects that plague technologies. Siri babbles at random, the Nest thermostat converts the house into an oven, and the Roomba fails to clean the house to our satisfaction. Yet, we are not worried. Often, the failures of machines in these narrowly defined environments simply perplexes us, leaving us to ask, "How can it fail this badly? How can such simple tasks be so painful?" Unfortunately, warriors and their AI teammates will not

be compatible out of the box. Intuitively, this makes sense—even the Siri application on one’s iPhone spends at least some time learning its owners preferences. Even Watson, the AI-based Jeopardy contestant, was chided when it gave obviously wrong answers to questions put to it by Alex Trebek (Johnston, 2011). Admittedly, there is a significant difference between a robot that cleans your house, and a system that fights alongside humans. However, the way that these systems are designed has significant implications for warfighters that will rely on them. Specifically, the behaviors and dynamics that characterize high-performing teams are conspicuously absent in the design of human machine teams.

2. Don’t Just “Sprinkle Some AI on It”

It has been argued that AI will create revolutionary new capabilities that will enable the naval service to create new tactical dilemmas and challenges for adversaries (United States Marine Corps [USMC], 2018). A revolution in military affairs (RMA) is the term used to describe a paradigm shift in warfare, typically spurred due to the widespread adoption of innovative technologies (Sloan, 2002). These shifts have the potential ability to render previous doctrine and proven concepts obsolete.

In informational, economic, and military contexts, nations or blocs that develop advanced artificial intelligence capabilities will realize increased productivity and will enhance their ability to innovate, allowing them to outcompete their adversaries (Russell, 2019). Agreeing, DOD established the Joint Artificial Intelligence Center (JAIC) and has recognized that AI is a revolutionary, enabling technology essential in creating and maintaining a competitive advantage over adversaries (Department of Defense [DOD], 2020). The DOD’s goals for deploying AI alongside people are simple: help warfighters make better decisions and ensure that our forces thrive in dangerous environments (Bray & Moore, 2021). With the anticipated advantages of human-machine teams, a new way of fighting and winning wars is possible (Singer, 2009).

Among calls for artificial intelligence to ensure the success of the enterprise are those for applications to Joint All Domain Command and Control (JADC2)—an internet of things-like web of sensor capabilities that will require artificial intelligence to process

sensor inputs into targeting data to close the kill chain (Hoehn, 2021). It would be even more desirable if AI capabilities could process said sensor inputs and automatically provide actionable recommendations to fires platforms. Similarly, initiatives for senior leader decision support and making data-driven decisions require some artificial intelligence algorithms (DOD, 2020). Project Overmatch, and the swarms that will enable distributed maritime operations, requires close teaming between artificial intelligences and the humans they support (Office of the Chief of Naval Operations, 2020; DON, 2021). In the DOD, the implications are that artificial intelligence will be deployed across the defense enterprise, from senior leaders down to the tactical edge (National Academies of Sciences, Engineering, and Medicine [NASEM], 2021). Unfortunately, the record of accomplishment for humans and artificial intelligence-enabled systems operating alongside each other leaves something to be desired.

When the USS Vincennes accidentally engaged an airliner, Iran Air Flight 655, in 1988, the one constant in the situation was that its Aegis combat system correctly tracked the airliner's kinematics throughout the seven minutes from liftoff to shootdown (Dotterway, 1992). Undoubtedly, many factors contributed to the incident with the USS Vincennes; however careful study demonstrates that the ad-hoc team structure in the combat information center set conditions for information overload in critical crew positions. The correct information was ultimately disregarded because of the Aegis contact classification limitations and the way that the human crew interacted with the system. Subsequently, this led to target misidentification and a tragic engagement (Dotterway, 1992). This situation also points to the trust that humans have in artificial intelligence. There must always be an override function for the decision maker as AI will never replace human intuition. As a result, establishing parameters is critical to maximizing the benefits of human-machine teaming.

Patriot missile batteries demonstrated that they were just as dangerous to friendly aircraft in the 2003 Second Gulf War (Bode & Watts, 2021). Using radars to track objects and algorithms to identify them, systems usually provide their outputs to operators via visual display screens. Although, perception of any situation comes from the fusion of information collected by a system and a human's understanding of the environment,

humans in this environment limited the scope of their perception to the Patriot's sensors (Kennedy, 2021). As a situation unfolds, the subset of information presented to a person helps them to refine observations, orient on the problem at hand, and then hopefully, to decide and act (Kennedy, 2021). Interestingly, track classification reliability challenges with Patriots were well known prior to the Second Gulf War. Yet, this limitation was not compensated for until much later (Bode & Watts, 2021).

Because the training data set lacked specificity to prevent false identifications, the system fell back on what it knew—in essence, everything in the sky is some kind of a missile and therefore, a threat (Bode & Watts, 2021). Operators need to remember that artificial intelligence is brittle. Learning occurs based on a set of training data, and systems interpret the world solely with an algorithm that is updated with that data. When data lies outside of the training data, the system is prone to fail because it cannot generalize what it knows with what it sees (National Academies of Sciences, Engineering, and Medicine [NASEM], 2021). A similar problem exists with biased data, like that previously used in Patriots. If everything the system learns is biased towards missiles, it should not be a surprise when the system behaves in biased ways (Han & Choi, 2021). This is where the human usually comes in and helps the machine to understand the exceptions that it sees. If the human is in a time critical situation, and their information is limited to a display screen, it is hardly surprising that they will orient incorrectly on the situation. This happened with Patriot missiles in 2003. Consequently, three instances of friendly fire resulted—two coalition attack craft were downed, and one friendly attack craft purposely fired on a Patriot that had targeted it (Bode & Watts, 2021).

Such cases indicate that further attention to how DOD is placing artificial intelligence alongside humans is required. Because humans and machines work differently, the DOD must play to the strengths of each partner. Strategists and future operating concepts envision artificial intelligence supporting their human partners, but a key element is missing. Deciding, communicating, and acting at machine speeds is not possible when teams of humans and machines are cobbled together in ways that prevent each partner from playing to their strengths. Deliberately designing systems to support the different use cases and teams that rely on them will be essential.

3. Different Flavors of Artificial Intelligence

One of the best attributes of artificial intelligence is also one of its biggest challenges. One of the goals of machine learning is to create a learning agent that can handle learning and adaptation happen on its own; without continuous human oversight (Domingos, 2015; MIT Horizon, 2022d). Because an AI system learns and is iteratively growing in capabilities, humans tend to think of artificial intelligence as a future system (NASEM, 2021; MIT Horizon, 2022b). When it is constantly learning and growing, understanding how it arrives at its recommendations will be a challenge. Error boundaries could become opaque, and human partners could lose awareness of where the system is operating outside its capabilities (Bansal, Nushi, Kamar, Lasecki, Weld, & Horvitz, 2019). People come to implicitly trust the system, or personify it, and believe it to be more capable than it actually is (Singer, 2009). Some drivers figured this out the hard way when their Tesla’s autopilot failed in catastrophic ways (Burke, 2022). Many think of AI as magic, a black box that can solve any task, however they are machines that have specific use cases and limitations.

Generally speaking, artificial intelligence is given a goal, they work towards the completion of that goal, and they learn the preferences of users over time (Russell, 2019). Artificial intelligence agents in use today are akin to Apple’s Siri, Amazon’s Alexa, or iRobot’s Roomba. Named “narrow-artificial intelligence,” these systems have a fixed use case. Generalizability is absent—performing tasks beyond the limited confines of its programming ends poorly. For example, Google’s AlphaZero easily defeats experts in games of StarCraft II because of its deep learning algorithm—deep learning takes large volumes of inputs and weights them across different “neurons” to learn connections between inputs—but that does not mean it can be dropped into the next service-level wargame and wow the participants with exceptional performance (Brose, 2020; Massachusetts Institute of Technology [MIT] Horizon, 2022). The all-in-one machine that can beat you at chess, drive the kids to school, and then do your taxes is called “general artificial intelligence” (Russell, 2019). General AI is a machine that can do many things equally well, including the exercise of creative, human-like thought (Russell, 2019). This type of artificial intelligence is likely decades away (NASEM, 2021). The distinction

matters when we consider how artificial intelligence will support units in the field. Intelligent agents have unique strengths, and are designed to do specific tasks.

4. Teams Solve Problems

The Department of the Navy (DON) is aggressively pursuing human machine teaming because it believes that integrating AI-enabled systems into operational/tactical teams will improve their performance (Department of the Navy [DON], 2021). However, simply placing an artificial intelligence system into a team of humans will not ensure success. Team structure and behaviors of its members matters. It is likely that the injection of a machine partner into a team will be different than adding a human team member (NASEM, 2021). For example, team trust, communications, coordination, and interaction dynamics could be subtly changed and lead to challenges in team situational awareness, agility, and performance.

Temporary in nature, teams are used to solve problems for organizations (Macmillan, Entin, & Serfaty, 2004). They must acquire, process, and act upon information, such as environmental factors, the task at hand, the limitations of their team, and the status of the team's goal (Macmillan et al., 2004). Members usually have unique roles and a broad range of skills (Cooke, Gorman, & Rowe, 2004). Winning happens when we align roles to the strengths of teammates. Team members can rely on each other because they know the strengths of teammates, and know who to ask for backup (Marks, Burke, Sabella, & Zaccaro, 2002; Smith-Jentsch, Kraiger, Cannon-Bowers, & Salas, 2009). Displaying patterns of communication, coordination, as well as interaction, a team's ability to use these and build an accurate mental model dictates their success and performance (Marks et al., 2002; Cooke et al., 2004). Unsurprisingly, there is a considerable amount of work that goes into the design of high-performing teams before we consider the integration of artificial intelligence. Much of these analyses and assessments are done intuitively by humans, however, what is intuitive to a human may not be easily programmed into a machine partner (Russell, 2019).

5. Centaurs as a Model

Centaurs are teams of machines and people that have partnered in chess and are an example that DOD could learn from (Epstein, 2020). Initially, the machine begins by sensing the positions of pieces on the board. Comparing this data to patterns it was taught during deep or supervised learning sessions, the machine has been trained to recognize, infer conclusions based on the board, and then provide recommendations (Epstein, 2020). The human evaluates these recommendations based on context and then acts. These centaur teams pair the machine learning, pattern recognition, and recommendation engines of an AI with the human's strengths (Stumborg, Brauner, Hughes, Kneapler, Leaver, Patel, Shields, & Collins., 2019). The human provides context, strategy, adapts to unexpected moves, makes inferences about the opponent's state of mind, and directs the action. Centaurs are so effective that even top-ranked opponents are handily defeated (Epstein, 2020).

Although chess games are not military operations, the principles that help centaurs win are still applicable. Unlike military environments, games such as chess have rigid rule sets with a fixed set of possibilities (Aycock & Glenney, 2021). Along with the player, the AI has a unique role in the team. Observing the environment, the AI provides piece movement recommendations that are tailored to the player's current situation. Humans orient on the new information, refine their strategy, and then decide on a series of winning moves. Both members of the team have a shared understanding of the board, the series of moves to deploy, and a common goal: winning the game. Players understand the boundaries where the system is most effective, and that its training data exposed the machine to countless, unbiased chess games. Even the communications between the player and their machine are optimized. The player updates the machine with real-time data and the machine can explain its recommendations to users in ways that avoid overload.

Machines have the ability to execute tasks much more rapidly than human counterparts, automate the processing of data, and operate in areas unfit for human partners (NASEM, 2021). Extrapolated to military use cases, this means that machine partners may be able to quickly assess intelligence, surveillance, and reconnaissance (ISR) data for target criteria or outliers, use these outputs to improve the quality of decisions or the targeting

cycle, and strike in areas where humans cannot safely operate (NASEM, 2021). Machines' extraordinary computational power allows them to quickly analyze data, extract patterns, predict outcomes, and convey them to a human decision maker. Such analyses could be helpful for decision makers at all levels of the enterprise. Machines have also proved that they can make tactical-level decisions that were previously unseen by human opponents (Brose, 2020; Scharre, 2018). This was demonstrated during games of go, where the machine was making moves in obscure parts of the board, areas that were seemingly unimportant—unimportant that is, until they trapped the human player many moves later (Brose, 2020). Although great at tactics, machines are less capable than humans when it comes to grand strategy and creative thinking (Epstein, 2020; Russell, 2019). Adding the tactical strengths of machines with human strategic thinking presents a potential competitive advantage over potential adversaries.

6. We Need Models Grounded in Team Performance

Warfighters expect their teammates to understand their role, how their tasks contribute to the whole, and where they may back them up when perplex conditions emerge. Centaur teams demonstrate that this is possible with AI with the caveat that they must be designed correctly. Truthfully, artificial intelligence cannot do everything that a team needs to do—they do not have to. What if they could be used in a specific narrow role, such as an agent that could regulate the communications of a team, or one that could identify optimum patrol routes for its team? What if more than one system could support the team with unique roles assigned to each? Team designs should account for the inherent strengths of the machine partner. AI algorithms excel at optimization problems, and the systems exhibit goal-oriented behavior (Russell, 2019). This machine teammate could set team performance as its goal, and then take steps to help the team optimize its performance. If done correctly, the AI will assist humans in areas where they are weakest, helping them to perform at a higher level.

The first step in doing this is to develop a conceptual model that captures the desired behaviors of team members. Supplying models of behavior to designers, they may design machine partners that will help manned teams accomplish their mission without self-

induced friction or fatalities. Unsurprisingly, the Russian Federation and People’s Republic of China are pursuing competitive advantage over the U.S. military through the development and weaponization of artificial intelligence systems (Tangredi, 2021). This is more incentive to get it right the first time. Afterall, combat against peer adversaries is hard enough without Siri getting in the way.

This research will focus on identifying and understanding the communication dynamics in a human-machine team. Each interaction between the machine and its teammates conveys different meanings (Canan & Demir, 2021). It is imperative that this interaction is clear and understood by each teammate. In critical situations, too much interaction can be just as detrimental as too little and can cause information overload. Lastly, hyper specialization has at times limited decision making by confirming human biases in certain situations. When an analogous situation arises, the decision maker may erroneously rely on heuristics and apply previous, inapplicable experiences (Kahneman, 2011; Kennedy, 2021). Machine partners can help avoid this trap by presenting possibilities that would otherwise go unnoticed (Epstein, 2020). It is worth noting that machine partners could be trusted to the point of overreliance; at this point an overreliance on AI may make decision points opaque, causing humans to forego making an important decision (Liu et al., 2021; NASEM, 2021). Therefore, the aim of this study will be to pursue the development of a conceptual model that will account for the human factors and team dynamics that could enable an autonomous system to partner with humans in teams more effectively. As such, this study will emphasize the team performance dimensions of human machine teaming rather than computer science and software components of such a team.

B. RESEARCH QUESTIONS, METHODOLOGY, AND PATH

1. Research Questions

This research seeks to answer the following questions:

- What does it mean to be a teammate?
- How can machines be partners with humans in teams?

- How can the integration of machine partners into previously human-human teams improve the team's agility?
- Can machine partners be better understood with shared team mental models or interactive team cognition approaches?

2. Methodology

An exhaustive review of the literature will be conducted, and two applicable case studies will be analyzed. The goal of this study is to produce a conceptual model of human machine teaming and provide contextual, real-world knowledge about communications within human-machine teams. Given the limited number of lab-based experiments on human machine teaming currently available, this study will examine the data to identify broad themes and patterns.

The nature of human teams and of human team dynamics has been extensively studied. This area of literature is rich with findings that can provide details about human-human team dynamics; however, little has been written about the integration of machines into human-human teams. With the inclusion of machine partners into traditionally human teams comes the need for research into human-machine teams. This study will first characterize artificial intelligence and its projected impacts on warfighting. An analysis of the existing literature regarding machine-machine teams will be provided, followed by the more developed subjects of human cognition and human-human teams. This study will then characterize human-machine teams before describing the communications, coordination, and interaction dynamics of human-machine teams. The authors will then demonstrate how these dynamics may be linked to the concepts of team agility and performance.

C. PURPOSE OF RESEARCH

The purpose of this study is to explore communications, coordination, and interaction dynamics in human machine teams and explicate their potential impact on team agility and performance. Such exploration is necessary to develop an understanding of team dynamics as human machine team constructs become more prevalent in DON. This study

will produce conceptual models of human machine teaming that may inform the design of future systems. The results of this study could help the DON better understand the implications of integrating narrow-artificial intelligence capabilities into team constructs. Such knowledge would ultimately allow the DON to apply research findings to improve agility in human-machine teams.

THIS PAGE INTENTIONALLY LEFT BLANK

II. ANALYSIS OF THE DEPARTMENT OF THE NAVY ARTIFICIAL INTELLIGENCE STRATEGY

Artificial intelligence (AI) is forecasted to enhance the American way of life by stimulating economic growth as well as dramatically improving national security (The White House, 2019). It is believed that AI will affect all aspects of government, and could allow for significant advantages over the nation’s competitors (The White House, 2019). As a part of this effort, the Department of Defense (DOD) is pursuing AI capabilities as a means to enable the technological advantages needed to deter conflicts and win wars (Department of Defense , 2018). The DOD’s AI strategy asserts that DOD personnel are the department’s center of gravity, and are the primary target of AI systems (DOD, 2018). More specifically, the DOD aims to employ AI to empower its human assets (DOD, 2018). This interest is due to the assessment that combinations of humans and machines in human-machine team (HMT) constructs will be the “killer app” that will grant advantages over competitors (Singer, 2009; Lewis & Vavrichek, 2019; Scharre, 2021).

The Department of the Navy (DON) developed its own strategies and frameworks to deliver on this concept. The department supplemented the DOD’s overarching strategy with definitions, taxonomies and standards from the president’s Executive Order 13859 and the National Institute for Standards and Technology’s (NIST) U.S. Leadership in AI: A Plan for Federal Engagement in Developing Technical Standards and Tools (Lewis & Vavrichek, 2019). This study will analyze the DON’s AI strategy, and will limit its scope to the human machine teaming (HMT) line of effort (LoE).

A. DON’S PURPOSE

The United States Navy seeks the means to leap ahead of its competitors. The stated purpose of the DON’s strategy is to create a guiding framework that could be used to align with the DOD’s efforts to accelerate the adoption of AI (DOD, 2018). This adoption is envisioned to support and enhance the DON’s ability to conduct distributed maritime operations by leveraging human-machine partners (Office of the Chief of Naval Operations, 2020). It is believed that AI will enable new, revolutionary tactics and

capabilities, while augmenting and enhancing the department's legacy practices (DON, 2021a). These innovations are hypothesized to make the department more effective, efficient, and competitive in an era characterized by great power competition. More specifically, this strategy seeks to align the disparate use cases, resources, and goals of HMT employment within the DON so that the department may achieve a competitive advantage over adversaries in all domains and operational scenarios (DON, 2021a; Work, 2021).

B. DON'S PROBLEM SPACE

The DON's AI strategies are coherently integrated and possess common LoEs traceable to the *National Security Strategy* (2018). These strategies point back to one observation: the U.S. is not prepared to defend or compete against great power competitors in the era of AI (DON, 2021b). At a high level, the DOD is seeking AI-enabled improvements in all aspects of the enterprise. Thus, the problem space for DON's HMT initiative comprises all aspects of the enterprise from the tactical to strategic levels. There is a desire to leverage AI to solve innumerable departmental challenges and augment human partners (DOD, 2018; DON, 2021b). The strategy requires linkages between the enterprise's technological and process infrastructures, commercial and coalition partners, the DON workforce, intelligence, surveillance, and reconnaissance (ISR), information operations, logistics, as well as command and control (C2) functions (Legaspi, Mah, Hsieh, 2021; Tangredi, 2021). The integration and alignment of these areas, stakeholders, and technologies is argued to be essential to augment the performance of the DON's human capital with the superior analytical capabilities of machine partners (Office of the Chief of Naval Operations, 2020).

Even when narrowing down the scope of the strategy to focus on HMT, this represents an enormous and complex problem space. To achieve meaningful adoption of HMT as it is envisioned by the DON, strategy, technology, all aspects of the environment, as well as organizational processes, structure, and metrics will need to be carefully considered (Bergin, 2021). This scrutiny must include the key variables needed to enable these areas.

C. KEY OPERATIONAL VARIABLES AND A TECHNICAL FRAMEWORK

Deploying HMT in the use cases envisioned by the DON’s strategies requires a suitable framework to guide implementation. To this end, the DON’s intelligent autonomous system framework appears to provide a suitable foundation for the implementation of HMT (DON, 2021a). It will organize investments based on portfolios of capabilities, identify strategy gaps at high levels, and enables the sensible allocation of resources by resource sponsors (DON, 2021a). The framework is comprised of 3 broad categories that contain 7 top program elements, and decomposes many implementation variables into 4 levels. Although it sounds quite complicated, the graphical representation shown in Figure 1 summarizes the discussion comprehensively.

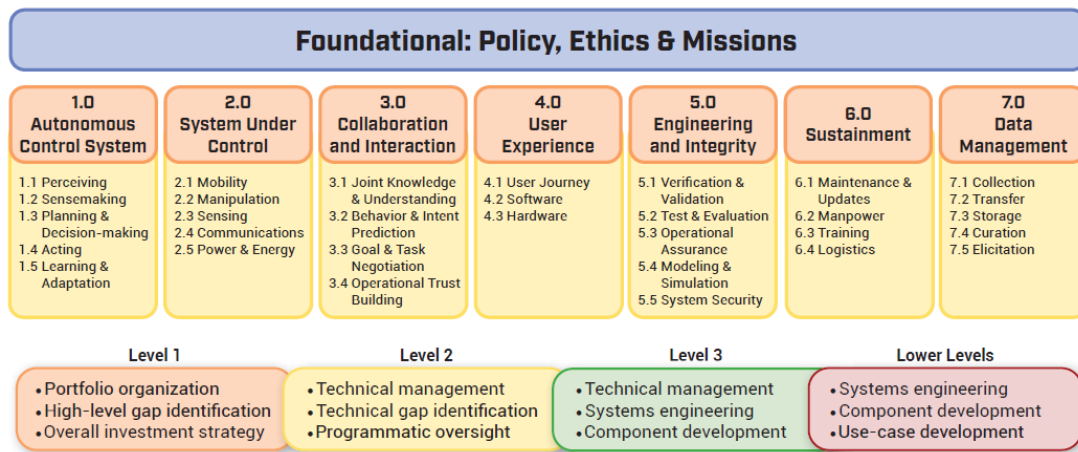


Figure 1. Intelligent Autonomous System Technical Framework. Source: DOD (2021a).

The key variables underpinning this framework were identified by the National Institute of Standards and Technology (NIST), in its attempt to develop and prioritize AI-related standards. These variables, named as standards by NIST, include concepts and terminology, data and knowledge, human interactions, metrics, networking, performance testing and reporting methodologies, safety, risk management, and trustworthiness (NIST, 2019). Trustworthiness is further decomposed into accuracy, explainability, resiliency, safety, reliability, objectivity, and security (NIST, 2019). Although these variables are

listed in the requirements documentation, there are many gaps in the definitions, as will be demonstrated later in this chapter. As complexity in the environment increases, these design variables become increasingly important to enable the performance of HMTs. DON intends for these critical design variables to enable a spectrum of interdependent tasks for humans and machines (DON, 2021b). This spectrum can be modelled as a “2x2 grid,” which demonstrates how these variables align to suit the inherent strengths of humans and machines. Figure 2 illustrates the relative strengths of the team members, given different operational environments.

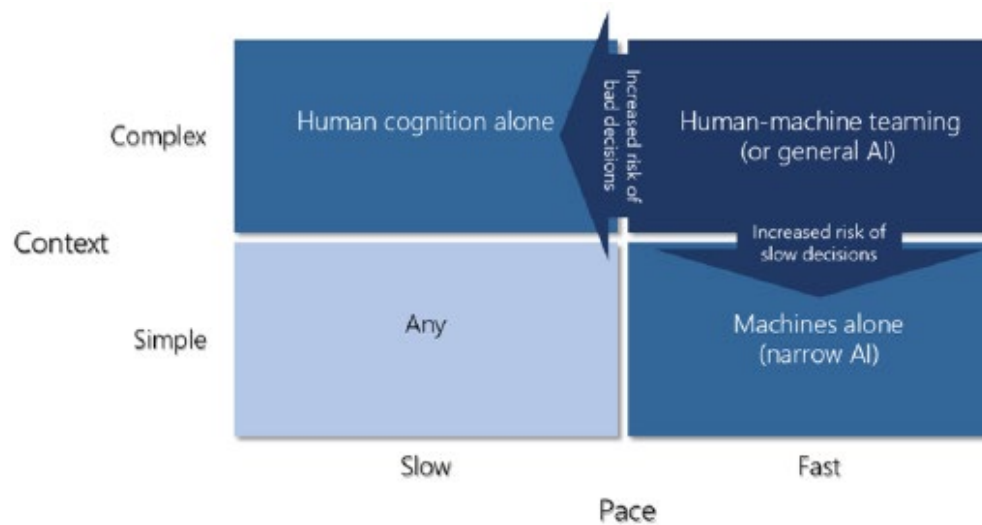


Figure 2. Artificial Intelligence Applications and Operational Contexts.
Source: Lewis and Vavrichek (2019).

The variables become operationalized into situation and time-based considerations, and are also represented in the team member design principles. As an example, the ability of an AI to explain the rationale behind its recommendations is believed to be an important enabler for human decision making in complex situations that require rapid decision making (NIST, 2019; Bray & Moore, 2021; NASEM, 2021). Both members of the team must be able to understand the situation, data inputs, the AI’s outputs, and most importantly, the human must be able to trust the machine teammate (Chappell, 2020; Bray & Moore, 2021; NASEM, 2021). Adhering to an implementation framework and deriving

the technical variables is not necessarily enough when considering the problem of technological adoption in a large enterprise like the DON. Adoption at scale has organizational design considerations as well (Kotter, 1978).

D. TECHNOLOGY ADOPTION

The DOD is very good at developing technologies; however problems and impediments sometimes arise when services try to adopt these new ideas and capabilities (Scharre, 2021). At all levels there is a reluctance to iterate and to embrace the Silicon Valley mantra of “fail early, fail often, fail fast” (Brose, 2020). Changing the culture of an organization may prove challenging, and cannot be done in isolation of other factors (Kotter, 1978). Therefore, for DON to make any progress in implementing the strategies, the department must be willing to embrace change as well as the risk of failure (Brose, 2020; DON, 2021a). Technology is a means, not an end in itself. Success in the adoption of the technologies discussed in the strategy hinges on the DON’s ability to characterize its problems, align technology to meet these challenges, and then spur adoption based on core principles outlined in its framework (Lewis & Vavrichek, 2019).

Many of the legal, safety, and ethics policies in place to safeguard against AI will need to be reevaluated. Weapons reviews, privacy considerations, as well as validation and verification mechanisms will need to be overhauled to address decades old policies (Lewis & Vavrichek, 2019; Singer, 2009; Scharre, 2018). To achieve the requirements needed to support AI, the DOD will need to revamp its organizational and leadership structures (DOD, 2017). The DOD took the first steps to change the frameworks surrounding AI with its *Autonomy in Weapon Systems Directive* in 2017 (DOD, 2017). The purpose of the policy is to prevent an AI-enabled system to use force in an undesirable situation (DOD, 2017). The enforcement mechanism of the policy is dependent upon rigorous verification and validation test events (DOD, 2017). Finally, a key contribution of this document is that it dictates the types of force that the weapon systems defined in the document may deploy, and which target sets they may engage (DOD, 2017; Clarke & Knudson, 2018). The implications of the policy are that humans will be required to be involved in decisions

regarding how, when, why, and under what conditions a weapon will be deployed (DOD, 2017).

There is a need to rapidly prototype, learn from these prototypes and then field them (Scharre, 2021). This is consistently a struggle with current acquisitions policies and practices (Brose, 2020). Many of the test and evaluation processes, procedures, and policies for AI do not exist and will need to be invented, codified, optimized, and distributed for enterprise use (DOD, 2018). A number of the standards necessary for the successful adoption of AI are not yet mature enough to guide the organization. Some of the standards, such as trustworthiness require guidance for accuracy, explainability, resiliency, safety, reliability, objectivity, and security (DOD, 2018). These are terms that are easily articulated as concepts, but require significant work to codify them into measurable standards. NIST understood this problem, and has been charged with fleshing out the necessary components of the standards (NIST, 2019).

The DON's knowledge of AI-enabled technologies will need to be improved (Schramm & Clark, 2021). Leveraging AI to make the workforce more effective will require end-to-end redesigns of training curricula as well as information sharing policies (DOD, 2018; Lewis & Vavrichek, 2019). Organizational business processes and capabilities will need to be redesigned. When it comes to partnerships and systems design, there will also be a requirement for common data types, structures, and complementary capabilities (Schramm & Clark, 2021).

E. WHY IT IS IMPORTANT? WHAT IS THE IMPACT OF NOT DOING IT?

As a nation, the failure to embrace AI-enabled capabilities will reduce access to global markets, spurring declines in economic prosperity, as well as lead to inevitable breakdowns in privacy (DOD, 2018). Arguably, competitor nations may not share the same outlook on the protection of privacy and civil liberties (Tangredi, 2021). DON may encounter operational challenges in future conflicts if it fails to study the ramifications of deploying humans and AI in close proximity. The ordered effects of such failure have already manifested themselves. High profile examples within the DON where human

judgement was unsuitably influenced by machine capabilities include the 2003 shutdown of an F/A-18 by friendly Patriot missile batteries and the 1988 shutdown of an Iranian airliner by the USS Vincennes (Bode & Watts, 2021; Scharre, 2018; Lewis & Vavrichek, 2019; Dotterway, 1992). The DON cannot afford to ignore the technological potential of integrating AI in manned formations; as its adversaries are certain to capitalize on this capability (Tangredi, 2021).

F. AI AND GREAT POWER COMPETITORS

The strategy indicates that competitors such as Russia and China are making significant military-related investments in AI (DOD, 2018; Tangredi, 2021). DOD asserts that the nation's competitors are likely to use these investments to erode American military dominance, and threaten global security, stability, and privacy norms (DOD, 2018). Both countries are already using AI in warfighting applications, intelligence, logistics, cyberspace operations, and in a variety of autonomous platforms (Hoadley & Lucas, 2020; Center for Naval Analyses [CNA], 2021a; CNA, 2021b; Tangredi, 2021). Although the U.S. uses AI in similar fashions, its acquisition process presents a significant hurdle if not updated (Brose, 2020). The pressure from competitors has spurred changes, but not quickly enough to keep pace with the changing AI environment.

G. IMPLEMENTATION AND STRATEGY GAPS

While the potential capabilities of AI are many, the DON is currently facing multiple gaps that cannot be bridged without a solid framework. The primary challenge is trust. On the surface, this issue may seem simple. However, when AI becomes an autonomous partner in an HMT the issue becomes even more complicated. Machine learning, which is at the core of some AI algorithms, is different than human learning (Howard, 2020). Machine learning consumes massive amounts of data to detect patterns whereas human learning often relies on inferential reasoning (Howard, 2020). The aforementioned information may seem insignificant on its own, however, this is ultimately the shaky bedrock on which trust in AI stands (United States Marine Corps, 2018). This gap is broadened even more in the military where trust is essential, and life and death decisions are made daily in combat (HQ USMC, 1997).

There are gaps in the kill chain in the form of whether or not to allow AI to engage human targets (Scharre, 2018). Recent experiments have exposed these gaps with disastrous consequences. During one exercise, a machine was tricked into identifying a school bus as an enemy tank and proceeded to annihilate it, simply because several pixels were spoofed (Howard, 2020). In this scenario, the machine was given broad engagement authorities without a human in the loop. Had a human decision maker been in the loop, then perhaps the mistake would not have been made. This example highlights how adversaries could easily expose these dynamics.

Perhaps the widest gap yet to be bridged is the gap in developing standards. Standards are crucial to the employment of AI in the military. As of 2019, more than 53 AI standards existed, however, as Jane Pinelis the Chief of Testing and Evaluation for the Joint Artificial Intelligence Center (JAIC) pointed out in 2021, “standards are elevated best practices and we don’t necessarily have best practices yet” (Waterman, 2021, p. 26; NIST, 2019). This statement indicates misalignments in DON’s AI strategy. As can be seen from Figure 3, it is not clear how AI standards should be created, finalized, and employed much less which ones to prioritize.

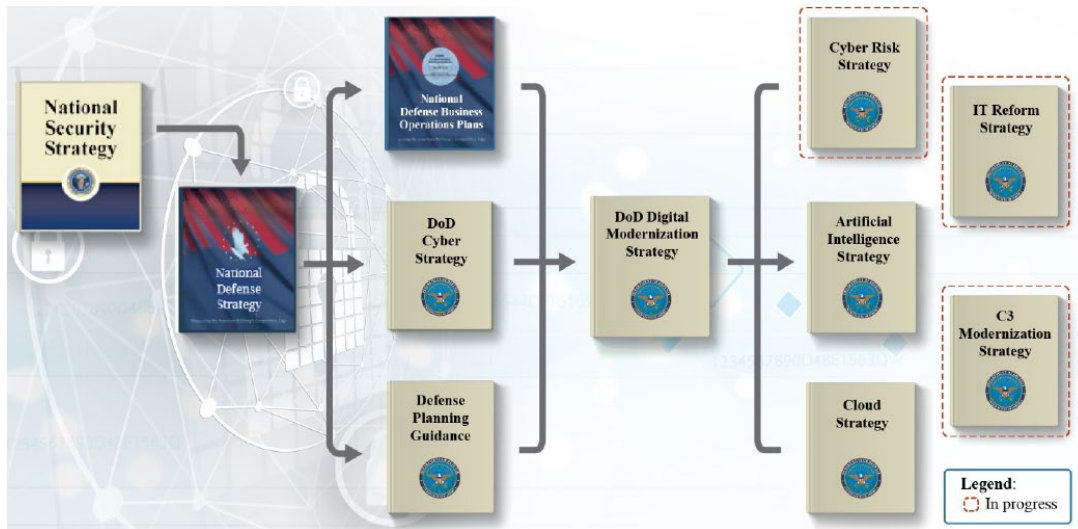


Figure 3. United States Artificial Intelligence-Related Strategies. Source: NIST (2019).

The main focus of the executive branch’s plan is to “ensure that technical standards...reflect Federal priorities for innovation, public trust, and public confidence in systems that use AI technologies...and develop international standards to promote and protect those priorities” (NIST, 2019, p. 3). It identifies nine areas in which AI standards should focus: concepts and terminology, data and knowledge, human interactions, metrics, networking, performance testing and reporting methodology, safety, risk management, and trustworthiness (NIST, 2019, p. 3). There were many linked standards comprising the first eight, however, trustworthiness standards were previously overlooked.

H. IMPLEMENTATION PLAN

The concept of operations for the AI strategy identifies needs for technological and organizational changes (DON, 2021a). The strategy calls for re-architecting cloud capabilities, data collections, as well as information sharing approaches (Lewis & Vavrichek, 2019). Many of the key organizational changes, technologies, and partnerships required to enable this strategy do not exist, or are mentioned as future initiatives that must be explored. Lewis and Vavrichek (2019) assert that successful implementation of the AI technologies in the DON will require leveraging technical expertise across the enterprise, understanding the cognitive processes of humans and AI partners, realistic test and evaluation events, as well as supporting personnel with frequent, yet substantial training exercises. An often overlooked, yet crucial mediator in the successful adoption and implementation for technologies such as these is education (Denning & Dunham, 2010). User education must be continuous, and its outcomes must be made relevant to the user (Harwick & Barki, 1994).

The strategy has several pillars and an appointed implementation mechanism. It prioritizes the fielding of AI systems that offload tedious cognitive and physical tasks. Eventual goals are to incorporate rapid, innovative, and responsible development of AI capabilities to support decision making and key military operations. The DON aims to leverage innovative capabilities developed by the tactical edge, but realizes that to do so it needs to put into place foundational building blocks, platforms, data repositories, tools, frameworks, and infrastructures to support edge AI practitioners (DOD, 2018). The

strategy relies on a change in DON's culture, skills, and employment approaches. It calls for a workforce that is trained to employ AI, capable of rapid experimentation as well as incremental development, and is risk informed (DOD, 2018).

The next pillar of the strategy is tangential. It calls for engaging commercial, academic, and international partners to realize global leadership in ethical, open-source AI capabilities (DOD, 2018). There is also a financial component in this LoE. It specifically acknowledges the need to entice partners with funding and long-term research opportunities (DOD, 2018).

The final LoE in the strategy is designed to support ethical and safety considerations. The DOD desires to use AI capabilities in a lawful, ethical manner (DOD, 2018). A tenet of this effort is robust, secure, and explainable AI as well as to pioneer test and evaluation approaches for these capabilities (DOD, 2018). In addition to operating the AI safely, there is a subordinate aspect of the strategy that calls for using AI to improve the qualities of decisions to reduce instances of unintended collateral damage (DOD, 2018).

The JAIC is intended to be the strategy's primary implementation mechanism. The JAIC is charged with creating feedback loops between AI research and warfighters' operational challenges, establishing the common infrastructure and foundation needed to leverage AI across the defense enterprise, facilitating AI planning, governance, ethics, and security considerations, and to grow AI expertise (DOD, 2018). This action arm will be essential to the successful adoption of AI in DON.

I. ANTICIPATED BENEFITS

The projected benefits are reducing risk to mission and risk to forces, improving naval decision making, and improving efficiency (Lewis & Vavrichek, 2019). Another forecasted benefit will be the ability to employ revolutionary tactics that will prove highly disruptive to the adversary (Lewis and Vavrichek, 2019). Better thread identification, cybersecurity, and overall readiness increases are all positive benefits on the battlefield. Advantages will be realized outside of combat as well. Recruiting will be more tailored to individuals, data analysis will be improved, and staffing requirements will be lessened

(Lewis & Vavrichek, 2019). As AI progresses and funding levels increase, unforeseen benefits will also emerge.

J. ORDERED EFFECTS

Misalignments in the DON's AI strategy, organizational strategies, and technology strategies may cause unanticipated effects that will reverberate throughout the enterprise (Bergin, 2021). The decision-making processes of AI capabilities are not well understood, and empowering them to operate in complex environments could have far reaching impacts. Requiring AIs to interface with each other without a human in the loop could lead to unexpected outcomes that increase risks to mission and forces (Bray & Moore, 2021; Scharre, 2018). As an example, an AI may flood the environment with requests and orders, artificially reducing the available quantity of supplies. Humans may turn to these same capabilities to introduce certainty and precision into decision making processes. This could lead to problems when the human does not fully understand the reasoning, programming, and logic that underpins their machine partner (Chappell, 2020; NASEM 2021). This introduces conditions where the human is not cognizant of the many factors of the decision space, and assumes that the machine partner's model is sufficient to support decisions. Such failures to incorporate important data increase the risks of making a poor decision (Chappell, 2020).

Aside from the challenges of automation bias, there is a very real possibility of skill atrophy that comes with reliance on technology to perform tasks (Chappell, 2020; NASEM, 2021). Machine partners will be designed to handle dirty, dangerous, tedious, or time-consuming tasks for human partners (Grosz et al., 2016). Depending on the scope of the task, humans may find themselves no longer capable of performing some tasks up to combat standards (NASEM, 2021). This is not as far-fetched as one would believe; after all, how many members of the DON can navigate via map and compass as opposed to being wholly dependent on capabilities such as the global positioning system?

Humans may experience other challenges when operating alongside machine partners. When conditions are optimal and an automated system is functioning as intended, humans tend to become distracted and fail to closely scrutinize the actions of automated

agents (NASEM, 2021). As a result, humans must devote more time to situational awareness building when presented with novel conditions (NASEM, 2021). The human partner's performance is negatively impacted when there are sudden increases of work and sensemaking that must be done (NASEM, 2021). In one such example, human drivers of Tesla vehicles repeatedly demonstrated inattentiveness while operating their vehicles (Morando, Gerhson, Mehler, & Reimer, 2021). Distracted drivers and a misunderstanding of the capabilities of the automated auto-pilot system were attributed to several high-profile vehicle accidents (Morando, Gerhson, Mehler, & Reimer, 2021).

Finally, failing to align the standards and variables that will be used to design AI with those of enabling technologies could create process breakdowns, forcing users to implement workarounds to complete their assigned tasks (Denning & Dunham, 2010; NIST, 2019). As an example, the 2013 vehicle wireless networking standard is used for the design of autonomous vehicles (NIST, 2019). This standard is nearly 10 years old, and identifies Wi-Fi requirements as opposed to autonomous vehicle communications requirements (NIST, 2019). Standards related to data formats, testing methodology, transfer protocols, cybersecurity, and privacy are examples of key areas where revisions are likely needed to support the use of AI in HMT (NIST, 2019, p. 10).

K. CHAPTER SUMMARY

The U.S. is looking to gain a competitive advantage on the battlefield. The DON AI strategy aims to be a vehicle to provide that advantage. The strategy and its technical framework align DON's overarching strategy with its organizational and technology strategies. It is theorized that by adopting AI-enabled HMT, casualties and collateral damage will be reduced, more informed and accurate decisions will be made, and national security will be improved (Stumborg et al., 2019). The strategy aims to make the desired advantages a reality but still leaves many questions to be answered. The DON has the tools, resources, and human capital to make the envisioned improvements but must provide achievable standards. This is essential given adversary investments in AI-based capabilities. In conjunction with AI, humans will complete the structure that is HMT and lead the way forward to a more secure future.

III. A REVIEW OF TEAM THEORY

The concept of teaming underpins all aspects of this study; as such it is important to understand the prominent theories and principles of teaming. Complex and difficult tasks, organizational dynamics, as well as the never-ending quest for improved performance drives organizations to form teams (Salas & Fiore, 2004). Teams allow organizations to leverage diverse expertise, skill sets, and additional resources to problems that may be too large for a single person to undertake (Espinosa, Lerch, and Kraut, 2004). While teams can bring many benefits to an organization, Salas and Fiore (2004) argue that organizations tend to focus solely on the benefits and that many teams are formed without due consideration to the factors that improve team performance. The nature and dynamics of human teams has been extensively studied. Military teams have even been described as a subset of action and performance teams due to their “distinguishing characteristics of skilled specialist roles, focused performance events, and improvisation due to the dynamics and unpredictable nature of tasks” (NASEM, 2021, p. 11).

Reams of literature provide details about human-human teams; however, little has been written about the integration of machines into human teams as a team member. With the inclusion of machine partners into human teams, the need for research into HMT becomes essential. This study will first describe human teaming constructs before reviewing HMT in chapter six.

A. SHARED MENTAL MODELS

Hinsz (2004) links the concepts of team agility, performance, and team dynamics to the concept of the mental model. Hinsz’s (2004) definition of a mental model is “... an individual’s mental representation and beliefs about a system, and an individual’s interactions with a system, with particular focus on how the individual’s interactions with the system leads to outcomes of interest” (p. 40). He further argued that mental models were generalizable to group constructs by replacing the representations of a system with those of a group (Hinsz, 2004). Van Ments & Treur (2021) go on to describe mental models as a type of blueprint or images of the mind in different forms. In his 2004 work, Hinsz

evaluated the sharedness of mental models, and how they related to the functioning of a team. A shared mental model can be defined as the shared representation of how a group's team members understand each other's work, responsibilities, the situation that the team is in, and how they perceive the problem structure of their ultimate goal (Hinsz, 2004; Fiore & Schooler, 2004; NASEM, 2021).

Shared mental models contribute to team situational awareness, and improve performance, especially in critical situations (Hinsz, 2004; Levine and Choi, 2004). For example, sports referee teams are composed of individuals who are responsible for different individual tasks in order for the referee team to be proficient (Sinval, Sinval, Maroco, Marques, Uitdewilligen, Maynard, & Passos, 2020). Prior to the start of a game or match, the referees coordinate with each other to determine who is responsible for each anticipated task (Sinval et al., 2020). This coordination helps them better facilitate the game as well as officiate with more confidence as each individual understands which tasks they are responsible for.

Hinsz (2004) argues that team performance is improved because members of the team have a similar understanding of a situation and will spend more time executing tasks as opposed to coordinating actions. This becomes possible as team members will be able to better explain and anticipate member actions, allowing them to manage dependencies and adapt to fill gaps more effectively (Espinosa et al., 2004; NASEM, 2021). Shared mental models also directly relate to the concept of meta-perception. Specifically, Hinsz (2004) demonstrated team performance improvements when team members understand the interactions between the group and its members. These interactions would ideally characterize which member of the group would complete what tasks, when, and under what conditions (Hinsz, 2004). This is the linkage between meta-perception, metacognition, and team shared mental models. These concepts integrate to enable teams to dynamically interact with each other to complete tasks.

B. INTERACTIVE TEAM COGNITION

A team's cognition does not originate at the team level, but is developed through the aggregation of each individual member of the team (Rentsch and Woehr, 2004). This

makes intuitive sense, especially in heterogeneous or distributed teams where information is processed through different lenses. Such processing is aggregated to form a meaningful understanding of its environment (Kennedy, 2021). Kennedy (2021) demonstrated that this group-level processing was the integration of many different projections, especially in military applications. This relates closely to the concept of metacognition, or how team members understand the ways that they process information (Hinsz, 2004).

Metacognition involves the group's tasks, and influences the ways that teams problem solve, incorporate information, make sense of their environments, as well as make valuations of team member expertise (Hinsz, 2004). Rentsch and Woehr (2004) assert that a team's effectiveness can be attributed to a level of shared cognition amongst its members. They further assert that shared cognition is improved by team-level training. The concept of metacognition also integrates with the concept of requesting backup from team members. Marks et al. (2002) related the valuations of team member expertise to the backup behaviors demonstrated by some teams. These valuations could be used to identify scenarios where backup is needed, and which types of assistance are required.

Teams act as a cognitive system when they are able to coordinate effectively, plan, think, and act as a team (Cooke & Gorman, 2009). A critical component of this is cognitive system is the knowledge of the inter-role responsibilities of a team's members (Marks et al., 2002). Team cognition may be found in the organization of information within team members when they attempt to accomplish a task. Team interactions use this information, including which roles should be engaged in which tasks at what times, to sequence the tasks needed to achieve a goal (Marks et al., 2002). The information and sequences of tasks allows team members to visualize each role's contributions towards the goal, identify any deficiencies, and engage in compensatory behavior (Marks et al., 2002; Smith-Jentsch et al., 2009). This interactive team cognition fuels coordination behaviors, which may improve team performance (Marks et al., 2002).

1. Team Member Schemas

Klimoski and Mohammed (1994) refer to the overlap in cognition between team members as schema similarity. They hypothesized that team member schema similarity

was linked to team performance. Ultimately, Rentsch & Klimoski (2001) related shared cognition to team outcomes. This makes intuitive sense, as similarities in team member schemas imply “compatible knowledge structures for organizing and understanding team-related phenomena” (Rentsch & Woehr, 2004, p. 15). Shared cognition would describe the degree to which the disparate schemas match, however, this does not speak to the accuracy of the team’s schema. Additionally, Rentsch & Woehr (2004) note that the accuracy of a schema refers to the schema of the of the assessed actor as it exists in reality.

Inaccuracies in schema could be mitigated by interacting with the environment external to the actor, or through communications with partners (Canan & Demir, 2021). However, communicating with partners when there are differences in schema could adversely impact team effectiveness until these schemas achieve congruence (Rentsch & Woehr, 2004). There may also be a condition where requesting clarification or executing compensatory backup behaviors may adversely affect members of the team who cannot absorb additional information or tasks (Smith-Jentsch et al., 2009). This could ultimately detract from team performance as members would succumb to overload and be unable to execute the tasks prescribed to their roles.

Rentsch and Woehr (2004) assert that team member interactions will cause them to develop a sense of the capabilities of the teammates in which they interact. Attributes such as competence, reliability, communicativeness, and cooperativeness could be assessed, and are likely to affect the team’s interaction dynamics (Rentsch and Woehr, 2004). Mathieu et al. (2010) determined that team members are less likely to rely upon each other when their partners possess incongruent attribute scores. This in turn could negatively impact team interactions and performance. These results are similar to other studies that discussed team backup behaviors (Marks et al., 2002; Smith-Jentsch et al., 2009).

Rentsch and Woehr (2004) discuss the concept of meta-perception, which may prove to be relevant in human-machine teams. Meta-perception is Actor A’s perception of how Actor B perceives Actor A. It is important to understand how the machine partner perceives the actors on a team. Armed with such understanding, human team members will be better able to understand how the machine partner characterizes its environment. According to Rentsch and Woehr (2004), meta-perception may also include a team’s task,

processes, equipment, personnel, and operating environment. In an ideal situation, the team should possess both congruent and reciprocal schemas, which would enable effective team cognition as members spend more time performing tasks and less time building understanding (Rentsch & Woehr, 2004; Hinsz, 2004).

2. Teamwork and Taskwork

Cannon-Bowers et al. (1993) noted that shared cognition is a byproduct of team-related issues. Meanwhile, Mathieu et al. (2000) asserted that shared cognition stems from task analyses and a model comprised of eight critical task attributes. In addition, Cooke et al. (2004) argue that shared cognition is a byproduct of team knowledge and the dynamic interplay of team processes. Interestingly, Cooke et al. (2004) determined that even though teams may have significant inter-positional knowledge, these teams did not display significantly improved team performance. However, this cross-training was strongly correlated with improved teamwork and taskwork performance metrics (Cooke et al., 2004). Innately, this makes sense as cross training in different roles helps members of a team understand how to execute a sequence of events before they understand who does them (Cooke et al., 2004).

Narrow or hyper-specialization leads humans to rely on past experiences and familiar patterns which have the potential to result in catastrophic consequences. Argyris (1976), who helped establish the Yale School of Management, identified this phenomenon as single-loop learning. Single-loop learning is when decision makers tend to go with the first common answer that presents itself (Argyris, 1976). For example, a 2015 study found that cardiac patients had a higher survival rate if they were admitted to a hospital when thousands of cardiac doctors were away at a national cardiac convention (Jena, Prasad, Goldman, & Romley, 2015). Many researchers attribute this to a reduction in common treatments, such as a stent, while the specialists were away (Jena et al., 2015).

Another phenomenon called “cognitive entrenchment” was identified by Erik Dane who is an organizational behavior professor at Rice University. Dane (2010) described it as when experts have a challenging time adapting to new rules as opposed to non-experts. For example, experienced accountants were asked to find deductions using a new tax law.

Following the study, the experienced accountants did significantly worse than new accountants (Dane, 2010). This example presents a dire warning for our current military team leaders. By adding a machine teammate unencumbered by traditional military experiences, a unique perspective is offered.

Cooke et al. (2004) determined that teamwork and taskwork are two measurable dimensions of team cognition. It was noted that taskwork knowledge usually preceded teamwork knowledge, and that teams whose members generally understood tasks from multiple perspectives tend to perform better than those who do not (Cooke et al., 2004). Team situational awareness was described by Cooke et al. (2004) to be another measurable attribute that contributes to team performance.

C. TEAM SITUATIONAL AWARENESS

In the quest to understand teams, one must also consider the concept of awareness. Awareness can be described as the knowledge that is “created through interaction between an agent and its environment” (Gutwin & Greenberg, 2004, p. 181). Along with this definition, it is assumed that awareness exists about the current state of an environment, which may change over time (Gutwin & Greenberg, 2004). Reconciling these changes requires some form of interaction between agents and the environment, and as a part of this, maintaining awareness is usually required to support some prescribed goal (Gutwin & Greenberg, 2004). With this definition provided, the linkages between concepts of awareness and teams may be explored.

It has been established that shared knowledge and interaction equates to higher team effectiveness (Cooke et al., 2013). Underneath team effectiveness is the idea of team situational awareness. Team situational awareness is the degree of cognizance each team member holds in order to accomplish their job (McNeese et al., 2021). In totality, team situational awareness becomes greater than the sum of its parts. This definition suggests that each team member brings their own piece of cognitive knowledge and when added to the knowledge of other teammates becomes team situational awareness (Mohammed et al., 2010). Gorman et al. (2006), pointed out that team situational awareness is directly tied to the interactions that take place as teams assess and make decisions. This viewpoint suggests

that team situational awareness is more than simply cognition. It also involves the coordination of communication involving the right people at the right time. Cooke et al. (2004) asserted that as team members interact through team process behaviors such as communication, coordination, and interaction, they transform individual knowledge points into a collective team knowledge base that is used to guide the team's actions.

Gutwin and Greenberg (2004) assessed that effective teams were able to maintain an awareness of how workspaces were shaped by team member interactions and noted that technological limitations may detract from team performance. More specifically, Gutwin and Greenberg (2004) discussed how technologies may not support shared workspace awareness, especially for distributed teams. When systems lack visual or other cues to indicate team and workspace interactions in a shared space, it is unsurprising that additional coordination mechanisms would be necessary in order to build awareness of objects, activities, or to anticipate team member actions (Gutwin & Greenberg, 2004). This implies that team members and system designers should be aware of the information that must be captured, communicated, and presented to others (Gutwin & Greenberg, 2004).

D. COMMUNICATIONS, COORDINATION, AND INTERACTION DYNAMICS IN TEAM CONSTRUCTS

It is not enough to bring machine partners into a team traditionally comprised of humans. The communications, coordination, and interaction dynamics present in a team contribute towards its effectiveness and performance. HMT must have a robust and resilient mechanism for control and synchronization to be effective in critical situations like those found in military operations (Fout & Ploski, 2018). The structure of the team has also been demonstrated to have effects on its communication and coordination requirements (Macmillan, Entin, & Serfaty, 2004). The team's structure has to account for the division of tasks, responsibilities, and resources as well as its leadership hierarchy (Macmillan et al., 2004). The environments that the teams operate in may be dynamic and are likely to be characterized with uncertainties in time, resources, as well as guidance and authoritative direction (Singer, 2009). These may all contribute to team dynamics. Specifically, research has proven that team dynamics in critical situations may be influenced by the operational environment, understanding of taskwork, and an

understanding of the team's physical and cognitive resource limitations (Mohammed, Rico, & Alipour, n.d.).

Macmillan et al. (2004) characterize teams as information processing nodes that must acquire, process, and act on information, such as environmental factors, taskwork, and team limitations, to achieve the team's goal. As described in the situational awareness section, teams must have a shared understanding of its agents, their environment, and their goal. This team situational awareness construct requires communications to occur in order to exchange information about the team's current tasks, structure, member responsibilities, the situation that the team is in, as well as its coordination requirements (Macmillan et al., 2004; Levine and Choi, 2004). Interestingly, communications and shared cognition were demonstrated by Levine and Choi (2004) to have effects on important team attributes including: trust, morale, satisfaction, and effectiveness.

Macmillan et al. (2004), introduce the concept of communications overhead as a hidden cost of team cognition, and relate team communications behaviors to team performance. Their argument was that the information exchanges, called communications, must be exchanged during the execution of tasks and that these consume cognitive or physical resources (Macmillan et al., 2004). This is related to Cooke et al.'s (2004) work about measuring team cognition. During their 2004 study, Cooke et al. determined that teams with consistent message content and communications flows demonstrated improved team performance. Consistent communications reduce variability within the team, meaning that fewer cognitive resources may be needed as the team predictably executes team processes.

One of the biggest hurdles to successfully implementing HMT is communication dynamics (Saenz, Revilla, and Simon, 2020). While machines only comprehend on a binary level, human teammates can decipher weak signals that may fail to trigger a response in an AI's algorithms. The human ability to place things into context and implicitly communicate improves team dynamics and is a human strength only accessible to machine partners through HMT constructs. Macmillan et al. (2004) note that for implicit communication and coordination to be effective, there must be a shared mental model. This raises the question; how would a team build a shared mental model with a machine partner?

These types of gaps in a machine's capabilities can lead to trustworthiness issues in the team, which may adversely impact a team's collective efficacy (Mathieu, Rapp, Maynard, & Mangos, 2010). Additionally, decision making matrices in gray area environments must be clearly defined as human input is critical in complex situations in which a machine may lack context (Saenz et al., 2020). Although this may sound like a fault, programming a machine partner to communicate about gray areas and to defer to a human teammate for its preferences could enable a machine to handle future exceptions more easily in the future (Russell, 2019).

This study will use the Espinosa et al. (2004) definition of coordination, which asserts that coordination is "the effective management of dependencies among subtasks, resources, and people" (p. 109). Generally, coordination mechanisms and dynamics are important to the effective operation of a team, and may be described as a process or an outcome in itself. Espinosa et al. (2004) noted that as teams grow and as the scope of work increases, the coordination mechanisms become vital due to the complexity of tasks and the team interdependencies. These vital coordination mechanisms are not fixed; they are dependent upon the tasks at hand, locations of team members, and the size of the team (Espinosa et al., 2004).

The explicit coordination mechanisms that teams use are dependent upon programming mechanisms such as schedules or formal plans as well as through explicit communications such as through utterances or written messages (Espinosa et al., 2004). Meanwhile, implicit coordination mechanisms are those mechanisms that teams are not purposely trying to coordinate through. Examples of implicit mechanisms include the concepts of team cognition and shared mental models that were discussed earlier in this chapter. Team members share knowledge about their tasks as well as each other, and will implicitly act to achieve their specified goal (Espinosa et al., 2004). Ultimately, the tools and means that are used for coordination are reliant on the team's interdependence model as well as its level of team cognition (Espinosa et al., 2004). In many instances it is also likely that team members will use implicit mechanisms more frequently than explicit ones as they become more familiar with each other, their tasks, and the context of their work (Espinosa et al., 2004).

Coordination mechanisms for teams are highly dependent on the task, team, and context variables that reside within tasks (Espinosa et al., 2004). Different combinations of coordination mechanism dynamics in changing contexts could result in the use of entirely different types of coordination, which could ultimately impact a team's dependencies and team performance (Espinosa et al., 2004). Stated differently, even teams that have excellent coordination may perform poorly if tasks are not well understood, team actions are not synchronized, or the strategy is not well suited to achieve the team's goals. In cases like these, where there are coordination problems, Espinosa et al. (2004) noted that a system's task, team, and context variables can be changed, which could mitigate coordination challenges more effectively than forcing the use of different coordination mechanisms.

Tasks have a strong influence on the type of dependency that is used (Espinosa et al., 2004). This makes innate sense when one considers the broad scope of work that a team may undertake. Teams may be employed in scenarios that range from highly mechanistic assembly line operations to specialized work that requires extensive, specialized knowledge, such as that of an aircrew or combat information center (CIC). Team dynamics and variables could also affect dependencies present in a team. Examples of these impactors include: team size, member skill levels, levels of familiarity with teammates, geographic dispersion, task uncertainty, and task load of teammates (Mathieu et al., 2000; Espinosa et al., 2004).

The contexts that a team operates in may also contribute to dependencies. When it comes to technological contexts, information technologies may change a team's roles, interactions, workflows, and information flows (Bergin, 2020; Espinosa et al., 2004). Organizational structures and factors may also impact team interactions, which could shape dependencies in tasks (Bergin, 2020; Espinosa et al., 2004). The geographic proximity of a team and the synchrony of a task can also impact team interactions and coordination (Espinosa et al., 2004). Specifically, teams that are separated in time or by distances may demonstrate coordination difficulties, arrive at a shared mental model, and may require additional coordination to achieve goals (Espinosa et al., 2004). Ultimately, Espinosa et al. (2004) noted that teams will likely need to possess and employ a repertoire of coordination mechanisms to complete their tasks as the context of the work is altered.

Designing machine partners to facilitate interaction and coordination dynamics in HMT can be described by Johnson's (2014) coactive design concept. Johnson (2014) asserts that machine partners should be designed in such a way as to account for team interdependence models, as well as to operate close to and continuously with human partners. Interdependence models may also clarify the interaction dynamics that exist in a team and are likely to demonstrate coordination functions that occur while the team is executing its tasks (Johnson, 2014). Interdependence models and relationships in team contexts have been extensively studied (Johnson, 2014), however it is important to note that as machines replace humans in teams, communication and interaction dynamics are also altered (Canan, 2021).

The disparate decisions, judgements, and communications between members of a team convey different meanings and represent an understanding of the environment in which the team operates. This understanding reflects the team's model of its cognitive system and environment (Canan, 2021). Team mental models could be influenced by communications, coordination, and interaction dynamics in some situations (Mohammed et al., n.d.). Although the literature is divided on the utility of shared mental models and their impact on team cognition and performance, there is a gap in relating shared mental models to HMT constructs (Mohammed et al., n.d.).

Despite the argument over team mental models in literature, teams must still be able to orient themselves as a key step in completing tasks. As teams orient themselves, collect information, develop situational awareness, make decisions, and implement those decisions, the dynamics of a situation should contribute to the information exchange between the team's agents. The team's orientation and its information exchange are particularly important given the substantial probabilities present in human cognitive systems (Canan, 2021). This can be most easily described as the fluidity of cause-and-effect relationships; more specifically, just because event A occurs does not necessarily mean that event B will occur, especially in team constructs operating in critical situations (Canan, 2021).

There is also an element of operational risk management that must be considered. If there is an undesirable outcome that could be realized, then human teammates may need

to assert control or otherwise adjust confidence intervals to prevent the machine from setting conditions for undesirable outcomes (Russell, 2019; Stumborg et al., 2019). Stumborg et al. (2019) specifically discuss modifying a machine's confidence intervals to loosen or tighten restrictions placed on a machine partner's actions. This ability to adjust a machine partner's confidence intervals could be based on a team's risk tolerance and could enable a team to dynamically adjust to its operating environment. This mechanism would allow human partners to engineer themselves into or outside of feedback loops.

As machines integrate with humans as agents in a team, there is the potential for additional uncertainty to be injected in human decision processes due to human mental indecisiveness (Canan, 2021). This is unsurprising as one of the innate capabilities of machine partners is their ability to ingest and identify trends in large volumes of data. If communication dynamics in the team are unbalanced, the machine partner could continually output more information than human team members could make sense of thereby causing information overload (Canan, 2021; Macmillan et al., 2004). As information overload occurs, it is likely that information paralysis will occur in the human partners (Canan, 2021). The converse could certainly be true in a team with poorly tuned communications dynamics. A team that does not exchange information with itself and the environment will find itself unable to mitigate uncertainty, and may suffer performance degradations due to challenges from mental models (Canan, 2021; Macmillan et al., 2004).

In ideal conditions, the different interactions that occur between members of a team and their environment are constructive and should help to mitigate the previously mentioned uncertainty. This may have the added benefit of building shared awareness amongst team members, enabling additional constructive implicit communications, which could result in a virtuous circle of more efficient communications patterns (Macmillan et al., 2004). Communications may be pushed to recipients without prompting, reducing the number of communications in a system, thereby reducing cognitive load and improving team performance (Macmillan et al., 2004). A potential area of interest could lay in information operations contexts. Information and communications challenges could degrade team performance and also increase levels of human anxiety, which may support a desired effect against a target audience (Canan, 2021).

E. TEAM PERFORMANCE AND AGILITY

Per Rentsch and Woehr (2004), one possible variable for team performance could be the team member characteristics that are most relevant to team functioning. When a team member harbors negative perceptions about a teammate's attributes, it could foster intrateam conflict and ultimately detract from team performance.

Human teams have proven themselves in uncertain environments where past experiences helped them to quickly make decisions based on intuition. Their ability to operate autonomously and interdependently has been a strength drawn upon by armies for centuries. By adding a machine teammate, a vulnerability is brought into question, as machines have yet to prove their ability to make decisions in irrational environments (Lawless, 2020).

Human teammates who have worked closely together can anticipate each other's decisions and actions. This leads to improved performance and team dynamics (Damachala, Javaid, & Devabhaktuni, 2018). This familiarity is often described as cohesion. Cohesion is an asset that can enable tight-knit teams to outcompete adversaries while serving as one of the intangible targets of an enemy system that should be shattered (HQ USMC, 1997). In military contexts, the concept of cohesion is emphasized as a critical team characteristic alongside these other important factors: team membership, team configuration structures, information flows, team goals, priorities, interdependencies, communication, and coordination (NASEM, 2021). When engaged in combat with a near peer adversary, a team's ability to quickly react and make decisions based on experiences has the potential to mean life or death. Marks et al. (2002) noted that team effectiveness can be severely compromised by team member mistakes, especially in challenging, highly inter-dependent, and time critical situations. Team interaction dynamics is undeniably an area that cannot be overlooked when introducing a machine teammate.

“Hot Groups” as defined by Lipman-Blumen and Leavitt (1999) are, “a lively, high-achieving, dedicated group, usually small, whose members are turned on to an exciting and challenging task” (p. 1). Hot groups are not typically planned, they are formed through the inspiration and dedication of human teams to solve a complex problem or defeat a

formidable adversary (Lipman-Blumen & Leavitt, 1999). When not restricted by bureaucracy, hot groups' inventiveness and ingenuity create dynamic solutions that have enormous organizational benefits. The driving force behind these groups is that they believe undertaking and completing the task, project, or mission will be for the betterment of their organization. Previous hot group members have reported that they experienced elevated levels of creativity and production, more than any other times in their careers (Lipman-Blumen & Leavitt, 1999). Their task becomes all they think and talk about. Many teams have even brought cots to the office so they can work throughout the night (Lipman-Blumen & Leavitt, 1999).

Collectively, hot group participants believe that the monumental task they are undertaking will make a positive impact on the world. The structure and size of a hot group is not perfectly defined. Typically, they range in size from 3 to 30 teammates, however, groups of over 15 are rare (Lipman-Blumen & Leavitt, 1999). Additionally, the lifespan of a hot group is short. They come together to solve a specified problem and then disband upon its solution (Lipman-Blumen & Leavitt, 1999). The significance of understanding hot groups is to attempt to predict how a machine teammate could affect the productivity of high-performing teams. Furthermore, learning how the machine teammate could help improve the agility and performance of a high functioning team such as a hot group will ensure their success (Song, Zurita, Gyory, Zhang, McComb, Cagan, & Stump, 2022).

The Oxford English Dictionary (2021) defines agility as, "the ability to think and understand readily and quickly; quick-wittedness, alertness; mental-dexterity" (p. 14). Agility in human-human teams can be described as the ability to dynamically respond to uncertainties and changing circumstances to form novel linkages between team members and facilitate the accomplishment of a team's goal (Mathieu et al., 2005). A more apt description could be adopted from Johnson's (2014) definition of resilience, "... the ability to recognize and adapt to handle unanticipated perturbations that call into question the model of competence, and demand a shift of processes, strategies, and coordination" (p. 16).

Teams that understand an enemy's purposes, strengths, weaknesses, and leadership characteristics are more able to achieve a desired effect against said enemy (Singer, 2009).

Doctrinally, there are warfighting functions and permanent staff members dedicated towards providing assessments of these attributes and exploiting them. What is less clear is how human-machine teams will dynamically communicate, update, and act with these attributes in mind. While there are C2 models for such a scenario, any capability integrated into a team construct must not introduce additional friction or uncertainty, or else it introduces the risk of eroding performance (Kennedy, 2021; Singer, 2009). The ability to distribute information into a larger network increases agility by enabling teams to “self-synchronize” their taskwork, bolstering team attributes that increase the team’s effectiveness (Singer, 2009).

Team interdependency is related to the performance of the best teams, whether in the civilian workplace or the military (Lawless, 2019). For human-machine teams to prove themselves able during conflict, the team’s agility must not suffer as a result of a lack of interdependency brought on by the machine teammate (Lawless, 2019). A flexible team will be able to dynamically navigate the challenges that they encounter. Thus, team agility can be achieved through adapting team interdependence relationships and structures to meet situational requirements (Johnson, 2014).

F. SUMMARY

This chapter aimed to frame how understanding teaming is essential to building a strong foundation for HMT. Shared mental models were discussed and how they present the ability to improve team situational awareness and performance, especially in critical situations such as combat. Next it was shown how teams act as cognitive system when they are able to coordinate effectively, plan, think, and act as a team. Building upon cognition, team situational awareness was discussed and how shared knowledge leads to higher team performance. Subsequently, team constructs were debated, focusing on how a machine can successfully be inserted into a human-human team and play a contributing part. Lastly, humans have demonstrated the ability to create agility in human-human teams by being able to anticipate the actions of teammates. Introducing a machine as a teammate may cause friction in that aspect and requires further study.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. A REVIEW OF COMMAND AND CONTROL CONCEPTS AND DOCTRINE

A. THE NATURE OF COMMAND AND CONTROL

A wealth of literature exists on the nature of C2. C2 can be described as the “means by which a commander recognizes what needs to be done and sees to it that appropriate actions are taken” (HQ USMC, 2017, p. 1-4). Kennedy (2021, p. 22) notes that an important aspect of C2 that is unmentioned in this definition is that it is a competitive activity where an “endless quest for certainty” is balanced against uncertainty and time. Command is an endeavor unique to the human partners of a human-machine team due to the limitations of current AI systems (Lange & Carreno, 2021). Specifically, only the human partner is capable of the creative and dynamic thinking needed to solve complex and novel problem sets.

Builder et al. (1999) explored the nature of C2 and resolved that it is the commander’s vision of military operations. He notes that this vision is articulated by a command concept, which provides all of the necessary direction that subordinate agents need to execute their assigned mission. In theory, a well-articulated concept would require less communication in the field, and would enable more effective team performance (Builder et al., 1999). Conversely, command concepts that are not clearly defined or prove untenable result in additional communication overhead and more frequent information exchanges with commanders (Builder et al., 1999).

The Marine Corps notes in its doctrine that command is a top-down flow of information and command concepts to subordinate agents, while control is the feedback that a commander receives (HQ USMC, 1996). The control feedback mechanism is specifically designed to assist commanders in determining if their decisions, command concepts, and the actions of subordinate agents have had the intended effect (Kennedy, 2021).

The Marine Corps’ doctrine calls for the use of maneuver warfare to the greatest extent possible (HQ USMC, 1997). Key tenets of maneuver warfare include a heavy

reliance on mission-type orders and commander's intent. Similar to Builder's (1999) command concepts, commanders issue mission statements that comprise the goal of their organizations. Their commanders intent provides additional guidance about said goal and relates it to enemy forces, friendly forces, and the environment. A core concept present in Marine Corps doctrine is that subordinate agents are free to improvise within the bounds of the mission and the commander's intent (HQ USMC, 1996). This freedom may be essential given the DOD's current operational concepts and its envisioned need for increasingly linked and distributed sensing and shooting platforms (Brose, 2020; Kennedy, 2021).

HMT constructs in military scenarios may be reliant on sensor information from these distributed sensor networks, and may need to aggregate disparate information components into a cohesive understanding of the environment. This understanding of the environment then serves as the context from which to measure the effectiveness of the team's actions with respect to their commander's intent and assigned mission.

B. COMMAND AND CONTROL SYSTEMS

C2 systems help commanders analyze a situation, choose a course of action, and quickly determine what assets he or she will employ. Although the Marine Corps defines C2, its definition fails to identify that decision making is a predominately human activity (Coakley, 1992) A C2 system is more than a framework, software, or technology attempting to guide decision making. Even with the ever-growing number of human-machine teams, the decision making still lies with the human, at least for now.

The Marine Corps lays its C2 foundation on the fundamentals of centralized command and decentralized execution (Kennedy, 2021). These mission-type orders allow a commander to convey his or her end-state and allow his or her subordinate commanders to determine how to get the job done. This permits increased tempo and agility as decision making is passed down to lowest level possible. This style of C2 promotes a decrease in required communication and is considered a more robust command style (Builder et al., 1999).

Undoubtedly, artificially intelligent machines will pair with humans to form teams on the battlefield. Commanding and controlling these teams will prove pivotal to the success they achieve. As it stands, commanders hold all responsibility for everything their unit accomplishes and fails to accomplish. This concept has stood at the core of military C2 principles. As Coakley (1992) pointed out, C2 is vitally important to the success of unit as a perceived lack thereof will surely bolster an adversary's confidence. Some may think tossing AI into a team might add to the challenge of communication in command. However, the authors would offer that the channel of communication is less clogged when ordering or tasking a machine. The machine only listens as programmed. It does not hear in analogous terms, slang, or dialect. In line with the theory of Shannon (Shannon & Weaver, 1963), it is the value of what is being transmitted, not how much. In other words, is the commander communicating clearly.

It is very possible for a commander to give a highly motivating speech to his troops prior to engaging in battle, even to the point where his warriors are excitedly ready to lay their life down for the cause. Unfortunately, without clear communication of a plan, they will simply be sacrificing themselves for nothing. However, if the commander's plan is clearly communicated and his focus is on his troops understanding the message, it is much more likely prove successful. Although his theory was written decades before the combination of AI and humans to form teams was conceived, it is still applicable today. It is unlikely that most commanders will ever be able to understand mathematical concepts of Shannon and Weaver's (1963) Mathematical Theory of Communication; however, they must grasp the concept that successful C2 is based on transmitting error-free communication.

Opposing clear communication is the current problem of a surplus of information. Builder (1999) points out that technology, while great, has inundated commanders with too much information. They are flooded daily with reports, charts, and spreadsheets to the point where making a decision becomes harder than without technology. We are now faced with the dilemma of applying more technology to help us decipher the massive amounts of data in hopes of presenting the right information (Builder, 1999). Paralysis by analysis, although cliché, seems to befit the impasse. Builder (1999) presents his ideal C2 system as one that

does not eliminate the technology, rather one that uses the technology to present ideal circumstances and performance measures as clear communication based on the information provided. After all, a C2 system includes all the resources, human and AI alike, which are essential to assist the commander in decision making.

C. SEMANTIC EFFECTS ON DECISION MAKING

Friction and semantic problems occur so frequently on the battlefield that they are recognized as an enduring attribute of the nature of war (HQ USMC, 1997). Such problems are often chalked up to the “fog of war,” an expression coined by the famous war strategist Carl von Clausewitz, however, here may be a more scientifically rigorous explanation (HQ USMC, 1997). Ultimately, the meaning of a communication issued on the battlefield may not be universally interpreted by its intended recipients; this semantic problem described by Weaver (1953) directly impacts the abilities of personnel at all levels to understand, decide, communicate, and act in their operational environment.

1. The Nature of Meaning and the Semantic Problem

Given the friction and disorder prevalent in C2 processes, it is important to discuss the nature of meaning and its relation to the information used in decision cycles. Weaver (1953) defined three levels of a communications problem that affects communications systems: Level A, the accuracy of technical means when transmitting messages; Level B, the precision by which a message conveys the desired semantic meaning; and Level C, the effectiveness of the message’s meaning in guiding the actions of intended recipients (Weaver, 1953). These challenges exist in a communications system, which by design, is capable of handling incredible volumes of data that could form any combination of messages (Weaver, 1953). The meaning of a message is subject to the channels by which it travels, and while traversing a channel, the message will encounter man-made and natural interference as well as distortion due to different interpretations of the possible meanings of a message (Weaver, 1953). This deformation occurs as a result of the users’ understanding of the situation, and is further shaped by their experiences, expectations, culture, and biases.

The semantic problem discussed by Weaver (1953) is the disparity of the transmitted and received meaning of a message and its influences on the actions of a recipient. Semantic errors can be reduced by the ensuring that transmitter and receiver possess and operate from a shared understanding (Shannon et al., 1963). A shared language can be built through an educational and doctrinal foundation and then reinforced with realistic training to ingrain an understanding in a wide variety of scenarios. The training regimen has the added benefit of breeding trust and familiarity between personnel in the chain of command, reducing the likelihood of semantic noise during the orders issuance and execution phases. A common dictionary of tasks and important terms may similarly give personnel a baseline from which to begin communications during execution.

This baseline should also be used in concert with published orders and commander's guidance to achieve a common understanding between commanders and subordinates. The friction, disorder, and rapid tempo of combat operations ensures that semantic errors will continue to exist in combat, however, these errors can be reduced to a manageable level when the commander articulates their guidance and ensures an adequate level of understanding of the orders to be executed (United States Marine Corps, 2001). The commander can also influence the use of the communications system as they exercise C2 to quickly assess the effects actions. By disseminating their information requirements and priorities amongst subordinates, commanders can prioritize the information they receive. This information is a competitive resource that can be processed by personnel to build meaning and generate knowledge (Department of the Navy [DON], 1995).

2. The Cognitive Hierarchy

The DOD employs a mixture of platforms with an array of organic and machine capabilities that can sense information about environments and targets, including activities and attributes (Schultz & Clarke, 2020; Kennedy, 2021). These capabilities acquire data in the form of raw signals and observations of the battlespace (DON, 1995). The raw data is then processed, where it is framed against experiences, expectations, culture, as well as biases, and output as information (DON, 1995). The information product is integrated with other pieces of information, compared to existing models, and assessed for accuracy and

value in a process that creates knowledge (DON, 1995). Through the application of military judgement this knowledge is converted into an understanding of the operational environment called situational awareness (DON, 1995). Such information may quickly change; thus analyses must support a commander's decision-making cycles and sustain a relative tempo beyond what the adversary can match.

As commanders make decisions, they pass down orders either verbally or written. If there is a barrier in the communication process, these orders may not come across as the commander intended. Semantic errors in the different phases of the cognitive hierarchy can distort the meaning of the collected data and lead to an incomplete or otherwise flawed understanding of the environment. Put simply, garbage coming into sensors and decision centers leads to garbage coming out. To ensure that subordinates possess the correct semantic meaning of the commander's understanding of the environment, the commander must continuously assemble, validate, and then publish his or her image of the operational environment (DON, 1995). Because the system may not always operate in ideal circumstances, additional errors in meaning are still possible and could lead to an inadequate understanding at all levels (Kennedy, 2021). For example, subordinates may misinterpret position and location information in a variety of circumstances or could inaccurately report information needed to feed decisions upstream. These could have unintended consequences as information flows through the cognitive hierarchy.

3. The Decision and Execution Cycle

The cognitive hierarchy's process of building understanding directly supports the decision making and execution cycle. The decision making and execution cycle changes ideas and intentions into action and occurs at levels of a warfighting organization (DON, 1995). This is a continuous and time-competitive activity that requires rapid decisions to generate disorder within an enemy system (HQ USMC, 1997). This concept is easily captured in the orient, observe, decide, act loop model. In the model, a decision maker senses the environment around them, derives meaning from the information to form an accurate mental model of the situation, then uses that awareness to make decisions that will lead to the desired outcomes, then implements those decisions (DON, 1995). The loop

portion of the model includes a feedback mechanism where decision makers observe and assess the effects of orders, deciding whether additional guidance and action is required (DON, 1995). These processes are vital to C2. The orientation a commander arrives at shapes the remaining actions within the cycle and is based on sensor inputs. From this orientation, the commander assesses the environment and begins to issue orders. Subordinates in receipt of the orders serve as part of a networked group of sensors that can relay information back to decision makers during execution.

A simple error with the semantics of a message during the C2 process could create adverse conditions on the battlefield. Misinterpreted information may cause a faulty understanding of battlefield conditions, leading to the issuance of orders that cannot effectively shape actions. Semantic errors in the feedback loop of the decision and execution cycle may cause a commander to needlessly commit reserves, withdraw forces, adjust priorities of fires, or perhaps worse, believe no further actions are required. Semantic errors can occur at all levels of the decision and execution cycle, and explain how the confused meaning and decisions introduced additional friction into the operational environment.

D. ORGANIZATIONAL STRUCTURES

There is no one-size fits all structure that allows an organization to easily mitigate every problem and environment that it will face. This holds especially true for warfighting organizations during the challenges posed by combat operations. C2 for air warfare, ground warfare, and naval warfare units all possess different organizational structures as their mission sets are different (Coakley, 1992). In other instances, C2 may be structured to serve geographical purposes or operational boundaries (Coakley, 1992). Regardless of the mission, the C2 structure of any organization exists to process information at various nodes and exchange communication (Builder et al., 1999). This hierarchy of reporting relationships thus becomes the chain of command.

Mintzberg (1981) delineates five configurations of organizational structures: simple structure, machine bureaucracy, professional bureaucracy, divisionalized form, and adhocracy. He based the five configurations on five component parts of an organization.

First, there is a person who comes up with an idea and sits at the *strategic apex* (Mintzberg, 1981). That person then hires workers to conduct the basic tasks and are known as the *operating core* (Mintzberg, 1981). As the organization grows, the *middle line* is hired consisting of intermediate managers (Mintzberg, 1981). Lastly, the organization may realize that it requires two types of staff personnel, the *technostructure* and the *support staff* (Mintzberg, 1981). The *technostructure* design systems that handle the planning and control of work (Mintzberg, 1981). The *support staff* is responsible for delivering ancillary services to the organization (Mintzberg, 1981). Together, these five parts perform all the functions of an organization. To be clear, not all organizations need each part. According to Mintzberg, “The central purpose of the structure is to coordinate the work divided in a variety of ways; how that coordination is achieved-by whom and with what-dictates what the organization will look like” (1981, p. 3).

1. Simple Structure

In a simple structure, coordination is accomplished by direct supervision (Mintzberg, 1981). In essence, the boss gives the orders. This may sound militaristic but it is not. Typically, there is the chief executive officer (CEO) and maybe a few top managers who oversee the operators conducting basic work. The archetypal entrepreneurial company is based on the simple structure dynamic. There is very little standardization, minimal planning, and negligible training (Mintzberg, 1981). The employees throughout the organization must maintain flexibility in an active environment. According to Mintzberg (1981), this is preferred as this maneuverability allows them to compete against bureaucracies. Their environment needs to remain simple along with the production system in order for the CEO to retain centralized control (Mintzberg, 1981). This centralized control allows them to remain agile and innovative, a trait unattainable by most bureaucracies. Most organizations begin as simple structures but most eventually evolve to bureaucratic form (Mintzberg, 1981). This shift to bureaucracy may be unwanted by the CEO but is a sign of growth and ultimately success.

2. Machine Bureaucracy

Machine bureaucracy is founded on the premise of the standardization of work and low-skilled, highly specialized jobs (Mintzberg, 1981). In contrast to a simple structure, a machine bureaucracy is analyst heavy to create and manage its systems of standardization (Mintzberg, 1981). Additionally, since the organization relies on these systems, the analysts actually achieve some informal power thus resulting in horizontal decentralization. Because of this, a hierarchy typically arises in the middle line to manage the specialized work of the operating core where conflicts tend to arise from the unyielding departmentalization and routine jobs (Mintzberg, 1981). This hierarchy of the middle line proceeds functionally to the strategic apex (Mintzberg, 1981). In essence there is horizontal centralization in the middle line, which in turns leads to centralized power vertically as the ultimate power lies at the top. In order to keep this power at the top, the machine bureaucracy still must remain fairly simple (Mintzberg, 1981). Examples of machine bureaucracies can be seen in mass production companies such as auto makers, government agencies, and even fast food (Mintzberg, 1981). Depending on where an individual resides inside a bureaucracy determines their like or dislike of its configuration. The majority of its players suffer through dull and repetitive work but bureaucracies still remain prevalent in today's mass-produced and mass-consuming world.

3. Professional Bureaucracy

Professional bureaucracy also relies on standardization but it is the standardization of skills as opposed to the work processes as in a machine bureaucracy (Mintzberg, 1981). Examples include universities, hospitals, and large accounting firms. In a professional bureaucracy, the skills of its operating core are much higher therefore these employees are usually given control over their own work resulting in a decentralized structure (Mintzberg, 1981). In other words, decision making is made throughout the organization. As the work being done is complex, it is still standardized thereby most of the operating core works independently of their coworkers. As Mintzberg (1981) noted, a colleague of his watched an open heart surgery take place in which the doctor and anesthesiologist never spoke a word to each other; the surgery was five hours long. This example demonstrates how a

technostructure is barely needed in a professional bureaucracy. As a result, minimal line managers are required and for the few that do exist, little management is required. They spend most of their time conjoining their units to the larger environment (Mintzberg, 1981). The support staff on the other hand, is usually quite large in order to support the well-educated and highly paid operating core (Mintzberg, 1981). This results in a dynamic power structure where the operating core possesses democratic bottom-up power; however, there is also top-down control maintained by the strategic apex in more of an autocratic fashion (Mintzberg, 1981). A professional bureaucracy provides a good foundation for companies who reside in stable but multifaceted environments (Mintzberg, 1981). Additionally, although standardization provides its greatest strength, it does also prevent a professional bureaucracy from being flexible (Mintzberg, 1981).

4. Divisionalized Form

Similar to a professional bureaucracy, an organization structured on divisionalized form consists of separate entities linked by administration oversight (Mintzberg, 1981). The main difference between the two being that a professional bureaucracy consists of individuals in the operating core being the focus whereas in a divisionalized form units called divisions make up the middle line (Mintzberg, 1981). These divisions exist because their product lines differ. The divisions are typically a result of the organization attempting diversify from their original founding market focus. For the most part, each division runs its own operations. There is typically a division head who, for the most part, makes semi-autonomous decisions. This may lead some to see this type of structure as decentralized but according to Mintzberg (1981), “a divisionalized structure in which managers at the heads of these units retain the lion’s share of the power is far more centralized than many functional structures where large numbers of specialists get involved in the making of important decision” (p. 9). General Motors is a great example of divisionalized form. Here, the strategic apex still maintains direct supervision in the form of a headquarters. The headquarters sets standards in the form of outputs but leaves operational decisions in the hands of division managers (Mintzberg, 1981). The headquarters in turn holds them accountable for reaching milestones and overall division performance.

5. Adhocracy

The final organizational structure is adhocracy. Briefly stated, adhocracy is based on an organization of intermingling project teams (Mintzberg, 1981). It is considered the most difficult to effectively operate as it can be non-standardized and complex. In an adhocracy, power is constantly shuffled while control and coordination happen via mutual modification typically through unceremonious communication of experienced experts (Mintzberg, 1981). It relies on specialists who are highly trained to work throughout the structure, not just in the operating core. Unlike a professional bureaucracy, these experts create new products or output by working together (Mintzberg, 1981). Managers inside an adhocracy typically control narrow lanes by orthodox measures. This should not be viewed through the scope of power breadth, but it is mainly due to the small size of project groups (Mintzberg, 1981). Also different is the fact that the managers usually work hand-in-hand with those they oversee. Because all input and ideas are considered, strategy does not always come from the strategic apex, rather it forms as new projects are brought on board and worked upon (Mintzberg, 1981). Examples of an adhocracy consist of tech startup firms, consulting agencies, and even the early years of the National Aeronautics and Space Administration (Birkinshaw & Ridderstrale, 2010).

E. BOYD'S OODA LOOP

Colonel John Boyd's observe, orient, decide, and act (OODA) loop has been a highly influential concept used by the militaries worldwide as a framework to achieve a competitive advantage over adversaries (Richards, 2020). Its premise is founded on the idea that success in warfare is based on the ability to out-think and out-pace the adversary (Osinga, 2007). On this notion Colonel Boyd created the loop, which is meant to begin at *observe* and then iterates until *act*. In the observe phase, an agent attempts to identify a problem or threat and an understanding of the surrounding environment (Richards, 2020). In the orient phase, the actor then reflects on the observations made during observe to consider what should be done next (Richards, 2020). The decide phase then builds take knowledge gain in the first two steps and makes an action plan considering all outcomes (Richards, 2020). Finally, in the act phase, actions are taken based on the decisions made

in the previous step (Richards, 2020). In essence the competitor that can achieve action more quickly will ultimately achieve an insurmountable advantage, which then leads to victory (Richards, 2020).

Although Boyd’s OODA loop has been the focus of many studies and is fairly well-known, he never sketched the image publicly until 20 years after he first began using the term (Osinga, 2007). Most people know the sketch as a simple loop with the four phases looping iteratively. However, when Colonel Boyd did finally sketch the loop it was actually not a loop at all (Osinga, 2007). This sketch can be seen in Figure 4..

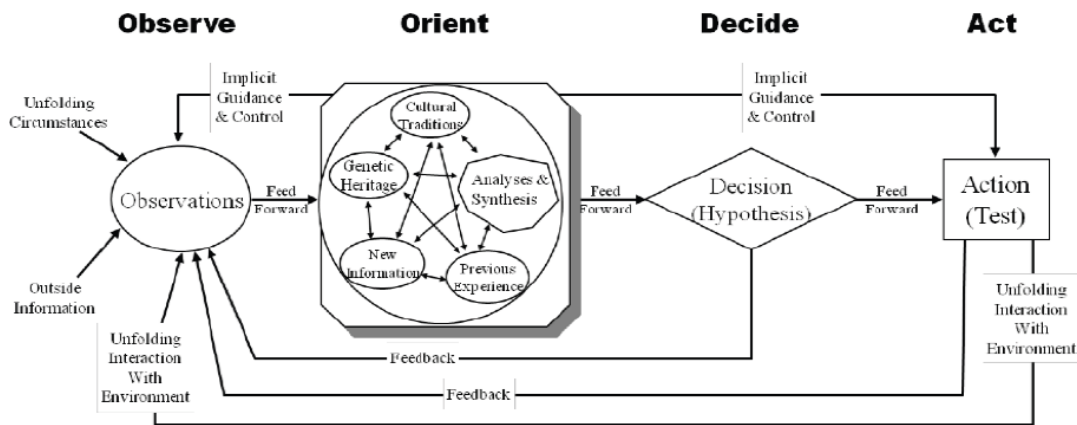


Figure 4. Colonel John Boyd’s OODA Loop. Source: Osinga (2007).

There have been many critics of the loop who state that it is too simple. They argue that a sequential loop cannot model an organization in conflict (Osinga, 2007). The authors of this study would offer that the majority of those critics likely do not know Colonel Boyd’s actual sketch, which involves more than the four simple phases. However, there are limitations. The OODA Loop model centers on its requirement for timely feedback between the actor, the environment, and external agents. A decision-maker may not receive the near-instantaneous feedback required by the model, and as such, the effects of one action may not emerge until after an actor has already made subsequent decisions (Aycock & Glenney, 2021). Additionally, sometimes it is advantageous to allow a course of action

to play out before deciding in an effort to create more favorable actions in the future (Osinga, 2007).

F. SUMMARY

In sum, C2 can be deduced to compiling communication via functional nodes allowing a commander or executive to make a more informed decision while ensuring those decisions are performed (Stanton & Baber, 2017). Macmillan et al. (2004) observed that pre-event planning can contribute to improvements in team mental models and situational awareness, enabling them to communicate more efficiently. These communication efficiencies were correlated with higher levels of team performance (Macmillan et al., 2004). Moreover, C2 is vital to success as a perceived lack thereof by an opposing force invites attack without reprisal (Stanton & Baber, 2017). Lastly, breakthroughs in technology has greatly extended the reach of C2, sometimes to the dismay of the middle line (Stanton & Baber, 2017).

THIS PAGE INTENTIONALLY LEFT BLANK

V. A REVIEW OF ARTIFICIAL INTELLIGENCE THEORY, CAPABILITIES AND LIMITATIONS

What do the spam filter for a modern e-mail system, Amazon’s online product recommendation system, and iRobot’s Roomba all have in common? They are all based on some form of AI (MIT Horizon, 2022a; NASEM, 2021; Domingos, 2015). These systems demonstrate some of the different capabilities resident within AI that will be discussed within this chapter. AI has been in use for decades and has been used to help humans complete tasks—routine and dangerous alike (Singer, 2009; Grosz et al., 2016; Scharre, 2018; NASEM, 2021). In the DOD, the routine tasks include processing intelligence, surveillance, and reconnaissance data, while some of the more dangerous tasks include autonomous operations within adversarial environments (NASEM, 2021). This chapter will provide an overview of AI, describe its intended benefits, and will conclude with some of AI’s current limitations.

A. WHY AI?

AI is a multi-faceted tool that could enable a significant competitive advantage over adversaries (NASEM, 2021). Such a competitive advantage is hypothesized to completely alter the practices and paradigms of warfare (Fout & Ploski, 2018). This new way of war is colloquially called an RMA, which can be defined as a change in how wars are fought, typically because of innovative technologies or employment methods (Sloan, 2002). RMAs can render previous doctrine and proven concepts obsolete, allowing competitors to gain a temporary advantage. Since the 1300s, military historians have identified at least 10 RMAs. Several notable examples, the English longbow, the cannon and cannon ball, the rifle, and the advent of precision fires have all made irreversible changes to warfare and tactics (Singer, 2009). Some of the new changes in tactics and capabilities potentially enabled by AI are revolutionary, including predictive targeting, automated mission planning, as well as the ability to decide and act at the rapid pace promised by machine speeds (NASEM, 2021; Stumborg et al., 2019)

Spam filters, credit rating systems, and product recognition systems are some of the more common systems that people commonly interact and reflect how AI is becoming an increasingly common part of everyday life (Domingos, 2015). The DOD seeks to mirror this trend, and is making plans to deploy AI at all levels of the enterprise (DOD, 2018). AI is well suited to developing timely, decision quality information, optimizing the use of military resources, and being able to quickly ingest and make sense of incredible volumes of data for the targeting of fires, or the forecasting of maintenance requirements (NASEM, 2021). This is important to note because AI systems are incredibly capable, and have a broad range of use cases.

Some of these use cases are specifically aimed towards enhancing the ways that humans function. DOD intends to integrate AI and humans to improve the performance of human teammates operating in human machine team constructs (Stumborg et al., 2019; NASEM, 2021). AI can do more than serve as autonomous wingmen or serve as ground-based fires platforms; they may also help teams manage goals, ensuring that they remain aligned as changes occur in team structures, tasks, and goals (Tirpak, 2021; Singer, 2009 NASEM, 2021). It is also possible for an AI to serve within a team as a communication and coordination hub, which would optimize team behaviors by connecting members of teams, prioritizing the flow of messages, and by helping humans clarify miscommunications (NASEM, 2021).

Underpinning the capabilities of AI is the fundamental assumption that it exceeds the capabilities of normal computers and most humans (MIT Horizon, 2022b). While this is certainly true in some instances, AI is only designed to function well in specific use cases (Russell, 2019; NASEM, 2021; MIT Horizon, 2022d). These areas will be discussed throughout this chapter.

B. THE DIFFERENCE BETWEEN AUTOMATION AND AI

At its core, AI is a broad label applied to a category of technologies with the express purpose of using computers to make automated decisions or predictions (MIT Horizon, 2022d). These technologies rely on large volumes of data as well as different algorithms to make sense of the data and perform their assigned task until a goal is reached (Domingos,

2015; NASEM, 2021; MIT Horizon, 2022d). One of the primary goals of these technologies is to create additional value for their users (Grosz et al., 2016; Russell, 2019). Unceremoniously, AI has been associated with computer systems performing tasks that used to require human skills and intelligence (MIT Horizon, 2022d). AI is subtly different from a similar capability, called automation. Automation is generally defined as category of technology that performs tasks independently, without the need for constant user inputs (NASEM, 2021).

It is worth noting that automation is not an “all or nothing” approach. Consider automation as a spectrum. On one side of the spectrum, humans may still oversee some parts of a specific task while foregoing supervision of others (NASEM, 2021). The other extreme is a fully autonomous system that requires no direct human oversight, and may display limited forms of self-directed behavior (NASEM, 2021). The latter is commonly referred to as “human out of the loop” (Bode & Watts, 2021). One example of this capability within DOD is the MIM-104 Patriot missile battery. Humans can directly control the acquisition and servicing of a target, or may simply set the system to autonomous mode and allow it to engage targets automatically (Bode & Watts, 2021). Since the definition of automation is very closely related to tasks that a human previously used to do, the capabilities and descriptions of automation will change over time (NASEM, 2021).

As AI functions similarly, characterizations of whether a system possesses AI or not will also change over time. According to NASEM (2021), AI is “a highly capable form of automation that could be used to sense and interpret situations, adapt to changes, prioritize and optimize based on goal changes, and improve its abilities based on learning” (p. 7). The learning capabilities resident within AI enable developers to set ambitious research goals for future systems. One such goal is to achieve the same level of intellectual capacity and memory as a human, including abilities to think creatively, place items into context, generalize, and to learn from its past experiences (Domingos, 2015; Russell, 2019; Stumborg et al., 2019; NASEM, 2021; MIT Horizon, 2022d). A system that can perform these tasks is decades or possibly centuries away (MIT Horizon, 2022b).

C. AI IS ALGORITHM BASED

The aforementioned goal of an AI that can perform at the same level as or beyond the capability of a human in a broad range of tasks is called artificial general intelligence (Stumborg et al., 2019; MIT Horizon, 2022d). Essentially, such a system will combine the capabilities of many disparate current AI systems into one agent. The AI systems that are deployed today are used in single, well-defined tasks, and are called narrow AI systems (Stumborg et al., 2019; MIT Horizon, 2022d). Simply because a system is narrow does not mean that it is incapable of performing complex tasks. Because of the increasing capacity of current processors and data storage capabilities, systems are becoming ever more capable (MIT Horizon, 2022d). For example, the autopilot and traffic aware cruise control modes of the Tesla vehicle are AI-based (Posada, 2021). Amazon’s Alexa and Apple’s Siri are examples of AI-based personal assistants that are gaining popularity, and may be found frequently in homes (MIT Horizon, 2022a). All of these AI-based capabilities require sophisticated algorithms as well as large data sets and skilled developers for training (MIT Horizon, 2022c).

In some algorithms, such as narrow rule-based systems, developers are required to manually define inputs, processes, and outcomes (MIT Horizon, 2022c). Systems such as these are also known as expert systems because the logic deployed in the system follows that of a human expert. (MIT Horizon, 2022d). Following the pre-programmed instruction sets, an AI interprets the process inputs and matches them to what it expects based on the rules built into it (MIT Horizon, 2022d). One drawback to expert systems is their inability to learn new information, or to accurately generalize in unfamiliar contexts without manual updates from human programmers (MIT Horizon, 2022d).

In an attempt to bypass the limitations of rule-based systems and their heavy reliance on humans, machine learning algorithms were introduced. In a machine learning system, the system is given a goal and numerous examples of a task, event, or classification (MIT Horizon, 2022d). Unlike rule-based systems, machine learning systems learn patterns from training data and apply them to reach the goal (MIT Horizon, 2022d). Ultimately, the system’s accuracy grows as the number of data points increases—some open-source systems are trained on billions of data points (MIT Horizon, 2022c). The utility of the

machine learning system is that because of its rigorous training, it may detect patterns that a human may have otherwise missed, and it is able to adapt to changing environments and circumstances (MIT Horizon, 2022a).

Supervised learning is the first machine learning that will be discussed in this study. Supervised learning occurs when a human developer directly shapes the information and purpose that an AI receives during its training program (MIT Horizon, 2022d). Within this machine learning technique, a human has organized and described that training data that will be ingested by a machine learner (MIT Horizon, 2022d). The descriptions of such training data are commonly referred to as labels (Domingos, 2015; MIT Horizon, 2022d). Similarly to a rule-based system, the AI is provided with its desired outcomes and conclusions during the learning process (MIT Horizon, 2022d). A supervised learner excels at tasks centered towards classification or regression—classification is the ability of the AI to sort disparate items such as images, while regression is the ability of the AI to perform statistical operations (MIT Horizon, 2022d).

A similar use case applies to unsupervised learning-based systems. An unsupervised learner is the inverse of a supervised learning system. More specifically, an unsupervised learning AI system is trained on data that has not been organized or labelled, and it is not issued a specific conclusion to draw from this data set (MIT Horizon, 2022d). Essentially, the system is tasked to conduct an exploratory data analysis, learning the data for itself (MIT Horizon, 2022d). As it discovers underlying patterns, the system will invent groupings and present its conclusions to users (MIT Horizon, 2022d).

This differs from reinforcement learning. With a reinforcement learner, the developer provides the AI system with potential conclusions, and trains the system on partially labelled data (MIT Horizon, 2022d). Unlike the aforementioned learning algorithms, reinforcement learning develops intelligence by learning from trial and error—learning with this algorithm focuses on sequences, not single events or data points (MIT Horizon, 2022d). Given the set of sequences and conclusions, reinforcement learners will learn from reward signals in their environment to make decisions as a means to reach their programmed goal (Russell, 2019; MIT Horizon, 2022d).

Armed with a basic understanding of the algorithms discussed in this chapter, it is possible to gain an understanding into the final algorithms—deep learning and neural networks. These machine learning methods are dependent upon complex data sets; however, they can handle increasingly complex tasks (MIT Horizon, 2022d). The autopilot and traffic aware cruise control modes of Tesla vehicles are an excellent example of what deep learning and neural networks can do for users. The algorithms inside the Tesla allow it navigate roadways and deal with the complexities operating a motor vehicle on congested roadways (Posada, 2021). The neural networks at the core of deep learning interpret patterns based on mathematical units called neurons (Domingos, 2015; MIT Horizon, 2022d). These neurons function interdependently, and allow the algorithm to interpret the patterns and interactions resident in a data set (Domingos, 2015; MIT Horizon, 2022d). The basic structure of a neural network is comprised of interdependent inputs and outputs that yield a field output. An illustration can be found in Figure 5.




Figure 5. Simple Neural Network Design

To initiate the learning process, a developer provides the system with a series of instructions, and the AI will assign different weights to neurons in the input and hidden

layers according to the instructions (MIT Horizon, 2022d). Ingesting the data, the neurons in the input layer send their outputs to the interdependent neurons in the hidden layer for processing (MIT Horizon, 2022d). Although Figure 5 depicts three rows of neurons, it is important to note that there can be many layers of hidden neurons. Adding layers of hidden neurons enables the AI system to perform increasingly complex tasks as well as understand intricate relationships and phenomena (Domingos, 2015; MIT Horizon, 2022d). If there are additional rows of hidden neurons, then the outputs from the previous row will be processed and fed as inputs into the next row until the final layer is reached (MIT Horizon, 2022d). At the final layer, the neural network provides users with its conclusions, based on the processing steps previously discussed (MIT Horizon, 2022d).

Because of the amount of processing done in neural networks, they are incredibly useful for discovering hidden or complex patterns from large quantities of data. This utility comes with a cost—developers may not be able to explain how the system’s processing led it to a specific conclusion (NASEM, 2021; MIT Horizon, 2022c). Certainly, the average user will not be able to trace the inputs and derivation of an unexpected or incorrect conclusion as information flows across a neural network (MIT Horizon, 2022c). Essentially, the learning processes of neural networks become opaque, limiting the ability of developers and users alike to understand what and how the system is learning, turning some neural network systems in black boxes of intrigue.

Failing to understand and visualize the learning process means that developers and users will encounter challenges when using these algorithms in real-world conditions. It becomes difficult to understand the system’s boundaries and also to predict under which conditions it may break down (Bansal et al., 2019; MIT Horizon, 2022c). Such opacity and difficulty also may mean that repairing the problem could be difficult, especially since neural networks are challenging to scale up and take a considerable amount of time to train (Domingos, 2015; MIT Horizon, 2022c). Interestingly, training these systems is an incredibly energy intensive task—this could prove to be another obstacle to repairing and training such algorithms in expeditionary environments (MIT Horizon, 2022c).

D. THE LIMITATIONS OF AI

While AI is capable of outperforming humans in some areas, as of the time of this study, it is still not able to outperform people in all areas. This is because AI systems generally lack the ability to reason and to make sense of their surroundings through context (Russell, 2019; MIT Horizon, 2022c). Many algorithms are specifically designed to capitalize on pattern recognition, and as such, cannot use reason to infer cause and effect relationships (Domingos, 2015; Russell, 2019; Aycock & Glenney, 2021; NASEM, 2021). This creates a second order effect where it cannot accurately forecast the future state of a system because it does not understand the conditions upon which data points were created, and therefore does not understand how a system is altered through the effects of actions (NASEM, 2021; Scholkopf, Locatello, Bauer, Ke, Kalchbrenner, Goyal, & Bengio, 2021). This may be because AIs rely on feedback mechanisms regarding decision results and learning processes (Domingos, 2015). NASEM (2021) noted that while AIs are learning, their initial adaptations will likely result in poor performance. While narrow-AI algorithms are trained to do one single, limited task exceedingly well, generalizability and exception handling are absent because these systems are hyper-focused in a single area, and their capabilities do not necessarily transfer across domains (NASEM, 2021; MIT Horizon, 2022c).

Amazon's Alexa or Apple's Siri are AI-enabled personal assistant applications that provide users with the impression of human-like skills and capabilities. While they are able to understand verbal queries and commands when uttered in ideal circumstances, the responses they generate come from a complex pattern-matching algorithm that trained the systems to match queries and commands to common phrases, and to yield an output that was associated to the utterance (MIT Horizon, 2022a). AIs such as these rely on statistical analyses as opposed to human-like cognitive processes (MIT Horizon, 2022a).

Recent high-profile crashes involving Tesla vehicles demonstrate how narrow AI systems lack the ability to place items into common-sense contexts. Tesla vehicles have mistaken false markings on the roadway to be indicate of changing travel lanes, duct tape on speed limit signs for entirely different speeds, tractor trailer containers for the open road, and have even mistaken the moon for traffic signals (Posada, 2021; MIT Horizon, 2022c).

Because AI systems are incapable of independent reasoning, they are easily fooled by challenges in contexts where the inputs, processes, and outputs are seemingly obvious.

This challenge is referred to as “brittleness” (NASEM, 2021, p. 8). In the aforementioned examples with malfunctions in Tesla’s AI, the system struggled to make sense of data that was presented in a slightly different context than it was trained on. In the case of Tesla, there have been more than a dozen high-profile accidents where the vehicle’s AI has caused a collision with first responder crews (Posada, 2021). Such brittleness implies that humans may need to adjust how they operate around systems and be prepared to rapidly assume tasks because of limitations within their AI systems (NASEM, 2021). Because the systems are limited to select use cases and do not generalize well, humans may become burdened with monitoring systems for breakdowns in common sense.

Aside from limitations in a system’s ability to apply common sense, modern AI-enabled systems require vast volumes of training data to operate. This makes sense as an AI is designed to derive conclusions from its environment based upon generalizations learned in training (NASEM, 2021). Teaching a system about the world requires a significant number of data points. As previously mentioned, Google’s proprietary computer-vision algorithm is trained on approximately 3 billion data points (MIT Horizon, 2022c). Compounding the training data challenge, machine learners tend to make more accurate generalizations when these data points are extracted, or otherwise representative of, real-world data (MIT Horizon, 2022c). Given that algorithms are designed for very specific purposes, collecting, labelling, and maintaining training datasets could prove to be problematic for the DOD, especially given the diversity of its intended use cases.

Certain algorithms, such as machine learners, are able to ingest data points and dynamically learn over time (Domingos, 2015; NASEM, 2021). This constant learning could the adjust behaviors and error boundaries for a system in non-obvious ways (Bansal et al., 2019; Russell, 2019; NASEM, 2022). As a system’s logic, strategies, and capabilities are changed in real time, developers and users may experience difficulties when trying to understand how training processes contribute to an AI’s observed processes and outcomes (NASEM, 2021). Certain functional and non-functional design requirements, such as

transparency and explainability may combat some of these challenges and will be discussed later in this chapter.

There may also be significant privacy considerations for the data used to train algorithms. Algorithms have been deployed in the medical field, where they are used in tasks ranging from identifying trends in health records to identifying anomalies in x-ray scans (Fusco et al., 2021; NASEM, 2021; MIT Horizon, 2022a). AI-enabled systems are also widely used for analyzing biometric data for use cases such as facial recognition (Domingos, 2015; NASEM, 2021; MIT Horizon, 2022a). Because AI systems must have data for the training process, developers must seed the algorithms with at least some data to ensure that it will function as intended (MIT Horizon, 2022c). This raises an important privacy-related questions: How does one protect the privacy of data in conditions such as these, given the need for algorithms to access and integrate real-world data? Where did the images, records, or other data from the training set come from? What would happen if the data was poisoned by an adversary or biased?

Such questions have strong, obvious ties to propaganda and information operations topics. There is also the insidious problem of human biases as they relate to AI-enabled systems. Assuming that a developer would be able to remove all traces of biases from algorithms and their training data, humans may incorrectly use an AI in ways that reinforce their own biases (Liu et al., 2021; MIT Horizon, 2022b). In a sort of confirmation bias, humans may know that an AI is impartial and unbiased, but may assume that the AI's are correct in ways that align with their own preconceived notions (Kahneman, 2011; Liu et al., 2021; MIT Horizon, 2022b). If there is a conflict between the conclusions of the AI and the human, the human may also exhibit automation bias, where they believe that system is more correct or capable than it actually is (Russell, 2019; NASEM, 2021; MIT Horizon, 2022b).

The generalizations and conclusions drawn by an AI system are dependent on the data used to the train the system (Choi, 2021; MIT Horizon, 2022c). Data sets must contain an adequate number of data points, which must be both accurate and relevant to the context that the AI will be operating in (MIT Horizon, 2022c). When systems are trained on inaccurate or incomplete data sets, it may lead blatantly incorrect conclusions, or it could

lead to incorrect answers with errors that are subtle and difficult to detect (Stumborg et al., 2019; Bode & Watts, 2021; MIT Horizon, 2022c). There may also be biases introduced from the way that inputs were used to build associations (Russell, 2019). Biased data may lead to non-obvious biases within the systems, and it could also manifest as improper characterizations of an AI's operating environment.

As an example, Papernot, McDaneil, Goodfellow, Jha, Celik, and Swami (2016) demonstrated that some AI systems, such as deep learning systems can be led to misclassify sensory data inputs, leading to unexpected outcomes. In Papernot et al.'s (2016) work, the authors observed systems misclassifying roadway stop signs for yield signs and rifles for helicopters, amongst many others. Papernot et al. (2016) adapted an adversarial approach, and purposely crafted images that would be perceived as correct by humans, but would be misinterpreted by machines. Thus, biased training data may result in an AI system that presents with its own biases.

Biases may also emerge from the structure of the algorithm. Some of the leading AI algorithms and models available from companies such as Google or OpenAI are built upon previous versions of an algorithm (MIT Horizon, 2022c). As new models are being layered on top of previous, foundational models, a model's previous biases may be carried forward into the upgraded version (MIT Horizon, 2022c). As an example, the developers of OpenAI's GPT-3 model programmed the algorithm to account for 175 billion different parameters, growing the number of parameters nearly 17 times from the previous version, GPT-2 (MIT Horizon, 2022c). Even though this represents a significant improvement over the previous version, and the algorithm is capable of handling increasing levels of complexity, it is possible that some of the hidden biases present in the previous foundation will emerge in unanticipated ways within the new version of the algorithm (MIT Horizon, 2022c). Given that many of the market-leaders in AI offer their algorithms to the public, sans-training data, it is possible that biased algorithmic structures may be promulgated (MIT Horizon, 2022c).

Uncovering biases may become difficult as algorithms grow in complexity. Akin to human biases, there will likely always be biases present within training data (MIT Horizon, 2022b). The aforementioned "black box of intrigue" challenge was discussed in

this chapter. If flawed data resides in the system it may lead to difficulties in uncovering why a system functioned the way that it did, or arrived at a specific conclusion. This problem exists with accurate and suitable data points. Called the “black box problem,” AI systems can solve increasingly complex problems based upon billions of variables, but they may not be able to explicate how they arrived at that conclusion or the why behind the answer (MIT Horizon, 2022c).

Such a lack of transparency is a problem for developers; however it may serve to undermine the trust that users place within a system (MIT Horizon, 2022c). Users may realize an erosion of trust in this condition because they can no longer follow the logic of the system as it operates (NASEM, 2021). High profile accidents, such as the recent string of Tesla crashes may similarly undermine a user’s trust and confidence in their systems (Posada, 2021; MIT Horizon, 2022c). The aforementioned privacy considerations could also adversely impact a user’s trust in their systems (MIT Horizon, 2022c). Explainability transparency, and trust may have significant, positive impacts on human machine teaming, and as such will be discussed in the following chapter.

E. THE BENEFITS AND CAPABILITIES OF AI

AI has many beneficial capabilities that could be used within private, industry, and defense contexts to great effect. Beyond assisting humans with dangerous, dull, or dirty tasks, AIs can recognize and act upon human speech, perform tasks such as predictive maintenance, and even help mitigate the effects of information overload (Grosz et al., 2016; Stumborg et al., 2019; MIT Horizon, 2022a). While the aforementioned limitations certainly apply to AIs and their capabilities, these new capabilities could signal that a new way of fighting and winning wars is possible (Singer, 2009). If AI is adopted in the correct manner, warfighters may be able to realize the benefits of AI and use its capabilities accordingly.

Systems now have the ability to interpret and process written and spoken languages in much the same way as humans do. Using natural language processing, AI systems can recognize an utterance or writing and take some action upon it (MIT Horizon, 2022a). Although natural language processing algorithms may struggle with regional accents,

technical jargon, signal noise, or background crosstalk, this still represents an excellent capability (MIT Horizon, 2022a). Liu et al. (2021) discussed that natural language processing could be used to help AIs explain visualizations and outputs. It is also possible to deploy natural language processing as a means to assist developers with debugging and improving AI models (Liu et al., 2021). The natural language processing applications and capabilities available at the time of this study far exceed Amazon's Alexa or Apple's Siri.

This capability similarly works with written communications (MIT Horizon, 2022a). Miller (2021) demonstrated the intersection of machine learning and natural language processing with DOD's novel big data analytics application, Advana. Using Advana, Miller (2021) was able to summarize documents, identify critical aspects of policies that had changed from previous versions, and deconstruct the document as a series of nodes and edges to see how one document relates to another. Other applications for natural language processing include social media sentiment analysis (Hirschberg & Manning, 2015). Additional algorithms may also be able to provide statistical analyses about demographics and relationships (Hirschberg & Manning, 2015). In a particularly fascinating instance, Canada's GPHIN system was able to mine data about a mysterious respiratory infection in December 2019, several months before the World Health Organization learned about COVID-19 (Robertson, 2020; MIT Horizon, 2022a). The latter two examples have obvious military intelligence and information operations use cases.

Using computer vision, AI systems can assist users by interpreting imagery. While natural language processing can interpret numbers and text, computer vision algorithms enable AI to understand different visual stimuli (MIT Horizon, 2022a). This algorithm is capable of interpreting hand-written text, and can convert it into electronic text (MIT Horizon, 2022a). Modern versions of computer vision enable facial recognition software, Tesla's autopilot and traffic aware cruise control modes, as well as AI-enabled analyses of satellite imagery (Ingle & Phute, 2016; Nachmany & Alemohammad, 2019; Posada, 2021; MIT Horizon, 2022a). The ability of AI systems to consume vast amounts of intelligence, surveillance, and reconnaissance data, and look for data points of interest is a capability of significant interest to the DOD (Brose, 2020; NASEM, 2021).

Although the discussion in the limitations section of this chapter identified that AI could not determine cause and effect relationships, and therefore struggled with predictions, there are some circumstances in which AI is effective in making predictions. One such area where AI can provide suitable forecasts for humans to rely on is in predictive maintenance (Hrnjica & Softic, 2020; NASEM, 2021 ;MIT Horizon, 2022a). Areas such as predictive maintenance contain a significant number of interrelated variables, and could be done by humans (MIT Horizon, 2022a). However, an AI is capable of evaluating the data, variables, and relationships at a faster rate than a human can, and may do it with a smaller margin of error (MIT Horizon, 2022a).

This could also be the case for systems that are difficult to model (MIT Horizon, 2022a). Certain AI algorithms may prove useful to humans when they are trying to model systems where relationships among variables are not well known or short-lived (MIT Horizon, 2022a). In cases such as these, AI may prove to be an accurate modeling resource (MIT Horizon, 2022a). To aid in predictions, AI systems may generate their own data sets, and use those data sets to identify the dimensionality of data, noise, and any emergent patterns (MIT Horizon, 2022b).

Akin to predictive maintenance, AI can help people make sense of complex systems and environments. Specifically, because AI may be useful in analyzing systems where there are interconnected and interdependent variables, they could help humans sift through the layers of complexity to identify patterns (MIT Horizon, 2022c). Modern AI systems are capable of working with large data sets to find patterns in datasets that may overwhelm a human counterpart (MIT Horizon, 2022c). This could help the humans in many cases ranging from identifying maintenance abnormalities before they become conspicuous to other cases such as combat operations, where short-term tactical implications must be quickly evaluated (Stumborg et al., 2019; Hrnjica & Softic, 2020; MIT Horizon, 2022a).

Finally, AI is beneficial because some algorithms are able to learn and adapt over time. Machine learning algorithms are able to improve through training processes that expose it to both training and real-world data, and do not require constant human intervention (Domingos, 2015; MIT Horizon, 2022d). When use cases are narrow and possess a well-defined set of rules, machine learning algorithms can learn through “self-

play” and learn to complete goals on their own (Russell, 2019; MIT Horizon, 2022a). As a result, AI-enabled algorithms may achieve levels of performance far beyond their human counterparts in some contexts (Stumborg et al., 2019). As an example, DeepMind’s Alpha Zero taught itself to play games of chess, shogi, go, and Starcraft II and handily beats grand-master ranked opponents (Russell, 2019; Brose, 2020; Aycock & Glenney, 2021).

F. CHAPTER SUMMARY

This chapter presented critical concepts necessary to understand AI. Alongside the DOD’s intended use case for AI, some of the use cases for industry and private citizens were discussed. Next within this chapter, a discussion about the differences between AI and automation followed, where the authors assert that AI is a more capable form of automation. Following this, important algorithms and their traits were explored. This high-level discussion about popular algorithms provided the background necessary to discuss the practical limitations and benefits of AI technologies. Although the specific design and selection of algorithms is beyond the scope of this study, the discussion in this chapter is essential context for the following chapter, where the authors begin to explicate how humans and AI-enabled machines may work together in team constructs.

THIS PAGE INTENTIONALLY LEFT BLANK

VI. HUMAN-MACHINE TEAMING

Wielding an AN/ZPY-3 radar, an automatic identification system, as well as electro-optical and infrared sensors, the unmanned U.S. Navy MQ-4C Triton flew into the operational area ahead of the manned P-8A Poseidon multi-mission maritime patrol aircraft (Kreisher, 2016). The MQ-4C was an unmanned asset, designed for intelligence, surveillance, and reconnaissance tasks, and was sharing information with the manned aircraft, allowing the vulnerable P-8A to become more familiar with the vessels and contacts of interest well before the manned craft transited the area (Kreisher, 2016). In this example of HMT, an AI-enabled system is operating within a team to perform what could be a dangerous task—reconnoitering a contested area. The team’s machine partner possesses a suite of sensors and data links that complement those aboard the P-8A, allowing the team to complete vital surveillance tasks.

While the team formed by the MQ-4C and the P-8A are a suitable demonstration of HMT, there are many essential elements that are overlooked in the example. The roles, interdependencies, and tasks of HMT structures are one area that could become much more complicated than this simplistic example (NASEM, 2021). Additional important elements for HMT include trust, team biases, as well as explainability and transparency in the AI partner’s processes (NASEM, 2021). This chapter will provide an overview of HMT including what it means for AI to be a teammate, HMT methods and models, design principles, as well as the effects of AI on team performance.

A. WHAT DOES IT MEAN FOR AI TO BE A TEAMMATE?

Teams are a goal focused, interdependent group of agents, where each teammate is assigned their own roles and responsibilities (NASEM, 2021). Adapting this definition slightly, NASEM (2021) defined HMT as a team encompassing at least one human and one or more AI systems that coordinate and collaborate to achieve some goal. Given a common goal, the ability to act interdependently, as well as assigned roles and responsibilities, the AI may be considered a member of a team (NASEM, 2021). The different roles and responsibilities should be divided based upon the capabilities and

limitations of the members on the team, in effect creating task specialization (Stumborg et al., 2019; NASEM, 2021). For example, human partners excel in their ability to think strategically, place items and events into context, and orient their teams (Stumborg et al., 2019). Conversely the machine partner's strengths lie in the ability to observe in much greater detail than the human partner as well as the ability to quickly sift through large volumes of data for insightful patterns (Stumborg et al., 2019).

Regardless of the strengths of the different members of the team, they are still required to work together to accomplish a goal. The interdependencies between positions on the team will require communication and coordination mechanisms to ensure that members of the team remain synchronized as they make decisions and interact to solve problems (Cooke et al., 2004; NASEM, 2021). As members of the team communicate, coordinate, and interact with each other, these interactions will generate information about the team, its goal, and environment (Canan & Demir, 2021). This information becomes explicit, and may be consumed by teammates to improve their understanding of the situation and environment that they are operating within (Canan & Demir, 2021).

While the exchanges of information may be possible through verbal methods or information networks, communications also occur through implicit means such as non-verbal mechanisms. High performing human teams practice implicit communications when they must exchange information but experience difficulties with direct communications (Gorman et al., 2020; NASEM, 2021). NASEM (2021) notes that behavioral cues are invaluable in human-animal teams, but have not been effectively explored within HMT constructs. This implies that the processing of team communications is limited to verbal communications, text-based inputs, or some other form of graphical user interface (NASEM, 2021).

B. WHY HUMAN MACHINE TEAMS?

The previous chapter noted that AI systems may complement humans in dangerous, dirty, or otherwise dull operational tasks (Grosz et al., 2016). In the case of dangerous situations, Wang and Kosowski (2019) noted that machine partners may ultimately reduce the probabilities of human casualties. The example used in the introduction of this chapter

demonstrates that the Department of the Navy (DON) is already using machine partners in this capacity (Kreisher, 2016). As will be discussed during this chapter, machine partners may also enhance the situational awareness of human partners, minimize the volume of tasks assigned to humans, and improve the performance of the team (Wang & Kosowski, 2019).

A goal of research into HMTs is to identify scenarios and environments where the human-machine team outperforms an all-human team or all machine-team (DOD, 2020). According to Stumborg et al. (2019), this goal is achievable. HMTs outperform teams of humans or teams of machines in some contexts (Chakraborti & Kambhampati, 2018; Bansal et al., 2019; Stumborg et al., 2019). This means that it is possible to combine the speed of machines with the cognitive abilities of a human to bolster performance and overcome the limitations of human partners (Stumborg et al., 2019; NASEM, 2021; Chakraborti & Kambhampati, 2018). The cognitive intelligence of a human is currently required to lend context to the options and recommendations generated by a machine partner (Stumborg et al., 2019).

Russell (2019) thoroughly proves this concept and demonstrates how the preferences and contextual abilities of a human partner are essential in achieving goals set for a machine partner. In one example, Russell (2019) describes the efforts of a machine to generate an optimal solution to several problems by adjusting its input parameters. In the example, the machine fails to consider that turning the sky orange would be an unacceptable outcome for the human partners. Although an unrealistic example, it demonstrates that there are conditions that a machine partner will not consider in its recommendations, thus requiring consultation of human partners in decision cycles. This discussion reinforces Liu et al.'s (2021) finding that humans can complement AI partners in various ways in support of different tasks.

C. HUMAN MACHINE TEAMING METHODS AND MODELS

1. Processes and Characteristics of Effective Human Machine Teaming

Some tasks may require a group of people to work interdependently to complete them. Teams come together to coordinate and perform work in such cases, and their results

are often greater than the aggregation of individual efforts on their tasks (Cooke et al., 2004; Bansal et al., 2019; NASEM, 2021). The principle that the performance of an HMT will exceed the sum of its individual parts was similarly reflected in Chakraborti and Kambhampati's (2018) findings. However, simply placing humans and machines into a team will not make them effective, even if they are in complementarily designed roles. This is because the pairing humans and machines in an HMT construct requires "sufficient levels of team intelligence, including the processes, knowledge structure, and behaviors necessary to promote effective teamwork ..." (NASEM, 2021, p. 11).

NASEM (2021) describes that HMT require at least two agents to possess shared goals, be assigned specific roles, and to demonstrate some model of interdependence to accomplish their tasks. Team members should be able to make decisions grounded in the context and status of the task they are performing, and use the knowledge gained from their roles to advance progress towards the shared goal (NASEM, 2021). Mental models are also important to the team, as they guide its ability to work interdependently. Mental models enable the team members to effectively coordinate, and allow individual team members to predict the future needs and actions of their teammates (Bansal et al., 2019; NASEM, 2021). When the team has an accurate mental model, members will be able to place actions in appropriate contexts, and work within environmental and situational constraints (NASEM, 2021). These attributes are similar to those found in chapter three of this study, and are indicative that HMT may rely on a common set of characteristics and definitions found in human teaming.

Distributed environments may also pose a challenge for shared situational awareness in HMTs. Gutwin and Greenberg (2004) noted that teams perform best when there is some semblance of a shared workspace. A machine partner could find itself operating alongside human partners in a virtual or through a geographically distributed manner. In these contexts, the machine partner will still need to cohabitate with partners in a shared workspace, and must be able to perform tasks and coordinate in ways that bolster team performance. Gutwin and Greenberg (2004) observed that when humans executed operations in these conditions, they required tools and software that could assist team members in the collection of information that fed awareness, as was able to both transmit

and display the desired information. The use of these tools and software to support shared workspaces may also support effective partnerships between human and machines.

2. Structures, Roles, and Tasks

Members in team structures may have different capabilities, responsibilities, and authorities; this is no different in human machine teams (NASEM, 2021). AIs that are integrated into team constructs may have different attributes, functionality, capabilities, and responsibilities from their human counterparts (NASEM, 2021). Although both entities may be thought of as having their own agency, the difference in these factors implies that a machine may not be as capable as a human counterpart (NASEM, 2021). Humans and machines may perform their tasks differently, and may not use the same language or cues (NSAEM, 2021). NASEM (2021) asserted that machine partners may not be as capable as humans, however in an excellent comparison, noted that humans and animals have a similar relationship in their teaming constructs, and that human animal teams are effective in their use cases.

The distributed maritime environments that the Navy and Marine Corps are planning to operate in may pose a challenge to HMT constructs. It has been acknowledged that the C2 of disaggregated and distributed forces is a significant challenge for commanders (Kline, 2017). Schramm and Clark (2021) assert that a machine partner may be able to assist a human staff in the control of forces in these complex environments if designed to do so. A team of humans, such as a staff, could publish tasks, constraints, priorities, metrics, and objectives to teams comprised of humans and machines in an auction format (Schramm & Clark, 2021). Machine partners could be designed to bid on these tasks based on their suitability to meet the requirements, with winning bids resulting in task assignments to teams (Schramm & Clark, 2021). Such a tasking method could allow staffs to adapt to denied and degraded C2 environments, and would still allow them to dynamically allot available capabilities to meet mission requirements. While this may sound farfetched, auction algorithms are in use in high-visibility organizations such as Uber, Lyft, and the U.S. Federal Communications Commission (Schramm & Clark, 2021).

The field of healthcare analytics offers some examples for comparison to human-machine teams in defense contexts. In healthcare analytics there have been recent efforts to move beyond the automated image classification of objects towards providing insightful analysis about the state of the object being evaluated (Grosz et al., 2016).

The machine's recommendations are still subject to the scrutiny of a human and are governed by a restrictive framework of regulations (Grosz et al., 2016). Gombolay et al. (2016) successfully demonstrated how a machine partner could help team members better schedule tasks and make decisions in medical contexts. More specifically, the AI optimization tool, named Nao, assisted doctors and nurses in a maternity ward with the scheduling of care for their patients (Gombolay et al., 2016). Astonishingly, Nao's optimization skills were qualified as "good decisions" by healthcare professionals approximately 90 percent of the time (Gombolay et al., 2016).. This collaborative teaming dynamic is an interesting parallel to the human-machine teaming concept in this study.

Although currently beyond the capabilities of most systems, there is an ardent desire for machine partners to be able to place observations and actions into situational contexts (Thilmany, 2007). This implies that such machine partners will be constrained by a complex framework of ethical considerations, policies, and rules of engagement (DON, 2021; Singer, 2009; Scharre, 2018).

Sycara and Lewis (2004) noted that teams are often comprised of personnel that possess unique specializations, and as a part of their roles, have differing responsibilities. This implies that a multi-role machine will be required, and that the design of such an agent will be quite complex. To overcome this limitation, Sycara and Lewis (2004) suggest the deployment of several narrow-purpose machines that could respond as needed to fill gaps within a team. This would allow for a dynamic, tailored approach to a team's task, environment, and structure (Sycara & Lewis, 2004). This construct could foster agility within the team, albeit with increased communications and coordination costs (Sycara & Lewis, 2004).

Fluid and meaningful interaction between teammates and maintaining interactive decision-making is required for teams (Canan & Demir, 2021). Each member of the team

is responsible for specific tasks that contribute to the overall success of the team. When a team is comprised of only machines, the same dynamics apply. Communication dynamics between machines are different as machines lack intuition when measured against the natural capabilities of a human (Russell, 2019). Thus, understanding how machines effect teamwork in a team construct is essential for understanding a team's communication processes (Canan & Demir, 2021). It is also important to understand that although communication is a vital part of team interactions, over-communication can negatively impact team effectiveness as well. As Canan and Demir (2021) pointed out, communication dynamics are indispensable and have several ways and means available for improvement.

Machines may operate in team constructs alongside other machines. An example of this deployment is automated swarming (Scharre, 2018). While some conceptual elements are common between teams of machines and teams of humans, there are some unique considerations to teams of machines. Machines will require the use of a teamwork model that is shared amongst all systems in the team (Sycara & Lewis, 2004). The ultimate shared goal of the team, future plans to achieve that goal, team communications protocols, as well as communications policies must be accounted for (Sycara & Lewis, 2004). Akin to human teams, the machines must also possess models of themselves as well as other agents deployed in the team (Sycara & Lewis, 2004). With these elements accounted for, Sycara and Lewis (2004) assert that teams of machines may match their capabilities to the specific roles identified within a particular plan. Machines may communicate with each other to inform members of commitments to a goal, the progress of tasks as they relate to the ultimate goal, as well as any barriers (Sycara & Lewis, 2004). When obstacles are encountered that prevent a machine from accomplishing the goal that it registered for, the machine may communicate a decommitment with other members of the team to indicate that the goal is not achievable (Sycara & Lewis, 2004).

These machines are communicating to build a common level of shared task understanding to improve the performance of the team. As Sycara and Lewis (2004) noted, it is easy for a machine to identify its own limitations, challenges, plans, and goals. Developing shared task understanding in a team requires evolving beyond the individual

level and accounting for the goals, plans, and actions of other team members, which is done through the use of models (Sycara & Lewis, 2004). When such models are properly developed and implemented, they will allow machine partners to identify when plans and tasks have been re-sequenced, halted, cancelled, or have otherwise failed (Sycara & Lewis, 2004).

Automated swarm teams have demonstrated that machine teams are capable of demonstrating complex behaviors and presenting problem solving capabilities from a given set of rules (Diukman, 2012). Swarm concepts have piqued military interest as swarm attacks have enabled flexible, robust, and adaptive strategies while minimizing probabilities of friendly force casualties. Understanding the communication dynamics of machine teams provides an insight into the effects of a machine teammate embedded in with small human teams (Diukman, 2012). The addition of a machine teammate will enable rapid, precise decision making in teams that have traditionally been solely comprised of humans.

3. Human Machine Team Interactions

One of the key tenets that underpins the Marine Corps' preferred warfighting doctrine of maneuver warfare is "commander's intent." This statement is the commander's personal vision of the purpose of the operation, and is designed to provide subordinates with context needed for independent operations (HQ USMC, 2017). This concept of intent is ubiquitous in team constructs that leverage mission tactics, but is a limitation in human-machine interactions (Sycara & Lewis, 2004; HQ USMC, 2017; Russell, 2019). Humans have challenges in conveying their intentions to machines, meanwhile, machines struggle to make their outputs decipherable by humans (Sycara & Lewis, 2004; Fout & Ploski, 2018). This could be a significant challenge in the HMT construct, where the machine is envisioned to occupy roles as value-added members of a team.

Chakraborti and Kambhampati (2018) noted that an AI must be designed to understand the humans that it will be interacting with. It is asserted that a machine partner will need to model the beliefs, desires, intentions, preferences, and expectations of its human partners (Chakraborti & Kambhampati, 2018; Russell, 2019; NASEM, 2021). This

aligns with the body of work for team mental models and makes intuitive sense; members of human-human teams automatically model these attributes when they are interacting with their human teammates (Chakraborti & Kambhampati, 2018).

A machine partner will need the ability to understand the human's mental models and capabilities when it communicates its own intentions and the justifications for its decision rationales (Chakraborti & Kambhampati, 2018; NASEM, 2021). An explainable AI capability could help to resolve mental model mismatches between humans and their machine partners, however Liu et al. (2021) noted that these reconciliations primarily occur from the machine partner to the human. This implies that the machine partner will be communicating to the human, who may or may not be able to participate in a constructive conversation with the machine partner to resolve inaccuracies and uncertainties. Mental model mismatches may be one cause for a machine partner to explain its rationale and attempt to update the human partner's mental model to that of the machine's (Chakraborti & Kambhampati, 2018). This begs the question: What if the machine's mental model is inaccurate?

While this challenge seems easy to overcome, human partners may struggle to overcome inaccurate mental models. Liu et al. (2021) note that one-way descriptions flowing from the machine to the human are seldom sufficient for humans to understand an AI's predictions. Interactions between humans and machine partners were demonstrated to be more effective than one-way communications (Liu et al., 2021). This implies that partners in the HMT will come to some kind of a consensus through these interactions. Liu et al. (2021) found that this interactive consensus building led to problems with over-trust of the machine partner's capabilities. Liu et al. (2021) also noted that when more information is provided to human partners, these humans more likely they are to agree with machine partners, especially when the AI's recommendations align with a human partner's biases. Team performance could be adversely impacted if a human partner begins to trust the machine partner's flawed recommendations based on inaccurate mental models. Bray and Moore (2021) asserted that machine partners will be able to perform autonomously as algorithms become more explainable and as confidence grows in their capabilities. While human partners have found that the addition of labels, confidence intervals, and interactive

communications with partners are useful in facilitating further interactions, these ultimately build undue confidence in a machine's capabilities, recommendations, and decisions (Liu et al., 2021).

This undue confidence in a machine's capabilities introduces an additional area of systems design. Human partners must have meaningful control over the systems that reside within their teams. Bode and Watts (2021) studied the integration of autonomous weapon systems, or weapons with automated features, and how these capabilities have altered human control. Bode and Watts (2021) noted that humans are unable to keep pace with machines because of the speeds that they operate in, the tasks that they perform are complex, and the stresses placed on the human from the system, environment, and situation. Stresses are placed upon human operators as they must manipulate the human system interface to make sense of their environment and trigger the desired actions. When the operator must break their focus to attend to a system notification, the cognitive load placed upon them due to sensemaking activities increases.

These factors effectively reduce a human's ability to control such systems and make useful targeting decisions, reducing the human member to a passive supervisor as opposed to an active one (Bode and Watts, 2021). Although their research focused on air defense, their research is generalizable because these factors align to conditions that users may experience during other defense use cases. Bode and Watts (2021) asserted that there are three requirements for humans to exercise meaningful control over machines: they must understand how the system functions and arrives at decisions; the system must be capable of achieving sufficient understanding of its operational contexts; and the human must be able to scrutinize decisions made by the system.

Human partners in team constructs are able to intuitively grasp the coordination and interaction dynamics as they work towards the completion of a goal. Conversely, a machine's ability to intuit progress towards a goal are dependent upon its model, which will likely require additional communications and interactions with teammates to remain valid (Sycara & Lewis, 2004). Achieving the minimum expected performance required to operate alongside human partners, such as performing different roles to fill positional gaps, as a value-added member of a team will require significant amounts of programming, and

will likely require the use of sophisticated algorithms designed to imitate functions of the human brain (Sycara & Lewis, 2004; Domingos, 2015). The programming of the machine partner must be done in such a way as to allow the development of models that can account for the tasks of its partners. These models would allow the machine partner to make inferences about team member intentions (Sycara & Lewis, 2004). If such intentions could be modeled and understood, the machine partner could proactively identify, organize, and store information that may support upcoming efforts as a means to improve team performance (Sycara & Lewis, 2004).

A limiting factor in HMT will be the speed that humans and absorb information flows and changes in the environment (Domingos, 2015). Feedback mechanisms could be used to indicate to a human partner that new information is available for consumption, or that an action has corresponded to a change in the interface or larger environment (Gutwin & Greenberg, 2004).

4. Design Principles Specific to Human Machine Teaming

In its review of HMT, NASEM (2021) noted that these teams need to be designed to support bi-directional communications about the team's goals and the larger purpose behind these goals. There must also be some level of explainability or transparency behind the decisions and actions of machine partners (NASEM, 2021). Unsurprisingly, NASEM (2021) asserted that there must also be operator-directed interfaces that could be used to guide the machine through dynamic periods should it become necessary. This recommendation mirrors that of Fout and Ploski (2018), who assert that such an interface would be essential to human-machine communications.

The implications of these requirements are that the machine partner must be able to understand, communicate, and interact with its human partners (Chakraborti & Kambhampati, 2018; NASEM, 2021). This implies that the machine may need to be capable of using natural language to build shared awareness, assist with planning, and to direct the human elements of a team (NASEM, 2021). With these capabilities, a human will be able to exercise meaningful levels of control over a machine partner (NASEM, 2021). The concept of meaningful levels of control could be distilled to the idea that

humans must be able to understand what a system is telling them, use these inputs to develop sufficient situational awareness, interact with a system, understand as well as predict the outputs of these interactions, and change them if required (NASEM, 2021, Bode & Watts, 2021). As an example, Bode and Watts (2021) described situations where the roles of humans have been minimized, and were unable to exercise an adequate level of control over a system. Bode and Watts (2021) argue that a human taking indications from a machine and then initiating actions to fire weapons systems without inputs into the rest of the targeting and engagement process is insufficient. Examples of this concept can be found in the MIM-114 Patriot engagements of friendly aircraft during the 2003 Gulf War (Bode & Watts, 2021).

Generating agility and performance in human-machine teams depends on identifying and selecting an optimal set of human and machine attributes (Stumborg et al., 2019). To achieve the desired end-state, optimization of these attributes may prove necessary. Quantifying and calculating the highest possible scores of desired human and machine attributes is an optimization problem that the machine partner is uniquely suited to execute (DOD, 2020). One example of a desirable attribute could be easily understood and adjusted confidence intervals. Stumborg et al. (2019) note that tunable confidence intervals may allow for dynamic control mechanisms that suit different risk appetites. This will be discussed in greater detail in the following sections. Additionally, Fout and Ploski (2018) asserted that human-machine teams should exemplify interdependence models oriented around common goals. Such interdependence between the humans and machines could be achieved by designing machine teammates with observability, directability, and predictability attributes (Fout & Ploski, 2018). Similar to Fout and Ploski (2018), the authors see benefit in applying Johnson's (2014) definitions of observability, predictability, and directability.

Observability is particularly important to team dynamics and will be defined as "making pertinent aspects of one's status, as well as one's knowledge of the team, task, and environment observable to others" (Johnson, 2014, p. 68). Predictability is the requirement that teammates actions should be predictable enough such that "others can reasonably rely on them when considering their own actions" (Johnson, 2014, p. 68).

Finally, directability will be defined as “one’s ability to direct the behavior of others and complementarily be directed by others” (Johnson, 2014, p. 68). These coactive design principles will be essential to support decision cycles, and will be discussed in the following section.

Bansal, Nushi, Kamar, Lasecki, Weld, and Horvitz (2019) discussed several design principles that may improve the performance of HMT. The study sought to evaluate the human partners’ mental model of the machine that they were integrated with (Bansal et al., 2019). The findings noted that machines designed with parsimony, non-stochasticity, and dimensionality enabled human partners to make better decisions in certain scenarios (Bansal et al., 2019). Ultimately, these design principles enabled human partners to understand the machine’s error boundaries and allowed humans to identify when the machine partner was in error or was likely to generate an error (Bansal et al., 2019). In the Bansal et al. (2019) study, parsimony was defined as how simplistically an error boundary could be represented. Non-stochasticity was defined as the number of features needed to effectively model the error boundary (Bansal et al., 2019).

Finally, dimensionality was defined as the number of features needed to describe the machine’s error boundaries (Bansal et al., 2019). It is worth noting that these design principles articulate that machines who require additional features, or areas of complexity, to describe the error boundaries will likely introduce model output interpretation challenges. Such challenges may complicate a human partners’ ability to understand when the machine is operating outside of its limitations and error boundaries (Bansal et al., 2019). This makes intuitive sense, as humans have demonstrated an inability to adequately understand how the interdependence of variables, elements, or features correlates to specific outputs in complex systems (Luhrsen, 2007).

It is recommended that AI systems, such as machine partners, be designed with parsimonious error boundaries, minimal stochasticity, limited task dimensionality, and the ability to integrate error-boundaries with previous versions developed by the model (Bansal et al., 2019). These principles are postulated to improve the performance of HMT assisting both partners in accounting for the weaknesses of the other actor. Bansal et al. (2019) assert that the success of HMT will depend on a human partner’s ability to trust the error-

boundaries of the machine partner. When human's mental model of a machine's error boundary is incorrect, the human is more likely to fail to predict when the machine partner will make an error and then fail to override the AI (Bansal et al., 2019).

For teams to be effective, its members must be able to trust each other. Mathieu et al. (2010) noted that team members are more likely to ask for and accept assistance from competent teammates that have excess work capacity. Fout and Ploski (2018) found that a machine would only be acceptable as a teammate if it could be trusted to perform tasks at least as effectively as its human teammates. Conversely, humans must also act according to how their machine teammate expects (Fout & Ploski, 2018). This becomes less problematic when human-machine teams are designed with observability, directability, and predictability attributes. An additional characteristic that could be of value in human-machine teams is Stumborg et al.'s (2019) concept of thick data. Thick data is described as "... the contextual information that puts a reality check on the solutions proposed by a machine teammate" (p. 29). Fout and Ploski's concept (2018) could be improved upon by allowing human teammates to incorporate these parameters so that machine partners do not waste decision cycles crafting solutions that do not meet situational constraints (Russell, 2019; Stumborg et al., 2019). Essentially, human and machine partners could improve decision cycles through experience operating with other and learning the "preferences" of their partners (Russell, 2019).

A critical challenge in operating in a human-machine team will be the tendency of human partners to implicitly trust machine partners without first validating that the machines are able to reconcile "the drift between the entity being modeled and the models..." in their environment (Stumborg et al., 2019, p. 36). Stumborg et al. (2019) note that this concept of adaptability is important as the actors and environments being modeled change over time, mandating that the team re-evaluate data, models, and assessments. Humans that interact with machine partners must understand the operating conditions that the system was designed for as well as the system's performance envelope (Lange & Carreno, 2021). Merely understanding a system's design specifications and performance envelope is insufficient; humans must have sufficient feedback mechanisms to understand how the system is functioning in relation to its limitations (Lange & Carreno, 2021).

The human tendency to surrender their personal control and defer to a machine is a frequent occurrence and has been the cause of operator fatalities in several instances (Fridman, 2018; Russell, 2019; Scharre, 2018; Shalvey, 2021; Burke, 2022; Posada, 2021). Despite limited use in combat operations, humans have already been observed anthropomorphizing their machine partners (Singer, 2009; Wallach & Allen, 2009). This habit of implicitly trusting machines could prove dangerous, especially in critical situations (Zhang, McNeese, Freeman, & Musick, 2021). Unwarranted trust may enable human partners to offload the wrong kind of tasks to a machine. Researchers studying the issue have noted that under-loading a human partner could be as harmful to performance as overloading them (Lange & Carreno, 2021; NASEM, 2021). Lange and Carreno (2021) argue that humans must remain engaged in the operation of a machine partner to maintain sufficient situational awareness and to make informed decisions. This makes intuitive sense; uninvolved team members will need to spend time building situational awareness about the interactions between the team and its environment to best fill in the gaps (NASEM, 2021).

Since armed autonomous capabilities are already able to identify hostile actions and automatically engage perceived hostile actors, failing to remain engaged and trusting the machine could result in undesired outcomes (Scharre, 2018; Bode & Watts, 2021). The safety and ethical considerations, as well as the ordered effects, of engaging such hostile actors may be beyond the capabilities of a machine partner (Scharre, 2018; Wallach & Allen, 2009). The decisions and considerations for engaging hostile actors are familiar to human partners, many of whom receive detailed rules of engagement guidance and ethical training for such cases. Humans will likely need to support a machine teammate by applying this framework of training and guidance when a machine partner is operating near the edge or beyond its performance envelop.

When Chakraborti and Kambhampati (2018) asserted that machine partners will need to model human partners' beliefs, desires, intentions, preferences, and expectations of its human partners to build suitable mental models, they overlooked the notion that while humans are capable of performing this modeling, it can often be a significant struggle. Christensen et al. (2020) noted that machines are currently much worse at building these

models than humans are. Human behaviors may be variable, their motivations non-obvious, and may even seemingly be contrary to goals (Russell, 2019; Christensen et al., 2020). Machine partners will need access to data that describes human partners, team roles, and interaction scenarios (Christensen et al., 2020). Merely having access to this data is not enough; machine partners will need an array of tools and frameworks to analyze data and operate in ways that contribute to team performance (Christensen et al., 2020). Christensen et al. (2020) estimated that this is a long-term endeavor made more difficult by the unrealistic expectations of machine capabilities. Arguably, the DON's envisioned use cases and framework of guidance outlined by Tangredi and Galdorisi (2021) are guilty of this.

As of the time of this study, there are no legal frameworks that place limits on the violence that autonomous systems can inflict on humans in combat scenarios (Scharre, 2018). Similarly, there are struggles with conveying fundamental warfighting concepts to machines. More specifically, how do designers impart the "laws of warfare" to autonomous systems? How do they define what proportionality and ethical treatment mean to a machine? (Singer, 2009; Scharre, 2018; Domingos, 2015). Current DOD policy requires commanders and operators to exercise appropriate judgement over the use of force, and they must remain involved in decisions regarding the employment of such systems (NASEM, 2021). A poignant example comes from Scharre's (2018) work. During Scharre's (2018) deployments to Afghanistan it was common for adversaries to use civilians or fighters wearing plainclothes to conduct reconnaissance and spotting operations. During one particular operation, Scharre's (2018) team was nearly compromised by a child who encroached on the team's hide site. Had the child been spotting for the enemy, this would have been a violation of the rules of engagement. The human team members decided to exercise restraint and avoid an engagement. Scharre (2018) asserts that this scenario may have a different outcome if machines are imparted with the ability to understand rules of engagement without larger situational contexts.

As humans and machines continue to partner in team constructs, there will be another challenge that emerges. How do we deal with the idea of "white lies" that may happen between humans and machines in HMT? Chakraborti and Kambhampati (2018) designed an algorithm capable of manipulation and deception, and noted that human

partners were decidedly split in their reactions to machine partners that lied to them in critical situations. While some humans were strongly against being lied to, others were agreeable to the premise so long as it positively improved team performance (Chakraborti & Kambhampati, 2018). Interestingly, humans generally did not see any issues with deception or manipulation of machine partners so long as it was for the “greater good” of team performance (Chakraborti & Kambhampati, 2018). Deception and manipulation of partners has been generally acknowledged to be an example of poor team citizenship behaviors in the traditional human-human teaming construct (Chakraborti & Kambhampati, 2018). While these behaviors could improve the performance of HMTs, operating alongside machine partners capable of such behaviors could prove challenging for human partners that will need to trust machines in critical situations.

D. EFFECT OF AI ON TEAM PERFORMANCE

One limitation of AI is that it is unable to identify specific causes and effects, especially in complex systems (Aycock & Glenney, 2021). This could limit an AI partner’s ability to develop suitable recommendations for human partners. This may result in the AI providing humans with ranges of possible actions as opposed to specific recommendations that are tailored to match its operational environment (Aycock & Glenney, 2021). Many humans may be limited in their ability to understand cause and effect chains in complex systems (Senge, 2006). This means that AI may only marginally improve decision-making in HMTs facing novel or dynamic situations (Aycock & Glenney, 2021). One of the fundamental assumptions of HMT is that pairing humans and machines together in team constructs will improve the decisions of human partners due to the superior analytical capabilities of a machine partner (Stumborg et al., 2019). This marginal improvement of decision-making runs counter to the DON’s goal of increased confidence in cause and effect relationships (DON, 2021b).

E. CHAPTER SUMMARY

This chapter explored some of the pertinent HMT design principles and considerations. Several important findings emerged from this review that may prove insightful to HMT. It has been asserted that teams of humans and machines performing

together in an integrated HMT construct will be able to perform more effectively together than either element operating in isolation (Stumborg et al., 2019; NASEM, 2021). This assertion relies on the ability of the human partner to make sense of and predict the behaviors of its machine partner (Bansal et al., 2019). Another important factor within HMT will be the trust relationship with its machine partner (NASEM, 2021). The ability to make accurate decisions based on the insights of machine partners is similarly important (Bansal et al., 2019; Chappell, 2020). Finally, it may prove important to exercise meaningful levels of control over machine partners (Bode & Watts, 2021; NASEM, 2021).

VII. METHODOLOGY

A. VIGNETTE DEVELOPMENT

1. Operation Provide Comfort: AWACS, Black Hawks, and F-15 Crews

The researchers will attempt to develop insights from case studies involving operational use of machine capabilities in team contexts. On April 14, 1991, two UH-60 U.S. Army Black Hawk helicopters were mistakenly shot down by two U.S. Air Force F-15C fighter aircraft. The incident was heavily investigated by military and civilian personnel alike to learn from and correct the mistakes made on that fateful day. This study aims to further contribute to that effort through the lens of applying HMT interaction. The authors will examine the existing data and information gathered through the investigations and explore how HMT and team cognition could be applied to future scenarios. Although the background and previous events will be discussed, the primary focus of the analysis will be of the roughly one hour of team interactions preceding the actual shutdown of the two Black Hawk helicopters. The goal is not to state whether HMT team cognition would have prevented the accident but how it could be applied to reduce miscommunication during team interactions. Accidents will always be an operational hazard; however, the authors aim to understand and report on how HMT can positively influence team cognition. The specific details regarding the Operation Provide Comfort incident will be described in further detail in subsequent chapters.

2. USS Vincennes Shootdown of Iran Air Flight 655

The shootdown of Iran Air flight 655 by the USS Vincennes in 1988 has been thoroughly studied. This study contributes to the volumes that have been written by examining the events of the shootdown through an interactive team cognition lens. This unique perspective accounts for the information and actions available to the different teams that were present during the event. The authors will use these well-studied interactions as a starting point for the insertion of machines into team constructs. The authors narrowed the scope of the vignette to focus on the air defense interactions that occurred during the seven minutes between 1017L and 1024L on 3 July 1988. Thus, the complex operational

environment and details regarding the surface engagement factors are largely removed from the vignette. This does not imply that they are unimportant or irrelevant, but that including them in the scope of the analysis would not provide additional insights into the performance of the team. The pertinent details that influenced the combat information center crew of the USS Vincennes are described in the following chapters.

B. APPROACH

1. Conceptual Systems Dynamics Models

The models discussed during this study were developed through a comprehensive examination of computer science, cognitive psychology, and military science literature. This multi-disciplinary analysis focused specifically on artificial intelligence, interactive team cognition, military C2, as well as DOD technology strategies. These topics were framed against specific, well-documented case studies to enable the authors to develop insights regarding the operational use of AI capabilities in team contexts. This qualitative research identified prominent areas of intersection that would benefit from systems dynamics models.

One of the foundational assertions of systems theory is that something with two or more related elements is a system (Luhrsen, 2007). These systems can be mapped through a series of related relationships that may be decomposed into inputs that feed processes and ultimately result in outputs (Senge, 2006; Luhrsen, 2007). Although causes may be hard to identify and behaviors of such systems can be difficult to observe, the elements of the system dynamically interact to exchange matter, information, or energy (Senge, 2006; Luhrsen, 2007). In some cases, the underlying elements of a system have the capacity to change or adapt, modifying their structure. Such systems are called complex dynamic systems. Complex dynamic systems can be described by their structure, interconnectedness, adaptations, inputs, processes, and outputs (Luhrsen, 2007; Allspaw, 2015). Although these elements are abstractions of phenomena, they allow users to diagnose problems, analyze behaviors, and model the effects of changes (Senge, 2006; Sterman, 2000).

Especially on the battlefield, “a military action is not the monolithic execution of a single entity, but necessarily involves near-countless independent and inter-related decisions and actions being taken simultaneously throughout the organization” (HQMC, 1997, p. 13). As these organizations act, their actions in one area could have consequences in across the battlespace (HQMC, 2017). Thus, the actions of actors, complex adaptive systems themselves, on the battlefield trigger responses from other actors, demonstrating that “everything is connected to everything else” (Sterman, 2000, p. 4). The actions that give rise to complex adaptive behavior continuously alter the environment, leading to additional patterns of goal and environmental assessment, problem framing, decision-making, actions, and the observation of results (Sterman, 2000). The interactions between these two agents are generating additional complexity and pushing the system into disequilibrium with the express goal of manipulating a target system to change into a desired state (Luhrsen, 2007). Thus, adopting a systems dynamics approach is appropriate given the warfighting contexts of the vignettes used in this study.

Using a conceptual systems dynamics model for human machine teaming, this study will explicate employing AI-enabled machines into team constructs. This approach will explore the impact of narrow AI machines and their effects on communication, coordination, and interaction dynamics for members operating in a team construct. As identified in chapter three, there are many studies that examined the communications, coordination, and interaction dynamics of human-human teams and the relationship between those dynamics to team agility and performance. However, there is a lack of knowledge about such behaviors and cause-effect relationships that comprise the behavior of the human-machine team as a dynamic system. Akin to Khan, Shiwakoti, and Stasinopoulos (2022), the authors assess that the unintended consequences for deploying AI-enabled machines inside of manned formations as teammates represent significant technological challenges and complexity. A systems dynamics model accounts for short and long-term behaviors in modeled systems, allowing the authors to navigate the costs and benefits to teams stemming from the deployment of AI-enabled systems into team constructs (Senge, 2006; Borges et al., 2021). Essentially, the emergence of human-machine teams has created a significant gap in the field that warrants exploration.

A causal loop diagram model is an effective means to capture the dynamic performance of teams of humans and machines at the team level. Specifically, systems-dynamics based on causal loop diagram approaches have been used in the past to model specific aspects of human machine teams (Wang & Kosowski, 2019; Borges et al., 2021; Khan et al., 2022). It is especially well-suited given the non-linear behaviors demonstrated by teams, and has been both analyzed as well as validated against many complex systems ranging from business supply chains to the cybersecurity ecosystem (Sterman, 2000; Khan et al., 2022). As noted by Khan, Shiwakoti, and Stasinopoulos (2022), this qualitative method can easily support quantitative analyses once the required data is available. Although there is currently a shortfall in the amount of quantifiable data in this area, the authors believe that the increasing emphasis on human-machine teams will lead to the generation of quantitative data. In fact, NASEM (2021) asserted that there are more than 150 areas in human-machine teaming that require further study. This indicates that volumes of data could become available in the coming years.

A causal loop diagram-based approach is further supported by the many components of human machine teaming. Specifically, communications, coordination, interaction, agility, and team performance are all highly interconnected, non-linear, and depend on feedback loops. Therefore, changing any one of the variables in any of the causal loop diagrams can result in unpredictable team performance as well as unintended consequences. Using the USS Vincennes vignette as an example, altering the structure of the combat information center team by assigning new roles to its members subtly alters the communications, coordination, and interaction dynamics, which led to a breakdown in performance as well as the engagement of a civilian airliner (Dotterway, 1992).

Further, teams are limited by the performance of their “weakest link.” When one agent becomes task-saturated or experiences information overload, there could be a cascade of unpredictable effects that may impact the performance of the team. It is believed that communication, coordination, and interaction dynamics will be of importance to human machine teams, however, these concepts may be difficult to quantify in human machine team constructs because they are not yet widespread in the defense enterprise

(Stumborg et al., 2019). A systems dynamic model may allow the authors to explore the relationships between these team dynamics despite the absence of quantitative data.

Finally, although the deployment of AI-enabled systems into human team constructs is occurring, there is still a great amount of uncertainty about the attributes that characterize an effective HMT. Intangible concepts such as trust, team interdependence, mental models, and commander's intent may need to be programmed into machine partners (Johnson, 2014 ; Chakraborti & Kambhampati, 2018; NASEM, 2021). Given the novel deployments of these human machine teams constructs, the risks these intangibles pose are blended with traditional organizational and technological risks. This uncertainty can be modeled through a qualitative systems dynamics approach, and may indicate some of the most pertinent variables and conditions that improve or detract from team performance.

2. System Dynamic Approach Using Causal Loop Diagrams

A causal loop diagram model is a tool commonly used in systems dynamics because it allow users to synthesize, comprehend, and graphically depict the behavior and relationships of complex, non-linear systems (Khan et al., 2022; Kim, 2000). Causal loop diagrams illustrate the relationships between variables present in the model, which allow users to conduct both analysis and informed discussion (Khan et al., 2022). These relationships manifest through a series of feedback loops that identify the underlying behaviors of the system under study (Sterman, 2000; Khan et al., 2022; Senge, 2006).

The feedback loops used in this study will be described as reinforcing or balancing loops (Senge, 2006). Reinforcing loops indicate that there is an amplifying, or self-reinforcing relationship present between variables. These loops compound growth in one direction with even more growth (Kim, 2000). As an example, increased information reporting to an already overwhelmed anti-air warfare coordinator further contributes to information overload. Balancing feedback loops are the inverse. Balancing loops exist when goal-oriented behaviors are present (Senge, 2006). These balancing processes attempt to maintain some goal or target equilibrium, and can be thought of as self-correcting loops (Sterman, 2000). Expressed differently, balancing loops resist additional increases in one direction (Kim, 2000). As an example, the number of simultaneous

communications that may exist on a half-duplex single channel radio net in a combat information center is fixed. When multiple messages must go onto the medium, these messages must be balanced in accordance with priority, system capability, as well as open space on the medium.

Delays are the final building block that are also modeled in systems. Delays are defined as the separation of an action and its consequences (Senge, 2006). Unrecognized or unobserved delays can lead to instability and process breakdowns, as well as goal overshoot (Senge, 2006). These effects are more prominently felt when there are longer delays between an action and its consequences (Senge, 2006). When one does not perceive a response to an action, the action is often performed until a response is observed (Senge, 2006). The challenge is that the behaviors of systems with delays are usually not consistent, simple, nor predictable (Kim, 2000). As an example, when the captain of the USS Vincennes orders the engagement of a hostile aircraft, there is a time lag before the aircraft is engaged. The crew in the combat information center must turn the firing key, the missile must fire, and then the missile must transit to its target. Assuming a successful engagement, the delay between action (engagement) and consequence (destruction of target) may take several minutes. Combined, delays as well as reinforcing and balancing loops represent the fundamental building blocks of system archetypes (Senge, 2006).

3. System Archetypes

Accurately depicting the behavior of a complex system is difficult (Sterman, 2000; Senge, 2006). According to Sterman (2000), math and data-oriented fields have struggled with satisfactorily explaining the behavior of complex systems; similarly, professionals facing complex problems often have limited cognitive capacity or time pressures and simplify their problems with uncomplicated, intuitive mental models. Generally speaking, humans do not handle even simple dynamic complexity well because of a tendency to rely on heuristics, underestimate delays, ignore feedback loops, and misinterpret the conditions of stocks and flows (Sterman, 2000). As demonstrated by Luhrsen (2007), the systems that warfighters interact with are seldom simple. Warfighters will be interacting with complex systems fraught with time pressures, conflicting resource demands, and under extreme

conditions—this is a significant problem when one considers that most humans’ mental models drift further from reality as dynamic complexity increases (Luhrsen, 2007; Sterman 2000).

Thankfully, there are certain patterns and structures that commonly occur in the world. These generic structures are called archetypes, and they suggest that many problems that are perceived to be unique may actually be analogous to a relatively small number of these patterns (Senge, 2006). These archetypes may be considered a starting point, and could be leveraged to diagnose problems as well as suggest solutions that attack the fundamental root of the problem (Kim, 2000; Senge, 2006). For example, after the symptoms of a problem emerge, people spring into action and implement fixes that they believe will address the problem. Eventually, a series of unintended consequences will emerge, which usually prompts another bout of problem framing and quick fixes (Kim, 2000). Unless the underlying source of the problem is treated, people are merely treating the symptoms, and the insidious problem remains (Kim, 2000; Sterman, 2000; Senge, 2006). This is a common systems archetype called “fixes that fail,” and it describes how efforts to cure the symptoms of a problem often exacerbate the problem, causing symptoms to worsen over time (Kim, 2000). Professionals usually find themselves fixing the unintended consequences of previous “fixes” until the root cause is addressed (Senge, 2006).

Both case studies present the authors with organizational structures, the roles and responsibilities of team members, as well as the adaptations exercised by the teams in pursuit of their goals. There are also numerous inputs, processes, and outputs present in each situation. This historical information constitutes data that can be applied to qualitative systems dynamics models. More specifically, each of these cases appears to demonstrate systemic behaviors that align to common archetypes. These archetypes will enable the authors to diagnose breakdowns in team processes and behaviors, and will inform the proposed insertions of AI into these teams.

C. MODEL DEVELOPMENT

How does one diagnose the fundamental problem that led to the shutdown of an airliner and the loss of 290 souls? How did the teams that were supposed to enforce a no-fly zone end up shooting down friendly helicopters on a peacekeeping mission? These incidents exposed significant problems with teaming behaviors in complex and dynamic situations. Using these two incidents as the scene for a thought experiment, the USS Vincennes incident and Operation Provide Comfort allow the authors to model the performance implications of human machine teams in very-well studied contexts. Because models are a useful tool for examining complex environments and diagnosing fundamental problems, models could demonstrate how to integrate human machine teams in ways that could mitigate team performance challenges (Sterman, 2000). The models in this study account for the organizational and social contexts, and thus are a suitable tool to influence the design of capabilities, teams, and organizations (Sterman, 2000).

This study uses causal loop diagrams to illustrate the relationships that exist between the different variables in each vignette (Kim, 2000). This approach makes systemic structures discernable and enables the authors to analyze and discuss systems without quantitative data (Khan et al., 2022). Once the authors of this study depicted key scenarios as causal loop diagrams, feedback loops were identified that inhibited team performance. Due to the fact that feedback loops provide a holistic view of key parameters and relationships in the entity that they model, it will be possible to identify specific team behaviors and processes in each scenario, and describe how their contributions led to both tragedies (Khan et al., 2022).

Because the authors hypothesize that poor team performance directly contributed to the outcomes in each vignette, team processes and behaviors will be the focus of the modeling efforts. The conceptual models depict interconnected attributes that influence team performance. Team performance is ultimately shaped by the understanding of a goal, shared mental models, as well as the communications, coordination, and interaction dynamics (Mathieu et al., 2005; Mohammed et al., 2010; Gorman et al., 2020; Marks et al., 2002; Canan & Demir, 2021; van den Oever & Schraagen, 2021). Each of these elements may predominate in different situations, thus it will be difficult to infer which

attributes will be most important to human machine teaming due to a lack of empirical data in the field. However, because each of these variables are dynamically related, they may have important linkages to a team's agility and performance that may not be immediately obvious without a conceptual model.

Inputs, processes, and outputs are often used to model team performance, and as discussed in this section, they are also suitable for modeling dynamic systems (Sterman, 2000; Espinosa et al., 2004). As an example, the input elements of the model may be comprised of task, team, and contextual factors that drive work arrangements, task dependencies, and the sharedness of mental models (Espinosa et al., 2004; Lim & Klein, 2006). The different dependencies are then managed by teams through a variety of disparate coordination mechanisms (Espinosa et al., 2004). A team's effectiveness in managing emerging dependencies is linked to the sharedness of a mental model and the coordination mechanisms that are used, which ultimately influences the performance of the team (Espinosa et al., 2004; Lim & Klein, 2006).

It is important to note that models are abstractions of the real world, and as such, this study attempts to simplify the systems under study to easily address the specific team-related problems, as opposed to mirroring every detail in the systems (Sterman, 2000). Only the essential features of these complex systems remain, yet these features will be adequate to explain how interactions from within the system generated the outcomes in each vignette (Sterman, 2000). The causal loop diagrams used in this study will consist of nodes and edges. Nodes will represent variables, while edges will characterize relationships (Khan et al., 2022). Arrows will be used to map the variables in the causal loop diagram; reinforcing arrows will be marked with an "S," while balancing arrows will be marked with an "O." Using this design, the authors will model the "as-is" systems present in the vignette.

After the "as-is" models have been developed, the authors will seek out the strategically important nodes, identify the most impactful, strategic loops, and propose modifications to the relationships between the nodes and loops (Dotterway, 1992). These new models will be characterized as the "to-be" system, and will reflect envisioned cases where a human or machine teammate could be introduced into the system to alter the

dynamics of the system to achieve a desired outcome. Akin to the modeling efforts of Dotterway (1992) and Snook (2011), the “to-be” models will attempt to change specific nodes, as well as the relationships between nodes preceding it, and those directly following it.

D. CHAPTER SUMMARY

In this study, the authors demonstrate how conceptual systems dynamics models may be used to characterize the performance of human machine teams. This qualitative approach is appropriate given the uncertainty and relative immaturity of the field of human machine teaming, which has elements of computer science, cognitive psychology, and military science integrated throughout (Sterman, 2000; Khan et al., 2022). Through the use of causal loop diagrams, a series of models will demonstrate the “as-is” team behaviors that led to the breakdowns present in the historical analyses of the shutdown of two UH-60 Black Hawk helicopters during Operation Provide Comfort as well as the shutdown of Iran Air flight 655. These “as-is” models will be contrasted with “to-be” models where a machine teammate assists in the regulation of team behaviors. The models will be primarily based upon common system archetypes that emerge when systems are viewed as a whole. This holistic view will allow the authors to evaluate the consequences, both intended and unintended, that emerge over time. Ultimately, the authors intend to inform stakeholders, such as system designers, about the desired team behaviors present in high-performing human machine teams.

VIII. ANALYSIS

A. DESIRED TYPES OF HUMAN-MACHINE INTERACTION

Decision makers in combat environments are intimately familiar with the time pressures, uncertainties, requirements, and limitations that they must operate within (Sycara & Lewis, 2004). It is believed that a machine partner with adequate access to data could assist decision makers by providing them decision support to discern important data trends more accurately, support planning efforts, course of action discrimination, as well as the prediction of consequences from these actions (Sycara & Lewis, 2004). To avoid complications and ensure that machine interactions add value, the machine partner must have access to the same information available to human team members (Sycara & Lewis, 2004). When the general support to a team is considered, machine partners may help teams more efficiently access as well as disseminate information, communicate, monitor conditions, and plan actions (Sycara & Lewis, 2004). These observations could provide vital inputs to the decision maker that they could then use to orient a team, and will ultimately form the basis upon which actions are performed. To be useful in rapid decision-making, these inputs must be easily understood by a human partner, and must support the use of the human's mental models.

B. VIGNETTE 1: OPERATION PROVIDE COMFORT

1. Background

On April 14, 1991, two U.S. Army Black Hawk (UH-60) helicopters were mistakenly shot down by two U.S. Air Force F-15C aircraft inside a no-fly zone over Northern Iraq. As a result, all 26 souls aboard the two helicopters perished (Snook, 2011). To understand how Operation Provide Comfort came about it is important to understand the geopolitical climate in the Middle East at the time. In early August of 1990, Saddam Hussein ordered roughly 1,000 tanks to charge south in the invasion of Kuwait. At the same time, he directed his special forces to strike Kuwait City via the air while seaborne commandos closed exit routes to the south (Snook, 2011). When all was said and done, it

only took Hussein two days to conquer Kuwait and place his Republican Guard along the Kuwaiti border facing Saudi Arabia.

In response to Saudi Arabia's plea for help, the U.S. acted quickly and fiercely in what is commonly known as Operation Desert Shield (Scales, 1993). The U.S. deployed a vast array of resources on land, in the air, and sea. They created a line in which Saddam and his forces were not to cross or threaten Saudi Arabia with any further aggression. The U.S. then went to the United Nations with a request to demand that Iraq withdraw their forces from Kuwait. A deadline to withdraw by January 15, 1991, was given (Scales, 1993). Again, Hussein did not budge. The deadline passed and so on January 17, 1991, Operation Desert Shield began with a missile strike courtesy of a U.S. Army Apache helicopter (Snook, 2011). The precision airstrikes that followed over the next 40 days were unprecedented. Following the aerial bombardment, coalition ground forces hammered Iraqi ground troops for an additional two weeks. After this punishment, Hussein gave up and formally accepted United Nations conditions and withdrew his forces (Scales, 1993). At this point, combat had ended, but the humanitarian missions were only beginning.

Although the war had ended, hostilities inside of Iraq continued. Shiites in southern Iraq fought with Hussein's Baathist followers (Snook, 2011). The Kurds in the northern part of the country fought to gain control of multiple cities. Both had brief moments of success until Hussein began to slaughter thousands with the help of his Republican Guard (Snook, 2011). The Shiites in the south were pounded by artillery strikes and aerial gunships. Meanwhile, the northern Kurds were pushed up against the Turkish border (Snook, 2011). It is estimated that approximately half a million refugees were displaced while 2,000 were killed each day as a result of Hussein's assault and the harsh mountainous conditions (Scales, 1993).

These gruesome assaults once again took place as the world looked on in horror, most of the time via the six o'clock national news. The United Nations was again called upon to intervene. Resolution 688 was created to establish a safe zone along the Iraqi border with Turkey (Scales, 1993). The relief efforts would spur the inception of Operation Provide Comfort, which aimed to provide security of the more than sixty relief agencies (Scales, 1993). The mission statement for Operation Provide Comfort was to "deter Iraqi

behavior that may upset peace and order to northern Iraq. On order, respond with sufficient force to protect nation's interest should deterrence fail" (Rudd, 2004, p. 7).

Prior to understanding the mishaps of Operation Provide Comfort, it is important to understand the complexities of Combined Task Force (CTF) Provide Comfort. Miscommunication played an integral role in the shootdown of the two black hawks. Communication among one force is difficult enough. Communication among different branches and different countries proves to be exponentially more challenging. CTF Provide Comfort was created to organize and perform military support for the ongoing relief efforts (Snook, 2011). It was stood up by Lieutenant General Shalikashvili on April 18, 1991 (Snook, 2011). As its title indicates, it was a force made up of different nations including the United States, France, England, and Turkey. Among the U.S. forces, each branch of service was represented in some form or fashion (Rudd, 2004).

The CTF was split into two Joint Task Forces (JTFs) (Snook, 2011). The first task force, JTF Alpha, was made up primarily of U.S. Army special operators whose task it was to aid the relief agencies in distributing humanitarian assistance to the refugees (Snook, 2011). JTF Bravo's mission was more combat focused. It consisted of the 24th Marine Expeditionary Unit in addition to British and Dutch marines (Snook, 2011). They were tasked with creating the physical safe zone set forth by the United Nations' resolution (Rudd, 2004). Achieving this was no simple task but by mid-May they had successfully carved out an area measured at more than 11,000 square feet (Snook, 2011).

A key goal of CTF Provide Comfort was to occupy the airspace above northern Iraq. In order to do this, the task force established a "no fly zone" for the entire airspace north of the 36th parallel (Rudd, 2004). The task force patrolled this zone on a daily basis to ensure no Iraqi planes could attack or threaten the ongoing relief efforts. In September of 1991, the urgent goals of the humanitarian and relief efforts had been successfully accomplished (Snook, 2011). As a result, the majority of the CTF was redeployed back to their respective home bases. Following these redeployments, the ground component of the CTF was deactivated. A military coordination center (MCC) was established in its place (Snook, 2011). Now, the CTF consisted of three components: the MCC, the combined forces air component (CFAC), and the joint special operations component (JSOC). The

MCC was responsible for any ground missions. The CFAC exercised tactical control of all Operation Provide Comfort aircraft. The JSOC was responsible for conducting search and rescue missions in the event a coalition aircraft went down over Iraqi airspace (Snook, 2011). Geographically speaking, the CFAC and JSOC both operated from Incirlik Air Base in Turkey (Snook, 2011). The MCC, however, was split into two separate locations. They had a headquarters inside of Iraq in the village of Zakhu. The MCC headquarters was supported by an administrative section who were located on the Pirinlik Air Base, Turkey. This airbase is where the two ill-fated Black Hawks were located (Snook, 2011).

The two UH-60 Black Hawks, whose respective callsigns were Eagle 01 and Eagle 02, were one of the three teams involved in the friendly fire shutdown of Operation Provide Comfort. The other two teams consisted of an E-3B Airborne Warning and Control System (AWACS) Boeing 707 and two U.S. Air Force F-15C Eagle jets whose call signs were Tiger 01 and Tiger 02 (Snook, 2011). Each team will be described in further detail within the following paragraphs.

The AWACS aircraft consisted of a platform of sensor equipment that is aimed to conduct surveillance and communicate with other aircraft in its area of responsibility. It is commonly referred to as an aircraft control tower in the sky (Bowron, 1998). It is widely recognized for the awkward looking radar dome that sits on top of the plane. On this particular mission, the AWACS mission was to deliver early threat warnings to friendly aircraft as well as provide air control for all aircraft apart of Operation Provide Comfort (Bowron, 1998). During this operation, the aircraft contained 19 crewmembers who were split into two teams, the flight crew and the mission crew (Snook, 2011). The flight crew was a four-member team consisting of the pilot, copilot, navigator, and a flight engineer. These four teammates were located at the front of the aircraft and their call sign was *SAVVY*. The mission of *SAVVY* was to provide a safe flight (Snook, 2011).

The mission crew's callsign was *COUGAR* and they operated towards the rear of the aircraft. The mission crew was "responsible for the command, control, surveillance, communications, electronic, and management functions to include: the control and monitoring of assigned aircraft, sensor management, internal and external communications management, and onboard systems maintenance" (U.S. Air Force, 1993, p. 4). The mission

crew is further broken down into three teams consisting of the surveillance section, the weapons section, and the equipment technicians (Snook, 2011). These three teams fall under the direction of the mission crew commander. The equipment technicians are responsible for in-flight maintenance of the communications, avionics, and sensor systems. Their job is frontloaded as they are responsible for getting the systems on board warmed up and ready to go. Once the mission crew commander announces, “on station,” the surveillance and weapons sections take over (Snook, 2011). The surveillance section is led by the air surveillance officer (ASO) and their fundamental job is to deal with all non-Operation Provide Comfort aircraft (Snook, 2011). This includes unidentified aircraft that appear on their scopes. During this particular mission, their focus was any contacts below the 36th parallel.

The weapons section is led by the Senior Director (SD) and as would be expected, is responsible for air warfare (Snook, 2011). To help visualize the onboard physical positioning, the SD sat right between the MCC and the Airborne Command Element (ACE). *DUKE* was the callsign for the ACE. *DUKE*'s job was to serve as a liaison to the Mission Director on the ground. For future reference, the MD's callsign was *MAD DOG* (Snook, 2011). Serving underneath the SD were three weapons directors who each had specialized tasks. The first weapons director was titled the enroute controller. His job was to handle the flow of Operation Provide Comfort aircraft entering and leaving the no-fly zone as well as conduct radio checks with friendly aircraft outside of their area of operations (Snook, 2011). The tactical area of responsibility (TAOR) controller dealt with threat indicators and was responsible for the tactical control of friendly aircraft in the TAOR. Lastly, the third weapons director was the tanker controller. His job was to coordinate air refueling operations (Snook, 2011). In its totality, the weapons section was responsible for the tracking, controlling, and locating all friendly aircraft supporting Operation Provide Comfort (Snook, 2011). Notably, this was the first time this crew had flown together. This would play a pivotal role in the outcome as communication in this scenario would prove enormous.

The second team in this tragic event consisted of two UH-60 Black Hawks, also known as *EAGLE 01 and EAGLE 02*. As mentioned earlier, both aircraft belonged to the

MCC who operated out of Pirinlik Air Base in Turkey (Snook, 2011). Their mission consisted of general aviation support to the CTF (U.S. Air Force, 1993). The “UH” in their official nomenclature stands for Utility Helicopter (U.S. Air Force, 1993). As the name implies they are used for a wide variety of missions. On this fateful day, they were serving in a personnel transport capacity and were carrying 26 peacekeepers (Snook, 2011).

It is important to detail their communication capabilities as it was eventually determined that Identify Friend or Foe (IFF) failure was partly responsible for the shootdown (Perry, 1994). The Black Hawks standard avionics package consisted of IFF equipment, radar, and FM, UHF, and VHF radios (Snook, 2011). Prior to each flight, cryptographic keys would be used to install classified codes that enables a secure communication link with other friendly aircraft and controllers. One piece of equipment that the Black Hawks did not possess was called *HAVE QUICK*. At the time, *HAVE QUICK* was the best anti-jamming radio capability (Snook, 2011). Notably, both of the F-15s and the AWACS flight possessed *HAVE QUICK*. This absence of a communication method contributed to the lack of shared awareness between the F-15 and Black Hawk pilots (Snook, 2011). That being stated, the Black Hawks did possess IFF equipment that allowed it to be interrogated by other friendly aircraft trying to determine whether it was friend or foe. This interrogation is called a “parrot check” as once the aircraft is interrogated its “squawks” back with a code that lets the interrogating aircraft know that it is friendly (Snook, 2011). In order for the IFF system to work properly, the interrogator and the transponder must be loaded with the proper codes before takeoff.

In addition to its radios and IFF, the Black Hawks were outfitted with some protective countermeasures. They both possessed a system called the APR-39, which is a system composed of five antennas that is used to pick up radar signals (U.S. Air Force, 1993). The APR-39 alerted pilots that they were being “painted” (a term that means being searched or locked onto by another radar). If the pilots discover they are being painted, then they could use what is called an ALQ-144 system, which jams heat-seeking missiles through the use of infrared. In addition, they both possessed the M130 chaff dispenser, which is designed to disrupt radar guided missiles (U.S. Air Force, 1993).

The third and final team consisted of two F-15C fighters bearing the call signs of *TIGER 01 and TIGER 02*. These jets belonged to the 53d Fighter Squadron out of Spangdahlem, Germany (Snook, 2011). At the time, air crews from the 53d would rotate in and out of Turkey every six weeks housing their aircraft at Incirlik Air Base. When these crews worked out of Incirlik they fell under the operational control of the Combined Forces Air Component Commander (CFACC) (Snook, 2011). The two F-15C aircraft being flown that day were designed specifically for air-to-air combat and were carrying an assortment of missiles. In addition to their ordnance, the fighter jets traveled with an avionics platform that was similar to but not identical to the Black Hawks. The F-15s possessed *HAVE QUICK* radios, which to reiterate, the Black Hawks did not. Similar to the Black Hawks, the F-15s had IFF capabilities but in addition to the transponders the Black Hawks carried the jets possessed interrogators (Snook, 2011). This interrogation capability was the reason they would always fly the first mission into the TAOR on a daily basis, per task force policy (Rudd, 2004). The interrogators allowed them to sweep the area by searching for unidentified or unauthorized aircraft.

Air Force standard operating procedures (SOP) directed that F-15s fly in pairs (U.S. Air Force, 1993). This foundational sortie always consists of a lead and a wingman. Regardless of rank, the lead is always in charge. On this particular Operation Provide Comfort mission, the lead was a captain and the wingman was a lieutenant colonel. In addition to outranking his lead, the lieutenant colonel was the squadron commander (Snook, 2011). Those familiar with military rank structure and command relationships might find this interesting at a minimum, yet this is common procedure in the air community. So, on this mission the captain had command authority over the lieutenant colonel. Both were highly qualified and had the mission ended in tragedy, this debate would be less than an afterthought. Their mission on this day was similar to most days, sanitize the area to ensure there were no hostiles in the no-fly zone (U.S. Air Force, 1993).

2. Overview

As mentioned in the background, for the AWACS crew this was their first mission together. The majority of the crew began arising around 0430 on the morning of the flight

(Snook, 2011). They loaded their crew bus at 0600 in their billeting area for the ride to the airfield. At the airfield and prior to the flight, they received a few standard briefs as well as an intelligence update. After checks and preparation, the AWACS flight departed at 0736 (Snook, 2011). They ascended 32,000 feet into an orbit they would hold until the mission crew (*COUGAR*) had their systems up and running. Once the mission crew signaled they were operational, the flight crew (*SAVVY*) took them the rest of the way to the restricted operating zone one (Snook, 2011). At 0845, the Airborne Command Element (*DUKE*) notified the Mission Director (*MAD DOG*) on the ground that they were now “on station” (Snook, 2011). This notification cleared the way for *MAD DOG* to send friendly aircraft into the TAOR.

The Black Hawk crew’s morning started around 0515 with a short ride to Diyarbakir Air Base (Snook, 2011). A little after 0700, *EAGLE 01*, who was also the mission commander, convened a short briefing on the upcoming flight. It was a standard briefing consisting of an intelligence update, proper coding techniques, altitudes, and reporting procedures (Snook, 2011). Following the meeting, flight plans were dropped off at Operations and the crews headed out to their helicopters after a brief stop to draw their weapons (Snook, 2011). Upon arriving at the helicopters, the crews went through standard preflight inspections. At 0815, the pilots radioed air traffic control for permission to liftoff (Snook, 2011). It was at this time that one Turkish liaison boarded each flight. Per coalition agreement, all American military flights departing Turkey would have a Turkish officer on board (Rudd, 2004). At 0822, the two helicopters rose into the morning sun climbing to an altitude of 3,500 feet. At 0921, the Black Hawks made their first contact of the day with AWACS as they prepared to enter the TAOR (Snook, 2011). Prior to the radio contact, the AWACS tracking system showed the Black Hawks designator as “TY06.” Once the helicopters checked in, this designation was changed to a “EE01” by the SD indicating a friendly helicopter. However, at 0924, the tracking system lost the Black Hawks IFF signal and their symbology was stopped (Snook, 2011). This was the first indicator of possible communication problems between *EAGLE 01* and *02* and the AWACS flight. This was inconsequential at this point, however, as the Black Hawks had yet crossed into Iraqi airspace to land and pick up their peacekeeper passengers. At 0941, *EAGLE 01* and *EAGLE*

02 touched down in the Iraqi village of Zakhu where the MCC was located (Snook, 2011). Here they would pick up 18 passengers for the follow-on flight. At 0954, the Black Hawks departed Zakhu and radioed AWACS reporting they were now enroute to their scheduled meetings between United Nations and Kurdish delegates (Snook, 2011). About twenty minutes after takeoff, the Black Hawks flight path took them into a deep valley in between particularly rugged mountains. At this point, *EAGLE 01* and *EAGLE 02* once again dropped off AWACS radar (Snook, 2011).

The morning for the pilots of the F-15s began around 0630. Shortly after awaking and a quick breakfast, both pilots were taken by the duty van to the squadron at approximately 0730 (U.S. Air Force, 1993). After arriving at the squadron, the pilots sat down to receive briefings regarding the day's mission. Highlighting the intelligence brief was the report of the losing sight of an Iraqi surface-to-air missile site. In addition to this report, they were also informed about the shootdown of the Rwandan president the previous night (U.S. Air Force, 1993). Although both reports were unrelated to the shootdown that would occur later in the day, it is hard to argue that the events did not reside in the shared mental model of the two F-15 pilots. Increased tension in the minds of the pilots should not be discounted as they both knew there were threats potentially lying-in wait. Following the intelligence brief, the team reviewed their communication plan, looked at the list of scheduled aircraft flying in the TAOR that day, and their engagement plan should they encounter bogies (U.S. Air Force, 1993).

At 0900, *TIGER 01* and *TIGER 02* fired up their aircrafts. They promptly informed the Mission Director (*MAD DOG*) that they were ready to depart and asked for any changes to the Air Tasking Order (ATO). The Mission Director reported that there were no changes to the ATO and at 0920, the F-15s rolled to their arming location where the safety pins would be pulled from their missiles (U.S. Air Force, 1993). It was at this point that a check of their IFF systems was conducted. At this time, there were no issues and at 0935 both F-15s departed Incirlik Air Base in Turkey enroute to sanitize the no-fly zone in over northern Iraq (Snook, 2011). Shortly after takeoff, the F-15s would once again check their IFF systems per Air Force policy (U.S. Air Force, 1993). This was performed by having *TIGER 02* first interrogate *TIGER 01* on all four modes. Once completed, *TIGER 02* pulled ahead

of *TIGER 01* allowing him to conduct an interrogation using all four IFF modes. Both system checks were positively executed without inaccuracies (Snook, 2011). Following the IFF check, both pilots switched over to the enroute controller frequency to execute a radio check with the mission crew (*COUGAR*) aboard AWACS. After a positive radio check on the enroute controller frequency, both *TIGER* flights switched over to their HAVE QUICK radios that would be used upon entering the TAOR. The HAVE QUICK radio check was positively conducted ensuring both *TIGER 01* and *TIGER 02* had contact with AWACS via their anti-jamming communication devices (Snook, 2011). Prior to entering the no-fly zone, all communication checks came back positive. The last step in preparation entering the TAOR was to switch from their IFF code from 43 to 52. On that day, all friendly fixed-wing flying over Turkey would squawk on 43. Prior to entering the Iraq they were required to flip over to 52. At 1020, both F-15s passed through gate one and officially into Iraqi airspace. *TIGER 01* radioed the mission crew (*COUGAR*) aboard AWACS informing them they were now “on-station” (Snook, 2011).

Up until this point, all was as scheduled on this Operation Provide Comfort mission. However, at 1020 both *EAGLE 01* and *EAGLE 02* dropped off AWACS radar as they flew about 200 feet above ground in a low-lying valley between two mountain ranges (Snook, 2011). It was at this moment that the first critical mistake was made. The enroute controller removed the Black Hawks symbol from their scopes. This was the only visual reminder that *EAGLE 01* and *EAGLE 02* were inside the no-fly zone (U.S. Air Force, 1993).

Shortly after entering Iraqi airspace, *EAGLE 01*'s radar picked up something flying about 40 nautical miles southeast of his position (Snook, 2011). This showed up as a rectangle on his scope. To try and identify direction, altitude, and air speed, *EAGLE 01* locked on to the target at which point the symbol changed to a star (U.S. Air Force, 1993). *EAGLE 01* then used his interrogator to attempt and identify if the target was friendly. If the target was squawking code 52, a diamond would have replaced the rectangle. That did not happen. At this point, *EAGLE 01* switch the interrogator from “CC” to “auto” (U.S. Air Force, 1993). This change meant the interrogator would continuously monitor Mode IV. At the time, this was the secondary mode all friendly aircraft were directed to squawk. If this were the case, a circle would have replaced the rectangle on *EAGLE 01*'s scope (U.S.

Air Force, 1993). According to one of the investigations, the pilot of *EAGLE 01* explained that the star did briefly become a circle, however, it quickly turned back to a star (Andrus, 1994). This means that for a moment the scope identified the target as a friendly aircraft. The pilot would go on to explain that the interrogator had previously malfunctioned in a similar fashion (Andrus, 1994). This trust in identification capabilities should have caused red flags. At this point, *TIGER 01* had *TIGER 02* check his radar to verify the target. He did confirm with a “hits there” response (Andrus, 1994). This meant his radar was also picking up an aircraft.

At 1022, *TIGER 01* radioed the mission crew (*COUGAR*) aboard the AWACS to report the contact. Referencing “Bulls Eye,” a coalition geographical reference location in Iraq, *TIGER 01* told *COUGAR* he had a contact thirty degrees off of “Bulls Eye” at thirty degrees for fifty miles. *COUGAR* came back with a “Clean there,” response (Snook, 2011). This meant *COUGAR* was not picking up any contacts on his radar. After reviewing the data tapes in one of the investigations, it was determined that in fact the AWACS was not receiving any IFF returns at that time (Andrus, 1994). However almost immediately following this radio transmission, the AWACS started receiving sporadic IFF signals. Then at 1023, the “H” symbology once again appeared on the SD’s radar screen (Snook, 2011). The IFF signals would remain intermittent until 1028 when they became solid and visible without disruption (Snook, 2011).

Now uncertain of the target he was tracking, *TIGER 01* began to search for other contacts. While doing so he checked his map and noticed a road in the area of his contact. *TIGER 01* then asked *TIGER 02* if he thought the target could possibly be a vehicle traveling on the road. *TIGER 02* indicated he did not believe it was an automobile as his scope was showing the target traveling at 300 feet above the ground (Snook, 2011). Following the continued investigation of the target, *COUGAR* radioed back to the *TIGER* that he did now have “hits there” (Snook, 2011). The verbiage of this dialog would go on to be scrutinized as a “hit” typically refers to any blip on the radar while the term “paint” means an IFF return (U.S. Air Force, 1993). In a continued attempt to verify whether the target was friendly or not, the *TIGER* team once again attempted to interrogate target in Mode I and IV (Snook, 2011). Once again, he received no indication of friendly aircrafts.

After failing to make a positive identification via electronic means, *TIGER 01* proceeded to establish visual identification (VID) to confirm whether or not the target was in fact hostile. To do so, he would descend from his current altitude of 27,000 feet. At 1026, he was now within seven nautical miles. His onboard computer had automatically positioned his lock onto his target designator (TD) box (Snook, 2011). Even so, at this point he could not positively identify the target (Andrus, 1994). Then, at 1027, he visually established that a helicopter was in his TD box at a range of five miles out (Snook, 2011). *TIGER 01* contacted *COUGAR* to notify them of this identification and inform them to standby for VID of whether it is friendly or not. At this time, *COUGAR* placed a symbol representing an unknown target in the area they had previously been receiving IFF returns (Snook, 2011).

To establish a VID, *TIGER 01* conducted a flyby of his target from above. As he traveled over his target at a speed of 450 knots he quickly pulled up to look over his shoulder. At first glance, he thought he saw a shadow behind the helicopter but after a closer look realized that it was a second helicopter (Snook, 2011). At 1028, *TIGER 01* radioed *COUGAR* to report he had identified two “Hinds,” which is a NATO designation of a Soviet made helicopter. *COUGAR* responded back that he understood the transmission (Andrus, 1994). *TIGER 01* then directed *TIGER 02* to confirm his identification. On his approach of the targets, *TIGER 02* initially thought he saw a shadow behind the first helicopter as well before realizing a second aircraft. *TIGER 02* then verbally confirmed that he too identified two helicopters (Snook, 2011).

As the Black Hawks continued their flight to the southeast, *TIGER 01* and *TIGER 02* would fly back around to approximately 10 nautical miles before turning around again and back towards the Black Hawks. *TIGER 01* then informed *COUGAR* that both members of the *TIGER* team had identified two Hinds and were now “engaged” (Snook, 2011). This meant they were now preparing to use ordnance. *TIGER 01* then informed *COUGAR* that he was now “hot” meaning he had armed his ordnance and was preparing to destroy his target (Snook, 2011). He then told his wingman that he (*TIGER 01*) would engage the trailing helicopter while *TIGER 02* would shoot the leader (Andrus, 1994). *TIGER 01* then conducted one last IFF check on the target and received a negative response. Following the

negative response, *TIGER 01* fired one advanced medium Range air-to-air missile (AMRAAM), destroying its target. *TIGER 02* then fired one advanced intercept missile (AIM) 9 Sidewinder at his target. Once again, the missile found its mark and the fireball tumbled to earth. At 1030, the engagement was over and all 26 passengers on the two Black Hawks were dead (Snook, 2011).

3. Causal Loop Diagrams

a. A Breakdown in Communication by the AWACS Team

(1) Modeling the “As-Is” System

Analyzing the communication breakdown of the AWACS team illuminates a crew that was not ready for its mission. As a crew, this was their first flight together in-country (Snook, 2011). In addition to a lack of cohesion, the leaders of the team failed them at the most critical points on this particular mission. Most importantly and perhaps most destructive was the failure to communicate. On an individual level, each team member understood his or her role. However, at the team of teams level, their lack of communication would place them in a category known as a pseudo-team (Snook, 2011). A pseudo-team as described by Katzenbach and Smith (1993) is when the sum of the whole team is less potent than the sum of the individuals. In this case, it can be summed by using the cliché phrase of *passing the buck*. In essence, each teammate assumed someone else would communicate with the *TIGER* team and let them know there were friendly aircraft in the TAOR. Moreover, they were undertrained and not ready for the challenges they would face that day (Snook, 2011). In fact, the complications that arose were quite simple yet their inexperience and ineffective tracking of *EAGLE 01* and *EAGLE 02* caused them to fail to intervene when communicating with the *TIGER* team.

The failure to communicate on behalf of the AWACS team was heavily influenced by a lack of training together. Each member was qualified for their position but had not previously worked together. Using a *limits to success* archetype depicted in Figure 6, a lack of training inserted as a constraint severely reduced unit cohesion (Kim, 2000). “In a *limits to success* scenario, continued efforts initially lead to improved performance. Over time, however, the system encounters a limit which causes the performance to slow down or

even decline” (Kim, 2000, p. 21). A strong beginning as a team is foundational to the team’s overall strength, especially in air crews (Ginnett, 2010). This particular AWACS team was stood up quickly and in less-than-ideal circumstances. They had limited training time together that yielded uncertainty, a significant lack of trust, and a poor communication structure (Snook, 2011).

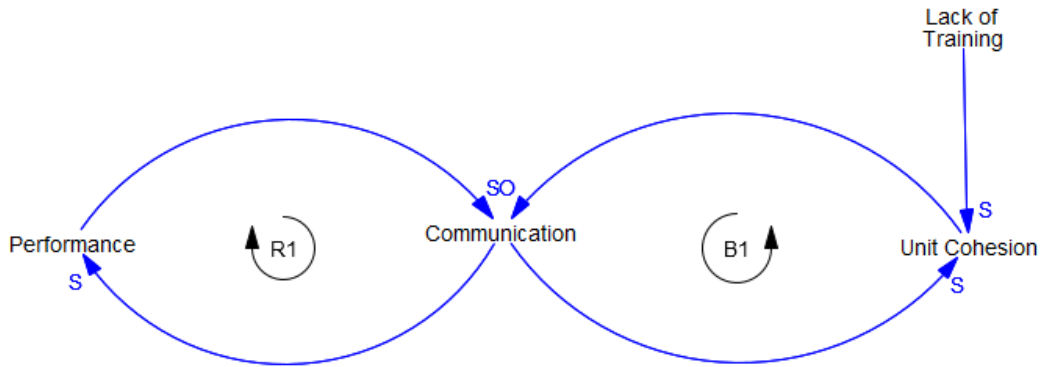


Figure 6. Failure to Communicate

Using Figure 6 to illustrate, over time communication strengthens performance, which in turn continues to improve communication as trust in the team improves. Trust has been shown to be positively correlated with team task performance, inter-team relationship commitment, and team satisfaction (Costa, 2003). Conversely, all three have been shown to reduce stress in teams (Costa, 2003). In essence, trust is propagated over time in teams through proven successful work whether in a training environment or on the battlefield. As trust builds, communication is strengthened and tacit knowledge is spread throughout the team (Nissen, 2014). The reinforcing loop shown on the left side of Figure 6 illustrates distinct but related dimensions that continue to improve over time.

This “as-is” model also goes to show that this improvement only continues if it is not impeded upon by a constraint. In the balancing loop on the right-hand side exemplifies that communication increases unit cohesion. As mentioned previously, unit cohesion is a time driven measurement that solidifies as more success is perceived by the team. When lack of training is inserted as a constraint, unit cohesion struggles to be a solid foundation

for a team. As an example, a basketball team that has played together for one year versus a team that has had no prior experience together is more likely to win given that both teams are of equal skill level. This is because the teammates know how each other plays. They know where their strengths and weaknesses are, they understand how each team member moves, and they can even implicitly communicate. This same form of unit cohesion is applicable to all teams, at times even more so in the military. Since the AWACS team had such little training as a unit, the cohesion bond had not formed prior to this Operation Provide Comfort mission. As a result, communication that may have been clearly understood by a cohesive unit was clearly not received by this team.

A brief review of the timeline further demonstrates this breakdown. After picking up their passengers in Zakhu at 0954, *EAGLE FLIGHT* informed the AWACS team they had departed for Lima (code word for their destination). The enroute controller acknowledges but it would be determined he did not actually know where Lima was located. He did possess a list of codes, which displayed this information, but failed to use it or ask a teammate for assistance (Snook, 2011). At 1012, *EAGLE FLIGHT* dropped off the SD's scope as they flew low through a valley. The ASO then input an attention arrow (indicates area of interest) in place of the lost helicopter symbology as a reminder to the SD that there were Black Hawks in that area. To be clear, there was no verbal communication to the SD by the ASO and when the SD failed to annotate this visual cue, and the attention arrow dropped off of his scope after sixty seconds (Snook, 2011). At 1020, the *TIGER* team entered the TAOR and asked AWACS for any updates. Although there was a computer-generated helicopter symbol based on air speed, heading and last known location, the AWACS crew failed to mention the Black Hawks to the *TIGER* team (Andrus, 1994). At 1021, the enroute controller completely removed the "EE01" symbol representing *EAGLE FLIGHT* (Snook, 2011). One minute later, *TIGER 01* reported a hit to the controller who reported no contacts on his scope. In what would seem to be an obvious place for further investigation and communication, the AWACS team simply does nothing (Andrus, 1994). At 1023, intermittent IFF signals began to reappear as the Black Hawks left the valley (Snook, 2011). At 1025, *TIGER 01* reported contact again and even though IFF returns were visible to the AWACS, no one aboard the E-3B aircraft identified

this to the F-15s (Snook, 2011). To make things worse, the controller then placed an “unknown” symbol on his scopes even though the “H” symbology was present on the SDs scopes (Snook, 2011). One minute later at 1028, *TIGER 01* visually identified two suspected Hinds. The controller simply responded that he understood and shortly thereafter the two Black Hawks were destroyed.

(2) Modeling the “To-Be” System

According to Kim (2000), a *limits to success* archetype is best used in advance of issues. In the case of a failure to communicate such as was observed on behalf of the AWACS team, critical thinking is required to solve the problem. In reinforcing loop one in the “as-is” model seen in Figure 6, communication strengthens performance, which in turn performance continues to strengthen communication. Unimpeded, this cycle continues but as the archetype suggests, it is only possible to succeed for a limited period of time (Kim, 2000). Balancing loop one begins to weaken communication as a lack of training is injected and weakens unit cohesion. The problem solver must identify why this lack of training exists. Is it because the unit has been so successful that training was deemed not critical, or was the lack of training a result of a bigger fundamental issue? Perhaps the larger organization has become lackadaisical as previous missions seemed to have went smoothly, even with limited training beforehand.

One potential solution to solving the problem symptom is adding a machine listener to the team. It must be noted that maximum training time together as a unit would likely be the most favorable fundamental solution; however, in the event that training prior to a mission cannot be accomplished, adding a machine listener may be able to prevent a tragedy such as the shutdown of *EAGLE 01* and *EAGLE 02*. This can be demonstrated by utilizing a *shifting the burden* causal loop diagram as seen in Figure 7 (Kim, 2000). In this type of archetype, Kim (2000) offers, “explore the problem from a different perspective in order to come to a more comprehensive understanding of what the fundamental solution may be” (p. 21).

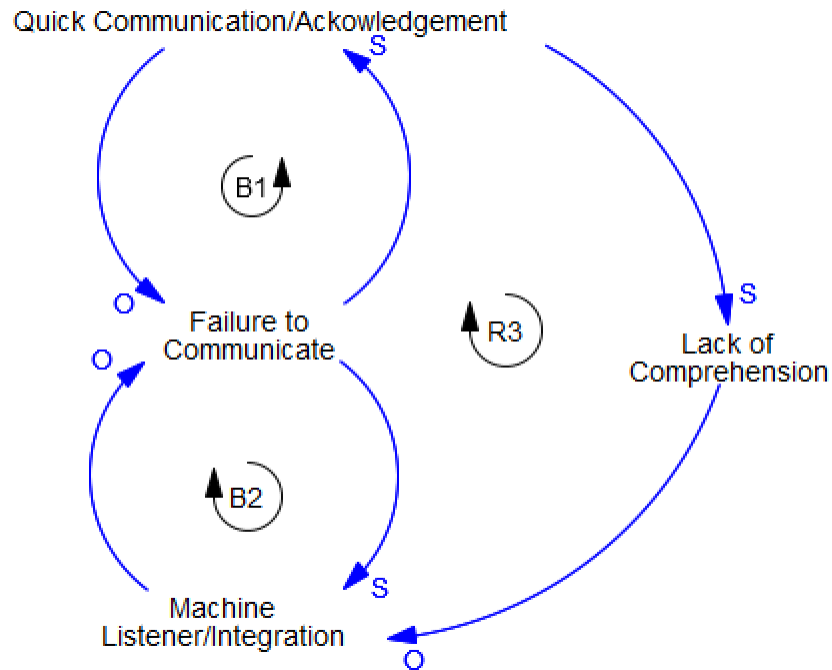


Figure 7. AI to Integrate Team Communications

In the causal loop diagram depicted in Figure 7, failure to communicate is the problem symptom. In the case of this particular Operation Provide Comfort mission it is mainly due to a lack of training time and togetherness seen in the AWACS team. This failure to communicate is symptomatically solved through quick communication that is ineffectively acknowledged. This ineffective communication results in a lack of comprehension throughout the team. At this point, the team is unaware of what each team member knows and does not know. Not only does this cause trust issues, but their shared mental models do not coincide. Inserting a machine listener into the loop to breakdown the lack of comprehension will help reduce the effects of a lack of communication. To be clear, it does not solve the problem of communication between the teammates. Rather, if cojoined with the sensor information gathered by the AWACS platforms (radars, radios, etc.) it can listen to ensure that this information is being passed within the team. Should it not be communicated, the machine listener can then insert itself either audibly or through computer screen notifications. For example, the Enroute Controller placed an attention arrow to indicate to the SD that the Black Hawks had dropped off of the scope (Snook,

2011). If the alert is not recognized within sixty seconds, it automatically drops from the SD's scope (Snook, 2011). It is possible that if a machine listener teammate was integrated into the system it could have reminded the SD of an unnoticed alert. Additionally, if a machine listener was onboard and was tracking the crew's radio contacts it is possible that it could have reminded the crew that there were in fact Black Hawks in the area of question when the F-15's first reported contact. Further yet, if it becomes possible for a machine listener to monitor all of the communications links in use, it could have notified the F-15s directly that friendly helicopters had entered the no-fly zone prior to the sanitization of the TAOR by *TIGER* team.

Machine listening interprets signals using machine learning to abstract meaningful knowledge from sound (Cella, 2017). In essence, it is a way of processing audio signals via a computer in a way that replicates the way humans process audio signals. As machine learning has rapidly progressed over the last twenty years, machine listening has been used to monitor heart statuses, used in the composition of music, and of course, smart home assistants such as Google Assistant, Apple's Siri, and Amazon's Alexa (Ahmed, Agarwal, Kurniawan, Anantadjaya, & Krishnan, 2022). Prior to machine listening, artificial intelligence was powered by computer visualization, language processing, and speech recognition (Cella, 2017). Having the ability to hear and interpret what is being heard is far different from being able to recognize speech and hear sounds. This is the major difference in speech recognition and machine listening (Cella, 2017). For example, most humans would know they needed an umbrella to go outside if they heard rain outside their window. Humans can hear footsteps on a marble floor and distill what type of shoe is worn by the passerby. These examples demonstrate the difference between listening and hearing (Oneto, 2022). It is worth noting that machine listening requires understanding advanced mathematical signals that are beyond the scope of this study.

b. Misidentification of the Black Hawks

(1) Modeling the "As-Is" System

On 14 April 1991, 26 souls aboard two Black Hawk helicopters tragically lost their lives as they were shot down over Iraqi airspace due to friendly fire. Initially, the pilot of

TIGER 02 (the wingman) was charged with 26 counts of negligent homicide while the pilot of *TIGER 01* (the lead) was not charged at all (Verhovek, 1995). The reason behind this decision was that *TIGER 01* identified the helicopters as hostile Hinds while *TIGER 02* did not confirm that the helicopters were Hinds and still proceeded to engage. *TIGER 02* would go on to recant his testimony at his hearing by saying he did positively identify the helicopters as Hinds. As a result, his errors were found to be “reasonable” and his charges were dismissed (Verhovek, 1995). Additionally, multiple members of the AWACS team were charged with dereliction of duty (Verhovek, 1995). However, after numerous investigations and hearings, only one individual stood trial. Captain Jim Wang, who was the SD aboard the AWACS flight was sent to court-martial but was eventually acquitted (Headquarters 8th Air Force, 1995).

Making sense of why the F-15 pilots misidentified the two friendly helicopters as enemy hinds is a difficult task. Focusing specifically on their incorrect visual identification of the Black Hawks points to a system archetype known as *escalation* (Kim, 2000). In this type of archetype, “one party (A) takes actions that are perceived by the other as a threat. The other party (B) responds in a similar manner, increasing the threat to A and resulting in more threatening actions by A” (Kim, 2000, p. 22). These actions led to the reinforcing loops as seen in Figure 8.

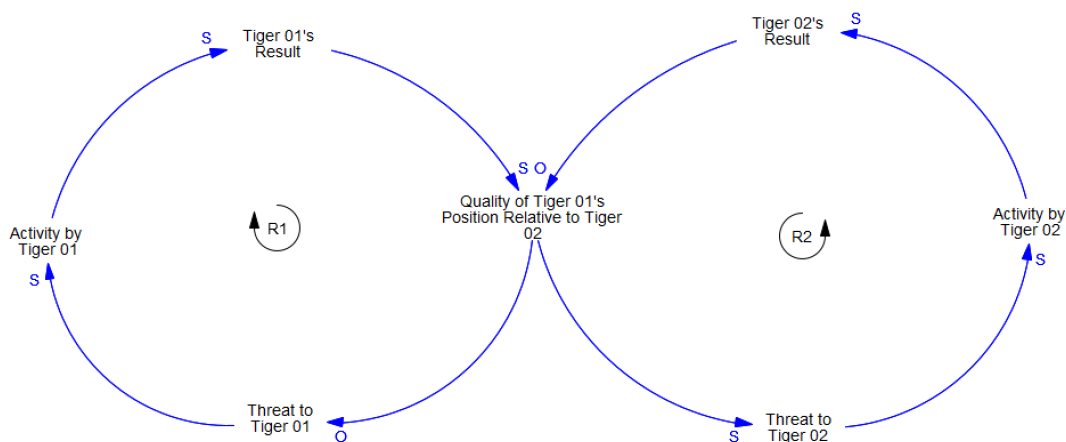


Figure 8. Misidentification of Black Hawks

As a result of multiple preflight factors, *TIGER 01* expected to see an enemy helicopter and his suspicions were further confirmed by his wingman. Both pilots had combat experience and both were well aware of the political climate of northern Iraq (Snook, 2011). Both had recently received intelligence briefs outlining the movement of one hundred thousand Iraqi troops into the area in which they were operating. Additionally, the TAOR was officially designated as a combat zone (Snook, 2011). Although the flight was in support of a peacekeeping mission, they were well aware that they were operating over enemy airspace. Most significantly, there was no mention of friendly helicopters operating in the TAOR. To be clear, neither *EAGLE 01* nor *EAGLE 02* were listed on the ATO (Snook, 2011). Additionally, *TIGER 01* and *TIGER 02* were listed as the first friendly aircraft to enter the TAOR (Snook, 2011).

The escalation archetype depicted in Figure 8 demonstrates how *TIGER 01*'s expectation of seeing Hinds is ultimately confirmed by the activity of *TIGER 02*. That being said, it must not be forgotten that multiple stimuli led to the shutdown such as the failed IFF system, the lack of integration of *EAGLE FLIGHT* into Operation Provide Comfort operations, and misleading communication from the AWACS team (Snook, 2011). As each factor is inserted into the balancing loops, the threat to *TIGER 01* grows until he is finally convinced of the enemy threat and mistakenly sees what is not there.

(2) Modeling the “To-Be” System

There were many missed opportunities to positively identify the Black Hawks as friendly, however, the last chance to get it right came as *TIGER 01* conducted a VID fly-by (Snook, 2011). This VID was no easy task for the F-15 pilot as he had to quickly drop down from over 10,000 feet (Snook, 2011). From a distance of seven miles, *TIGER 01* thought the helicopter may be road traffic (Andrus, 1994). His wingman, *TIGER 02* noticed a shadow thus eliminating the contacts as ground vehicles (Snook, 2011). It was not until five miles out that he confidently identified his target as an aircraft (Andrus, 1994). At this point in time, his electronics were telling him that his mark was not friendly. He then followed procedures and decided to make a pass to VID the aircraft (Snook, 2011). After he made a pass, he asserted that the helicopters were in fact Hinds. Keep in mind, his F-15

was traveling more than three times the airspeed of the helicopter (Andrus, 1994). This gave him a few short seconds to visually confirm that the helicopter was a Hind based off a silhouette book he was carrying (Snook, 2011). One problem with this attempted solution is that at the time pilots were trained by studying silhouettes from the bottom of the aircraft vice overhead. In retrospect, this seems trivial at best.

After *TIGER 01* visually identified the helicopters as Hinds, he directed his wingman to make a visual pass to validate his assumption. This interaction provides another example of where the social interaction of the two pilots broke down. *TIGER 02*, after making his pass, called “Tally two” (Snook, 2011, p. 89). During sworn testimony *TIGER 01* went on to say he interpreted this as a confirmation of two Hinds. *TIGER 02* conversely said he did not positively identify them as Hinds rather his call of “tally two” meant he did not see anything that was dispute what *TIGER 01* saw (Andrus, 1994). In retrospect it seemed at least a second pass could have been made due to the uncertainty. After all, the helicopters were no threat to far superior capabilities of the F-15s.

Using a causal loop diagram archetype of *fixes that fail*, employing an AI teammate to authenticate the identification of the unknown aircraft could potentially help avoid such a catastrophic error. In a *fixes that fail* state, “a problem symptom cries out for resolution. A solution is quickly implemented that alleviates the symptom, but the unintended consequences of the fix exacerbate the problem” (Kim, 2000, p. 22). As illustrated in Figure 9, the problem symptom is ambiguous identification. This is a result of multiple factors such as the intermittent IFF returns, lack of integration of the *EAGLE FLIGHT* team in the ATO, and a failure on behalf of the AWACS to communicate that friendly helicopters were in the TAOR. The “fix” that is attempted in the VID fly-by conducted by both *TIGER 01* and *TIGER 02*. The visual identification in this case is incorrect and therefore further reduces the ambiguous identification. Additionally, after a delay, the misidentification becomes an unintended consequence. The pilots think they have solved the problem but in reality, have only made it worse and the outcome is catastrophic. By inserting AI into the misidentification variable, it may intervene and intercept the misidentification of the helicopters. This intervention then becomes authentication of the problem symptom, which

is the ambiguous identification of the friendly aircraft. The end-state is a solution provided by AI that solves the problem of ambiguous identification.

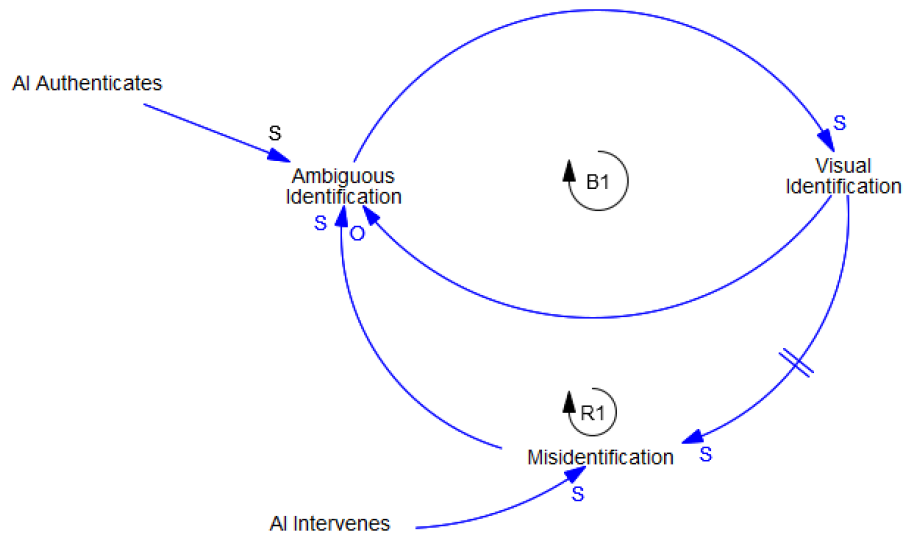


Figure 9. AI as Authenticator

AI as an authenticator is quickly growing among consumers, organizations, and even governments (Tawalbeh, Quwaider, & Lo'ai, 2020). In cybersecurity, it is used to enhance multi-factor authentication, it uses advanced algorithms in art recognition to detect forgeries, and is being rolled out in “smart cities” as a monitoring tool to help assist in a variety of decision-making processes (Rahman, Hossain, Showail, Alrajeh, & Ghoneim, 2021). Most suitable for use in the Operation Provide Comfort mission, neuromorphic vision sensors are bio-inspired instruments that replicate the functions of a human eye (Mondal, Giraldo, Bouwmans, & Chowdhury, 2021). These AI sensors “capture a stream of asynchronous events that pose multiple advantages over the former, like high dynamic range, low latency, low power consumption, and reduced motion blur” (Mondal et al., 2021, p. 878). As this capability is relatively new, there are some drawbacks as these cameras usually contain significant noise and provides low resolution (Mondal et al., 2021). Regardless, the potential for it to be a successful AI teammate warrants further study at minimum.

c. *Non-Integration of Eagle Flight*

(1) Modeling the “As-Is” System

Service rivalry in the U.S. military is long and complex. Each branch of the Armed Forces has its own goals and missions. As a result, non-integration became a hurdle that each branch had to jump over when executing joint missions. The failed 1983 rescue attempt of hostages in Iran, the Marine barracks bombing in 1983, and even the somewhat successful Grenada invasion of the same year are just a few examples of how non-integration cost friendly American service members lives (Allard, 1996). In an effort to quell anger at home, Congress enacted the Goldwater-Nichols Act in 1986. The overall goal was to shift the power from each individual service to a more joint military institution (Allard, 1996). Joint service became a prerequisite for those aspiring to promote to the rank of general officer. Although this was a step in the right direction, service rivalry still lay in the hearts of those who joined the military prior to these changes. It is almost no fault to them as they were bred in an environment of competing with their peer services for pride, utility, and money. Although non-integration is much more difficult to pinpoint than pilot error and the AWACS failure to communicate, this separatism between services did have a part to play in the shutdown of *EAGLE 01* and *EAGLE 02*.

Non-integration in the Operation Provide Comfort shutdown can be visualized in a causal loop diagram known as *Shifting the Burden* (Kim, 2000). This happens when “a problem is solved by applying a symptomatic solution, which diverts attention away from more fundamental solutions” (Kim, 2000, p. 21). In Figure 10, non-integration was initially born from service separatism. Separatism and short-term planning consistently drove non-integration in a cycle that saw only symptomatic solutions offered. Short-term planning was the quick fix when a fundamental solution was needed. Side effects, defined as coordination failures in the model, gathered as short-term solutions never solved the actual problem. These side effects then took away from the fundamental solution, which is integration. Some of these side effects were knowledge silos, manifesting as services not wanting to share information, unrealistic training, and vitriol between the services.

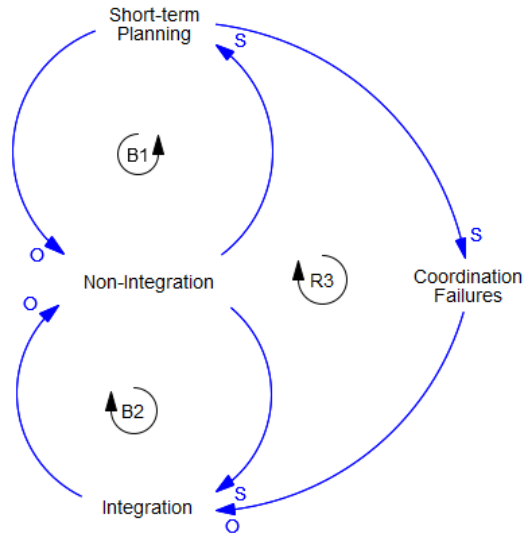


Figure 10. Non-integration of Eagle Flight

The fundamental solution of integration did not stand a chance as long as the services were competing with each other. Thus, as non-integration reached for a fundamental solution (integration), it was continually hampered by side effects and short-term solutions.

Non-integration was in fact pointed to by Secretary of Defense Perry (1994) in his findings. He identified that the F-15 pilots did not know that *EAGLE FLIGHT* would be operating the TAOR, that the Black Hawks were in the no-fly zone prior to it being sanitized, which was in violation of Operation Provide Comfort policy, and that the F-15s and Black Hawks were operating on different radio frequencies (Perry, 1994). Many studies have looked to identify how organizations integrate and differentiate. Lawrence & Lorsch (1967) studied how fragmented organizations present profound intellectual and interactive challenges when working as subunits. They would find that differentiation breeds segmentation and specialized knowledge in addition to attitudes and goals (Lawrence & Lorsch, 1967). This may still be observed today as Marines, Soldiers, Airmen, Guardians, and Sailors vie to prove their branch of services is the best. This interservice rivalry is not always a bad thing, but when non-integration interferes with mission accomplishment it can contribute to disastrous events such as in the Operation Provide Comfort peacekeeping mission.

Lawrence and Lorsch (1967) identified three dimensions in which subunits of an organization think and work differently. First is a difference in *orientation towards goals* (Lawrence & Lorsch, 1967). In the Operation Provide Comfort mission, the Army Black team served to placate those they were transporting while the Air Force team focused on thorough planning and implementation (Snook, 2011). Snook (2011) likens the Army aviators to jazz musicians while the Air Force personnel are more like classically trained musicians. These differences create few problems as they do not often play the same venues, but when they do problems are noticeable. On the day of the shootdown, the Army was aiming to please its passengers and entered the TAOR prior to the sanitation of airspace by the F-15s. Conversely, the F-15 pilots did not anticipate friendly air traffic in the TAOR prior to their arrival and were gathering their knowledge from a synchronization flow (Snook, 2011). The two teams' goals were not integrated and as a result proved fatal.

The second dimension provided is known as *orientations toward time* (Lawrence & Lorsch, 1967). When looking at the Operation Provide Comfort shootdown, the action event unfolded in less than 20 minutes. Even the members of the AWACS stated how surprised they were that the shootdown occurred after first contact was announced by the *TIGER* team (Andrus, 1994). *TIGER 01* would go on to testify that this was not out of the ordinary for him (Andrus, 1994). He made a VID of the Hinds and as such the rules of engagement were met. He thus proceeded to act as he was trained. It is worth noting that at over 500 knots, decisions must be made quickly and confidently.

The third and final dimension in which subunits of an organization tend to act differently is known as *interpersonal orientation* (Lawrence & Lorsch, 1967). This dimension refers to how teammates within subunits tend to interact with each other (Lawrence & Lorsch, 1967). This differentiation can be seen in the variance of the two Air Force subunits, the AWACS and *TIGER* teams. Fighter pilots are asserted to be flashy, flamboyant, and even a little cocky while members of the AWACS team are more technical as the majority of the crew is comprised of analysts (Snook, 2011). This differentiation may have factored into the reluctance on the part of the AWACS crew to override the fighter pilots in their identification of two Hinds (Snook, 2011).

(2) Modeling the “To-Be” System

It is easier to state in hindsight, however it seems illogical that the pair of helicopters were not listed on the ATO for the F-15s. As discussed previously, the fundamental problem is the fact that the Operation Provide Comfort task force was not fully integrated. Not to be mistaken, the Army and the Air Force did communicate and were technically a joint team. However, they were not integrated enough to provide a fundamental solution to the problem. Figure 11 provides a *shifting the burden* causal loop diagram that illustrates this type of archetype (Kim, 2000). As a reminder, in a shifting the burden archetype, “a problem is solved by applying a symptomatic solution, which diverts attention away from more fundamental solutions” (Kim, 2000, p. 22). In this case, non-integration is the fundamental problem symptom. It may be a result of service egos, poor communication, and knowledge siloes. In order to solve the problem, improvisation grows in the form of a symptomatic solution. The issue here is that improvisation in life and death situations, although a part of combat, is much more likely to fail when teams are not integrated (Peters, 2009).

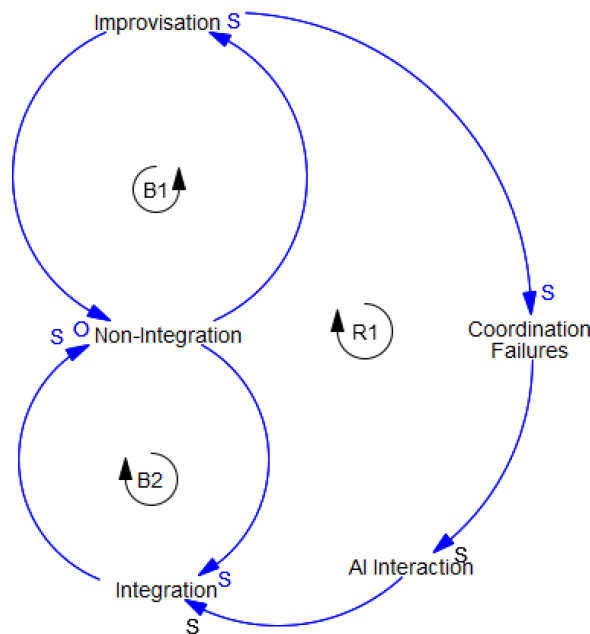


Figure 11. AI to Regulate Team Integration

As improvisation occurs due to non-integration, coordination failures appear as a side effect. These failures undermine individual training and performance as uncommunicated knowledge exists but fails to be distributed. Unable to act on this unobserved knowledge, decisions are quickly made without understanding the full framework. This is seen in the Operation Provide Comfort shutdown of the helicopters as the coordination failures prior to the mission resulted in the F-15 pilots being unaware of friendly aircraft operating in their TAOR. However, inserting a limited memory reactive machine AI interaction into this problematic causal loop helps to reduce coordination failures and thereby increases integration (Beck, Dibbern, & Wiener, 2022). Self-aware machines which, are not yet a reality, will offer even more opportunities (Beck et al., 2022). Deep Blue, IBM's supercomputer, which handled chess grandmaster Garry Kasparov is an example of a reactive machine (Beck et al., 2022). It is important to realize that reactive machines on their own possess no memory capability; however, limited memory machines do have the ability to remember the past through observations and programming (Beck et al., 2022). Combining the two types of AI would allow for a system to monitor the networks of the AWACS, *EAGLE FLIGHT*, and *TIGER* teams, which could help integrate communications. In theory, the machine teammate could introduce to the F-15 pilots the knowledge that friendly helicopters are operating in the TAOR, a crucial fact that was missed in the Air Tasking Order.

As technology progresses, an AI teammate could potentially play the role of an autonomous agent that has the ability to be performative and make decisions on its own (Russell, 2019; Beck et al., 2022). As Beck et al. stated, "what has been created in a socio-technical way by implementing patterns into a technical system, turns into a techno-social system once operating, where social agents in the organizational environment respond to the technical system and where the system may self-adapt to environmental changes" (p. 2). Until recently, AI has relied on algorithms, sensors, and code but now users are beginning to see AI as software systems that can operate on their own in the background or at the interface level with humans (Beck et al., 2022). This type of AI teammate could provide dynamic, particularly in a similar event such as Operation Provide Comfort.

C. VIGNETTE 2: USS VINCENNES CASE STUDY

1. Background and Overview

In May of 1987, the USS Stark was accidentally engaged by an Iraqi Mirage F-1 by a pair of Exocet missiles, killing 37 sailors (Dotterway, 1992). The resulting investigations and media inquiries left the ship and its leadership disgraced (C-Span, 1987; LaGrone, 2017). This engagement led to a sharp increase in U.S. naval activity in the Persian Gulf alongside an escalating anti-shipping campaign conducted by Iraq and Iran (Dotterway, 1992). Belligerents fought the campaign with sea mines, naval vessels, as well as attack aircraft (Dotterway, 1992). Between January and July of 1988, both belligerents executed a combined total of 69 attacks on commercial shipping in the Persian Gulf (Dotterway, 1992). Supplementing the number of significant actions in the battlespace, intelligence reporting indicated that there would be an increase in hostile activity against U.S. units in conjunction with the Fourth of July holiday (Dotterway, 1992). Thus, the USS Vincennes was operating in a dynamic, uncertain, and dangerous operating environment during the time of the incident.

Adversary tactics, techniques, and procedures (TTP) informed the rules of engagement issued to forces operating in the Persian Gulf. As such, Iranian forces were known to deploy ISR aircraft, such as the P-3, to cue surface attack capabilities (Dotterway, 1992). Hostile attack aircraft, including F-1s and F-14s, were deployed to the area of operations, and had demonstrated the capability to conduct combined arms attack in support of surface units (Dotterway, 1992). Intelligence reporting also indicated that hostile aircraft were not likely to fly routine patterns during the holiday period (Dotterway, 1992). These TTPs as well as the history of the area of operations were well known by the combat information center crew as they began operations on the morning of July 3, 1988.

Ocean Lord 25, a U.S. Navy helicopter, reported taking small arms fire from Iranian fast attack craft at 0945L on July 3, 1988, prompting the USS Vincennes and USS Montgomery to engage these attack craft (Dotterway, 1992). By 1013L, both ships reported that they were fighting a surface engagement against hostile Iranian gunboats (Dotterway, 1992). At 1017L on 3 July 1988, the SPY-1 radar onboard the USS Vincennes identified

an aircraft that took off from a joint military and civilian airfield located at Bandar Abbas, Iran (Dotterway, 1992). Unknown to the combat information center crew of the USS Vincennes was that this aircraft was Iran Air Flight 655, departing on its regularly scheduled flight to Dubai (Dotterway, 1992). Flight 655 was approximately 27 minutes behind its departure schedule, and was directed by air traffic controllers at Bandar Abbas to proceed to its assigned air corridor at an altitude of 14 thousand feet (Dotterway, 1992). Additionally, flight 655 was directed to emit IFF Mode III code 6760, a frequency that was appropriate for a civilian airliner.

From the USS Vincennes' perspective, this airliner flew lower than most other commercial aircraft previously observed, and slightly off center from the center of the air corridor (Dotterway, 1992). When it took off, the identification supervisor consulted the flight schedule, however, the IDS did not know the flight was 27 minutes late, leading him to determine that his track was not Iran Air flight 655 (Dotterway, 1992). Trying to positively identify track TN 4131, the identification supervisor hooked the track, and left the hook at Bandar Abbas for approximately 90 seconds. Unfortunately, this would cause sensors to remain fixed in this area, collecting on nearby data points while the track departed the area, and could give operators a false impression of their track number (Dotterway, 1992).

Interestingly, the USS Vincennes assigned track number TN 4474 to flight 655, while the USS Sides assigned TN 4131 to the contact (Dotterway, 1992). The Aegis system onboard the USS Vincennes ingested the track number from the USS Sides, and altered the original track number, TN 4474, to the USS Sides' number, TN 4131. It would be discovered later that this track number ultimately belonged to the block assigned to the USS Vincennes, and was an unreported link violation committed by the USS Sides. Unfortunately, there was no mechanism to inform operators that Aegis had performed an auto-correlation to merge target numbers (Dotterway, 1992). The common track number for TN 4131 ultimately changed three times in the elapsed time between detection and engagement (Dotterway, 1992).

Scanning his screens, the identification supervisor located onboard the USS Vincennes interpreted an IFF Mode III code 6675, while the onboard Aegis system stored

the actual frequency, Mode III code 6760 (Dotterway, 1992). The SPS-49 air detect tracker (ADT-49) recollected observing both Mode II and Mode III frequencies during the initial detection (Dotterway, 1992). Some members of the combat information center reported seeing IFF Mode II, while others reported seeing IFF Mode 3 (Dotterway, 1992). This may have been an after effect of the hook remaining at Bandar Abbas, and could have enabled a Mode II IFF transmission from an aircraft on the ground near Bandar Abbas to display on the identification supervisor's screens (Dotterway, 1992).

At 1022L, TN 4131 reached the 20 nautical mile decision point. The commanding officer inquired as to the status of TN 4474, not realizing that the track number had changed to TN 4131 (Dotterway, 1992, p. 40). Aegis showed that a member of the crew, sitting at the FC-1 console hooked TN 4474 for five seconds, showing a range of 110 nautical miles, an altitude of 11,900 ft, and an airspeed of 448 knots. (Dotterway, 1992). At the 20 nautical mile decision point, the air detect tracker, 49 ADT, observed an IFF Mode II on his remote control indicator, but not on his character read out (Dotterway, 1992). This observation was mirrored by members of the combat information center crew, who believed the contact to be descending at 1022L (Dotterway, 1992). As a result, several descending altitude calls were made over radio nets 15 and 16 (Dotterway, 1992).

Given the number of coalition partners present in the Persian Gulf, ships were continuously entering and leaving the data link network, carrying new track numbers with them. TN 4474 was in reality, a friendly A-6 that was accelerating and descending at the same time as TN 4131 was nearing the 20 NM decision point (Dotterway, 1992). This track number was brought into the link by the HMS Manchester, and was another unreported link violation (Dotterway, 1992). The kinematics of this aircraft matched what the combat information center crew of the USS Vincennes associated with TN 4131.

Near simultaneously, the USS Sides observed TN 4131 ascending to 11,000 feet (Dotterway, 1992). The USS Sides' air tracker was aware that the contact is a commercial airliner, and attempts to notify the tactical action officer (TAO), before being told "Shut up, you're making too much noise" (Dotterway, 1992, p. 42). The USS Sides' TAO identified that flight 655 was squawking Mode III, 6700 block, indicative of commercial aircraft (Dotterway, 1992). At this time, the determination of the USS Sides' commanding

officer was that the track was not a threat to the USS Sides based on flight path, lack of electronic emanations, historical precedent, and a lack of adversary anti-surface warfare capabilities (Dotterway, 1992). Unfortunately and despite being a critical member of the anti-air warfare team, the commanding officer's judgement was not available to the combat information center crew or captain of the USS Vincennes.

Despite interpretations that TN 4131 was accelerating and descending, it was actually ascending. In reality, at 14 NM from the Vincennes, Iran Air flight 655 was climbing through 12,000 feet at a speed of 382 knots one minute prior to missile impact (Dotterway, 1992). However, the international air distress operator recalled the contact being at a height of 7,700 traveling at a speed of 450 knots (Dotterway, 1992). The anti-air warfare TAO indicated that he "heard continuous reports of declining altitude." (Dotterway, 1992, p. 43). At 1024L, the Vincennes' system data indicated TN 4131 was at a range of 12 nautical miles, an airspeed of speed of 380 knots, and was climbing through an altitude of 12,000 feet. The identification supervisor observed TN 4131 at 445 knots and descending at 7,800 feet. (Dotterway, 1992, p. 44). Single channel radio calls were made by the combat information center crew during every minute of the engagement on both the international air distress and military air distress nets (Dotterway, 1992). The USS Vincennes did not receive a reply, and interestingly, the Bandar Abbas tower failed to relay international air distress warnings to Iran Air flight 655 (Dotterway, 1992).

From first contact to the USS Vincennes commanding officer's decision to engage, 3 mins 40 seconds elapsed. Most of the commanding officer's time during those minutes were consumed by a gun battle with Iranian gun boats (Dotterway, 1992). At 1024L, the air engagement concluded when the captain of the USS Vincennes turned the missile firing key, resulting in the shutdown of TN 4131 (Dotterway, 1992). Iran Air flight 655 was at an altitude of 13,500 feet, 3.35 nautical miles from the center line of its air corridor, at a speed of 383 knots (Dotterway, 1992).

Fascinatingly, the data stored in the Aegis system does not mirror the estimates and recollections of the combat information center crew aboard the USS Vincennes. At 1022L, the crew's mental model begins to rapidly diverge from actual events, which were recorded by Aegis (Dotterway, 1992). The majority of the combat information center crew

recalled that TN 4131 was descending and rapidly accelerating, while Aegis demonstrated that TN 4131 was steadily accelerating and climbing (Dotterway, 1992). The combat information center crew experienced significant difficulties in tracking the correct track designator, ultimately confusing TN 4131, Iran Air Flight 655, with TN 4474, a U.S. Navy A-6 Intruder (Dotterway, 1992). The anti-air warfare TAO, and the anti-air warfare coordinator were two critical crew positions that should have caught this issue; however, they were task saturated and unable to execute their assigned duties.

The performance of the combat information center crew aligns with Sterman's (2000) assessment that there is self-reinforcing feedback between peoples' expectations, mental models, and the evidence that they are presented with. The authors concur with Dotterway's (1992) assessment that given the ambiguities faced by the crew in the combat information center, and the time pressures felt by decision-makers, the downing of Iran Air flight 655 was assessed by senior leaders as a logical, though certainly undesired, outcome of stemming from interactions in a complex system (Apple, 1988; C-Span, 1992). This prompts the question, "how does one of the most sophisticated anti-air warfare platforms on the sea, manned by a combat qualified crew, accidentally shoot down an airliner?"

2. Causal Loop Diagrams

a. The Fate of the USS Stark Influences Current Decision Making

(1) Modeling the "As-Is" System

A series of investigations followed the Stark incident, which resulted in the commanding officer and TAO receiving non-judicial punishment as well as letters of reprimand (Black, 1987). Both were essentially disgraced, and forced out of the U.S. Navy. The executive officer was similarly punished. An additional output of these investigations was the promulgation of a new set of rules of engagement for U.S. Navy vessels operating in the Persian Gulf. These revised orders emphasized that commanders were responsible to defend their ships against perceived hostile intent, hostile acts, and emphasized that they may engage in anticipatory self-defense, or firing before being fired upon (Dotterway, 1992; C-Span, 1992). Interestingly, the crew of the USS Stark appeared to be saddled with conflicting priorities during this time. It appears that crew members in critical positions,

such as the TAO in the combat information center, were required to juggle administrative collateral duties while standing watch (Dotterway, 1992). The combat information center was also not fully staffed at the time of the incident; the intentional and unintentional manning shortfalls were acknowledged as readiness issues during subsequent investigations into the incident (Dotterway, 1992; Black, 1987).

The informal guidance given to units operating in the Persian Gulf in 1988 was that U.S. Navy ships cannot afford to take the first hit and cede the initiative, as well as the information narrative, to an adversary (C-Span, 1992). It was made abundantly clear that regardless of personal culpability, the command team was responsible for everything that happened or failed to happen aboard ship (Black, 1987). Thus, the fate of the USS Stark was weighing on the mind of Captain Rogers during the incident (Dotterway, 1992). Figure 12 discusses these pressures and illustrates how successful actions lead to rewards for the crew of the USS Vincennes.

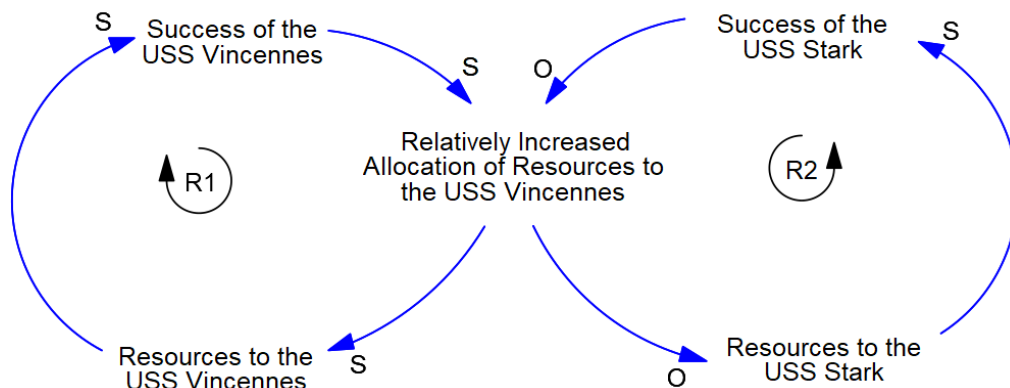


Figure 12. Avoiding the Fate of the USS Stark

This causal loop diagram is based upon the success to the successful archetype (Kim, 2000). Within this archetype, one entity is provided with additional resources, which ultimately increases the likelihood of success for that resourced entity (Kim, 2000). When the resourced entity is successful, it justifies providing more resources, which in turn sets conditions for success (Kim, 2000). In essence, this additional resourcing creates a single

winner, and limits the success of any entity that does not receive such resources (Kim, 2000).

While the USS Stark and the USS Vincennes were not in the same command, and were not deployed at the same time, the loss of life and materiel, as well as the public spectacle resulting from the incident were components of Captain Rogers' decision calculus during the incident (Dotterway, 1992). The promulgation of a clear, consistent set of rules of engagement enabled Captain Rogers with a broad set of authorities to protect his ship and crew (Dotterway, 1992). The crew of the Vincennes was also supported with a deluge of intelligence reports that cautioned a heightened state of readiness during the days nearing the U.S. Independence Day holiday (Dotterway, 1992). His ability to take positive measures and engage without being fired upon could have theoretically prevented the USS Vincennes from suffering the same fate as the USS Stark. Due to the fact that the USS Stark was not supported with the same unambiguous set of authorities as the USS Vincennes, and was not provided the same level of support in the form of intelligence reporting, it could be argued that it was not afforded with a critical resource needed to operate in an incredibly dynamic operational area.

As the crew of the USS Vincennes triumphed in its battle with Iranian fast attack craft, and was successful in defending itself against a perceived airborne threat, the crew was awarded with impact awards such as the legion of merit and the Navy and Marine Corps commendation medal (Moore, 1990). Because of the perceived success of the crew in an uncertain and dangerous deployment, key leaders were rewarded with key billet opportunities and were allowed to advance in their careers. The opportunity provided by these assignments permitted additional successes to the successful.

(2) Modeling the "To-Be" System

Correcting the reinforcing behaviors found in the success to the successful archetype is possible. The fix actions require a review of the goals and objectives that the actors within the system are working toward (Kim, 2000). With an understanding of the goals, an analysis as to why the underlying system is rewarding one team while punishing the other is required (Kim, 2000). Solutions to the problem generally require the removal

of one half of the archetype, and then optimizing resources to remove resource competitions (Kim, 2000).

The line of inquiry into the underlying conditions that enabled the USS Stark to be struck by a pair of missiles identified a crew staffing problem within the combat information center (Dotterway, 1992). Given the problem that multiple crew members were away from their watch sections and consoles, the fix actions used to address the success to the successful archetype would not provide for a fundamental solution. However, applying the shifting the burden archetype does provide for such a solution.

Within a shifting the burden archetype, problems are solved by treating their symptoms as opposed to applying solutions that combat the fundamental problem (Kim, 2000; Senge, 2006). Over time the side effects continue to manifest themselves and may grow considerably worse (Kim, 2000; Senge, 2006). This trend sets conditions for problem solvers to treat side-effects as the root-cause of the problem as they continue forego fundamental solutions (Kim, 2000). The remedy for this cycle is to deploy solutions that address the fundamental problems that will ultimately lessen, if not totally cease, related symptoms over time (Kim, 2000). These descriptions are reflected in Figure 13.

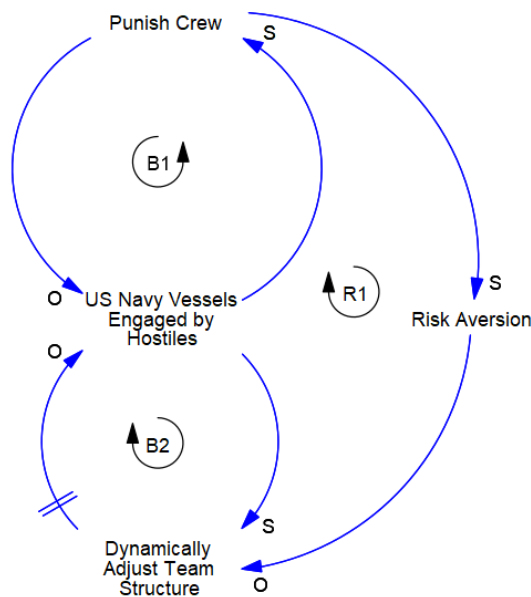


Figure 13. AI Regulates Team Structure

In balancing loop one, the USS Stark incurred fatalities after an accidental engagement by an Iraqi aircraft (Dotterway, 1992). This caused the U.S. Navy to impose punishments on the leadership of the USS Stark (Black, 1987). This promulgated additional guidance that U.S. Navy ships were to avoid taking the first hits in an engagement (C-Span, 1992). More specifically, the rules of engagement were altered, and certain self-defense authorities were emphasized after the engagement of the USS Stark (Dotterway, 1992; C-Span, 1992). Thus, the side effect of risk aversion emerged in reinforcing loop three. However, careful study of Dotterway's (1992) minute-by-minute account of the events onboard the USS Stark in the minutes leading up to its damage by Iraqi aircraft indicated that there was a fundamental problem with crew manning within the combat information center. Captured in balancing loop two, one possible fix action calls for correcting the problems within the combat information center as a way to mitigate manning challenges. A machine partner may assist with the fundamental solution to this problem: adjusting a team's structure.

At the time of the engagement, the operator for the combined antenna system as well as the close in weapons system was not present in the combat information center (Dotterway, 1992). Because this crewman was gone for an extended period of time, another member of the combat information center was dispatched to locate him (Dotterway, 1992). This posed a problem for manning within the combat information center because the weapons control officer and combat information center watch officer billets were combined into a single position (Dotterway, 1992). As the crewman responsible for self-defense weapons temporarily departed the team structure, this set conditions for task overload in critical positions essential for self-defense.

This scenario provides a unique opportunity for the injection of AI as a machine partner. While one could argue for the use of an AI teammate as a means to optimize resource distributions between subordinate teams, there is a more beneficial use for this AI teammate (Bansal et al., 2019; Russell, 2019; Stumborg et al., 2019; NASEM, 2021). The fundamental problem in this vignette was that the human team was unable to perform critical tasks given an ad-hoc team structure. A machine teammate could support the team by maintaining and disseminating information about current team structures, tasks, and

goals. In human teams, Cannon-Bowers and Salas (2001) noted that possessing similar information about tasks, organization, teammates, roles, responsibilities, and predictions are attributes of a team's shared cognition.

A machine teammate could support interaction dynamics in teams by mapping current team structures to teammate roles and responsibilities, and then communicating information about gapped capabilities to teammates. This would flatten the traditional communications hierarchy of the combat information center and improve the team's ability to overcome perturbations (van den Oever & Schraagen, 2021). Such action could enable the team to dynamically provide backup for missing roles and enable it to adapt both tasks and structures to meet situational requirements. With additional information about the organizational structure, tasks, capabilities enabled by personnel, the HMT will be able to continue to make decisions and make progress toward shared goals despite the presence of perturbations (Kennedy, 2021; NASEM, 2021). Akin to Gombolay et al. (2016), these decisions could be further supported if the machine partner was to present prioritized recommendations for task and crew member allocations alongside an assessment of decision quality.

If a machine teammate were able to bring such capabilities to the team, the team may be able to dynamically adjust team structure. More specifically, the roles, responsibilities, and authorities of the team members could be adjusted to meet the demands of the situation. This represents one way to attack the fundamental problem described in balancing loop two of Figure 13. As the team adjusts to the situation and environment, it may be able to handle dynamic challenges such as unexpected hostile engagements. Such actions would reduce the effects of symptoms and side effects documented in balancing loop one and reinforcing loop three.

b. Time Pressure Loops Erode Performance

(1) Modeling the "As-Is" System

One of the themes present in Admiral Crowe's (C-Span, 1992) and Dotterway's (1992) assessments was that the crew was under significant time pressure during the incident. The rapid velocity of Iran Air flight 655 relative to the USS Vincennes meant that

it approached Captain Rogers’ 20 nautical mile decision point in mere minutes. Admiral Crowe, the Chairman of the Joint Chiefs of Staff from 1985 to 1989, had the gift of hindsight when he sat in a simulated recreation of the events of the shutdown, and still noted that “those seven minutes happened awfully fast” (C-Span, 1992). Figure 14 discusses these time pressures, and places them into context within the larger system.

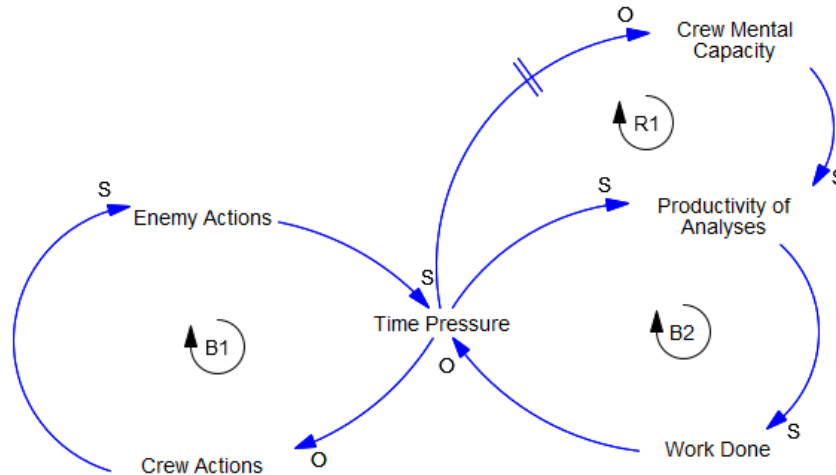


Figure 14. Time Pressure Loops

This causal loop diagram is based upon the limits to success archetype. In this archetype, continued efforts lead to increasing performance, however, the system encounters a limit that erodes performance or causes a decline (Kim, 2000). In the model, a range of activities occur within a fixed allotment of time. The mental capacity of personnel, both friendly and enemy, decreases as the time pressures increase. This limits the productivity of analyses, and the amount of work done. Ultimately, the amount of work completed decreases as time pressure increases.

When the crew in the USS Vincennes combat information center launched anti-air missiles against Iran Air Flight 655, they perceived track TN 4131 as an Iranian F-14 (Dotterway, 1992). The crew was under significant time pressure given the threat-based intelligence they possessed, observed adversary TTPs, and the decreasing range between TN 4131 and the ship. TN 4131 was observed lifting out of Bandar Abbas, a joint military-civilian airfield, at 1017L, and it crossed the 20 nautical mile engagement decision point at

1022L (Dotterway, 1992). During this time, the crew conveyed five warnings to TN 4131 over international and military air distress radio nets, validated flight schedules, and exchanged a series of communications with their higher headquarters (Dotterway, 1992).

Complicating the crews' analyses, Iran Air flight 655 departed from Bandar Abbas 27 minutes late, was 10 miles off center of its approved air travel corridor, and was flying at a relatively low altitude for an airliner in that area of operations (Dotterway, 1992). Members of the combat information center recalled observing inconsistent IFF modes and codes emanating from track TN 4131, but had to review different sources to validate the identity of the contact.

In balancing loop number one (B1), crew and hostile actors' actions are balanced by time pressures. In this context, the actions of both actors can be characterized as the tasks that they complete. The friendly actions and hostile counteractions presented in this loop explains how the USS Vincennes operated in its environment. More specifically, this demonstrates the reality that one agent's actions spur reactions and counteractions from other actors present in the battlespace (Kennedy, 2021).

In balancing loop number two (B2), the productivity of the combat information center crew reinforces their ability to perform analyses and execute tasks. The amount of work that can be completed is balanced by the amount of time available. As time pressure increases, productivity increases. It is worth noting that there is a limit to this linkage between time pressure and productivity. In the long term, continuous time pressures lead to burnout, which may erode productivity (Sterman, 2000; Senge, 2006). The combat information center crew in the USS Vincennes could only perform a limited number of actions between the 1017L takeoff period and the 1024L engagement.

Reinforcing loop number one (R1) compounded the effects of the crew's mental capacity and productivity with time pressures. In this loop, time pressures increase as the crew's mental capacity decreases. This means that as time progresses, crews' mental capacity and productivity ultimately decrease.

Although it may be argued the crew of the USS Vincennes were trained to handle this scenario, the combat information center crew of the USS Vincennes were in a novel

situation. The ship was fighting a gun-battle against high-speed Iranian attack craft, one of the most dangerous threats in the area of operations, and sustained a casualty in one of its weapons mounts (Dotterway, 1992). The combat information center crew was also in a non-doctrinal configuration that did not mirror the structure they used during pre-deployment evaluations (Dotterway, 1992). The team's ability to communicate and coordinate information were limited due to an inundation of traffic on its internal communications channels. Radio nets 15 and 16 were monitored by key leaders and were almost completely saturated, with ambiguous communications being frequently injected into the medium by multiple stations (Dotterway, 1992).

To the detriment of the entire team, the anti-air warfare TAO and the anti-air warfare coordinator were bogged down with a significant amount of tasks, and were unable to execute their normal responsibilities (Dotterway, 1992). These communications and coordination breakdowns adversely affected the crew's interactions, and led them to derive an inaccurate mental model of the situation. This was evident when Captain Rogers inquired about TN 4474, and was demonstrated again during the wildly divergent recollections of combat information center crewmen after the incident (Dotterway, 1992). Compounding the problem, the anti-air warfare coordinator did not hold the requisite set of qualifications, was not trusted by the TAO, and as a result, was relegated largely to pushing buttons as a console operator (Dotterway, 1992).

This non-doctrinal deployment of the crew members, in concert with the lack of trust between various team members may have detracted from backup behaviors that could have distributed the workload between different crew positions in the combat information center (Marks et al., 2002; Smith-Jentsch et al., 2009). The divergence of mental models, exacerbated by time pressures may have also eroded coordination processes, may have also contributed to poor team performance (Marks et al., 2002; Smith-Jentsch et al., 2009).

(2) Modeling the "To-Be" System

In Figure 15, the time pressures impacted crew members mental capacity, reduced the benefits of their analyses, and reduced the work that the team was able to accomplish. In attempting to model the "to-be" system, it is unrealistic to do away with time pressures

while two belligerents are operating against each other. Kennedy (2021) noted that time is the dominant factor in competitive situations, such as combat operations, that constrains teams as they make decisions. Thus, the time pressures cannot be assumed away, and corrective actions must improve other variables present in the system. Specifically, efforts to improve the combat information center crew’s mental capacity and productivity could improve the amount of work that they are able to accomplish and may partly, if not completely mitigate the adverse effects of time pressures.

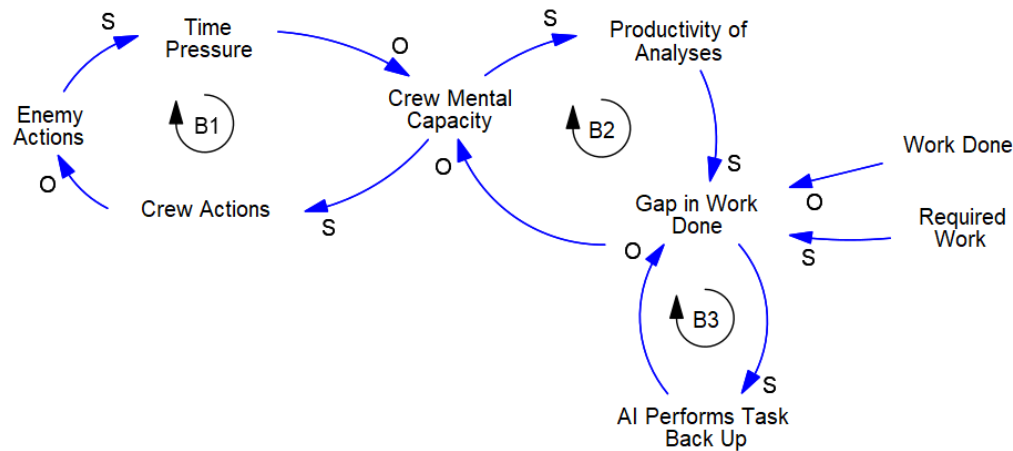


Figure 15. AI Back Up Behavior Completes Tasks and Improves Taskwork Capacity

Due to the increasing gap of uncompleted work as time pressures increase, teams fall further behind in tasks and productivity suffers. This is problematic given the series of tasks that a combat information center crew must perform to make sense of their environment. In Figure 14, incorrect and incomplete information was promulgated within the combat information center, and due to time pressures, this information could be validated by leaders who owned verification tasks (Dotterway, 1992). Thus, inaccurate information directly impacted the decision-making of the team’s key leaders (Dotterway, 1992). Figure 15 illustrates how a machine teammate could be one solution to address the breakdowns experienced by the team in Figure 14. The addition of a machine teammate into the team structure could improve productivity, accuracy of analyses, and may reduce the impacts of time pressures.

Balancing loop three of Figure 15 depicts a potential solution to the team's problem. Within this loop, the machine partner is able to assess the team's task backlog and act upon tasks that align with its strengths. The actions within balancing loop three feed into balancing loop two, enabling the combat information center crew to reduce the number of tasks within the backlog. Such task sharing reduces the burden on the human members of the team, freeing up mental capacity to focus on other important tasks. As tasks that are essential to sense-making and action are increasingly completed within the allotted time, the crew is able to mitigate the effects of perceived time pressures. Finally, by improving the team's mental capacity, the output quality of their tasks may be similarly enhanced.

Actions that the machine partner could perform may have outsized impacts on the team. For example, the machine partner could save crew members time by reviewing and displaying relevant flight schedules and air tasking orders, preventing crew members from having to manually search through records. Another example could be notifying the combat information center crew that the Aegis system took some action, such as correlation of a track, and then display information for tracks of interest on the large screen display. As the machine partner removes tedious, dull, or time consuming tasks from human team members, the crew is able to accomplish more work, and may retain mental capacity for more cognitively demanding tasks.

Improvements to the system are possible because the machine partner is executing team backup behaviors. Akin to teamwork behaviors found in human teams, the machine partner is monitoring the performance of the team and assisting team members as required (Smith-Jentsch et al., 2009). In Figure 14, the human team became unsynchronized over time (Dotterway, 1992). Marks et al. (2002) noted that synchronization and communications challenges are indicative of coordination problems, which may detract from the performance of the team. As the machine partner performs back up the behaviors depicted in Figure 15, it is also assisting the team by regulating its communication dynamics, which may prevent breakdowns in both interaction and coordination behaviors. Specific improvements to communications dynamics will be modelled in Figure 19. Ultimately, as the coordination behaviors improve, the performance of the team is likely to improve as well (Marks et al., 2002).

c. Ineffective Team Structure Prevents Limits Usefulness of Backup Behaviors

(1) Modeling the “As-Is” System

The USS Vincennes utilized a non-doctrinal organizational structure within its combat information center. The major deviations from doctrine were in the formation of the anti-air warfare TAO positions as well as the change in roles for the anti-air warfare coordinator. The U.S. Navy believed that such a setup would help the USS Vincennes, which was in charge of the force’s anti-air warfare functions, to track the complex air picture more effectively and efficiently for JTF Middle East (Dotterway, 1992). As a part of this organizational structure adjustment, the anti-air warfare tactical officer was delegated all force anti-air warfare responsibilities from the commanding officer of the USS Vincennes, which included the holistic U.S. Navy force as well as the USS Vincennes (Dotterway, 1992). This critical crew position was also responsible for the generation of air situation reports and coordination with higher headquarters (Dotterway, 1992). Increasing the task load, the anti-air warfare TAO was also responsible for ensuring the accuracy of the force’s air picture (Dotterway, 1992). This officer was also burdened with the anti-air warfare coordination functions—a task usually assigned to the anti-air warfare coordinator (Dotterway, 1992). As would be discovered during the investigation, this officer was given high-value, time-consuming tasks that are cognitively intensive.

Contrasting this with the anti-air warfare coordinator, it appears that there was a task imbalance. The anti-air warfare coordinator was responsible for the section colloquially called “air alley” (Dotterway, 1992). The air team, or “air alley,” was a three-sailor section comprised of the identification supervisor, tactical information coordinator, as well as the anti-air warfare coordinator (Dotterway, 1992). Normally, the anti-air warfare coordinator would be responsible for air coordination functions, designating targets with the fire control radar, validating aircraft flight schedules, and leading the air team (Dotterway, 1992). Since a change to the combat information center’s team structure removed some of the responsibilities normally assigned to the anti-air warfare coordinator, this sailor instead spent most of their time serving as a console operator (Dotterway, 1992).

The relationship between the critical gatekeeper positions of anti-air warfare TAO and anti-air warfare coordinator is illustrated in Figure 16. These two entities were the key leaders of the anti-air warfare mission for the USS Vincennes. Thus, their interdependency and influence on the larger team is also illustrated in Figure 16.

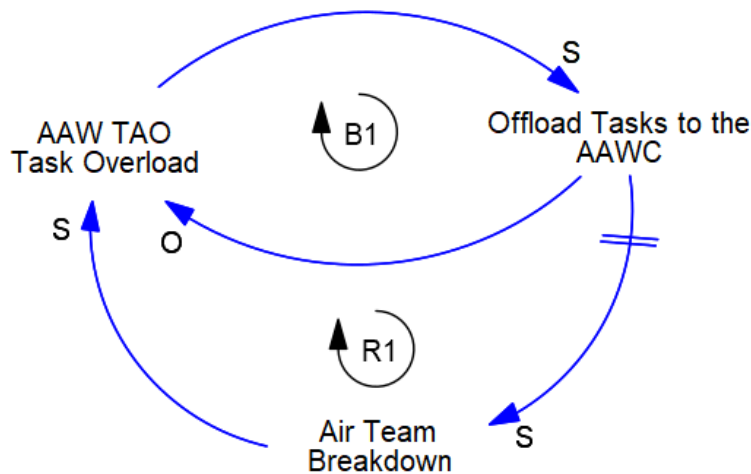


Figure 16. Ineffective Team Structure and Behaviors

This causal loop diagram is based upon the fixes that fail archetype. In the fixes that fail archetype, a problem emerges in the system, and a fix action is applied to address the symptoms of the problem (Kim, 2000). Fix actions are “quick fixes,” or ones that merely treat symptoms as opposed to the fundamental cause of the problem (Kim, 2000). While the symptoms appear to disappear after the fix action, the unintended consequences of the problem eventually exacerbate the problem, causing symptoms to re-emerge at their previous level, or in some systems, to manifest as worsening conditions (Kim, 2000). As such, the system eventually breaks down because of the incomplete nature of the solutions used to address the problem.

Interestingly, it was revealed during the incident investigation that the anti-air warfare coordinator was perceived to be a weak leader by the combat information center watch team, including those sailors of the air team (Dotterway, 1992). This operator was also characterized as inexperienced, and did not possess the requisite qualifications to hold the anti-air warfare coordinator billet (Dotterway, 1992). The revised tasks assigned to this

position, primarily console operation tasks, appear to support the capabilities of the sailor assigned to this role, however it contributed to an insidious breakdown within the combat operations center. When the anti-air warfare TAO became task saturated, his assigned tasks were delegated to the anti-air warfare coordinator. These additional taskings overwhelmed the anti-air warfare coordinator, who could no longer effectively lead the air team. This caused a failure in the primary function of the air team, which was supposed to validate the identification and track profile information collected by the Aegis system (Dotterway, 1992).

As the anti-air warfare TAO became consumed by the communication and coordination tasks, he began to offload the management of the air picture to the anti-air warfare coordinator. This practice is reflected in the B1 loop. This crew position was already occupied by console operations, leading the air team, validating flight schedules, and other tasks. Thus, information overload and analysis paralysis occurred in these critical gatekeeper positions (Dotterway, 1992). Because the TN 4131 kept closing with the USS Vincennes, additional warnings on distress radio channels were required. As TN 4131 closed within the 20 nautical mile mark, the anti-air warfare TAO was increasingly consumed by communications with higher headquarters, and the anti-air warfare coordinator was increasingly focused on trying to deploy the fire control radar to target TN 4131. Additionally, increased communications, coordination efforts, and interactions happened within the USS Vincennes combat information center as the TN 4131 closed with the ship. These compounding demands placed upon the air team resulted in ambiguous communications, the misidentification of air tracks, problems in effectively using the Aegis system, as well the failure to identify multiple tactical data link network violations. This cycle of continuous breakdowns is codified in reinforcing loop one (R1) of Figure 16.

Vital tasks such as the validation of the air team's track information and the development of a correct air picture failed to occur as a result of the overwhelm. In fact, because of the breakdown in the air team, the combat information center crew aboard the USS Vincennes was completely unaware of the presence of a U.S. Navy A-6 intruder conducting a surface combat air patrol 110 nautical miles away, instead mistaking its identifier, TN 4474, with that of Iran Air flight 655 (Dotterway, 1992). The

misidentification of TN 4131 and the misreading of its kinematics ultimately contributed to the shutdown.

Evaluating the system through a team dynamics lens, there are several factors that contributed to the poor performance of the USS Vincennes combat information center team. While it was believed that the ad hoc structure that was adopted for the deployment would allow the team to effectively manage the complex air picture, in reality, the non-doctrinal structure prevented the team from achieving optimum performance (Dotterway, 1992). The USS Vincennes and its combat information center crew had completed 8 major training exercises between October 1987 and May 1988 (Dotterway, 1992). The team structure, roles, and responsibilities were fundamentally altered in May 1988 when the ship was assigned to JTF Middle East. Altering this team structure changed the teamwork and taskwork dependencies, backup behaviors, as well as communication, coordination, and interaction dynamics. Disrupting the carefully tuned dynamics and behaviors of experienced teams may cause the team to fail to overcome novel challenges in its environment (Gorman, Stevens, Galloway, Willemsen-Dunlap, & Halpin, 2020). Gorman et al. (2010) asserted that these challenges generally evolve gradually, causing process adaptations over time, implying that teams will be forced to plan, decide, and act entirely new conditions.

The most evident result of the adopted ad hoc team structure was that it rendered doctrinally-ingrained dependencies essentially useless. When novel conditions emerge in the environment, experienced teams are able to adjust their communications, coordination, and interactions to handle these perturbations (Gorman et al., 2020; Gorman, Cooke, & Amazeen 2010). Gorman et al. (2020) demonstrated that inexperienced teams facing unexpected challenges frequently failed to effectively reorganize themselves to combat breakdowns, and as a result, struggled to achieve their goals.

When the primary gatekeepers and tenders of the air picture were no longer able to perform their assigned roles, they disrupted the interaction dynamics that were trained into the team. Established teams, such as those operating in the combat information center are very likely to possess a similar understanding of inter-role responsibilities (Marks et al., 2002). For example, as airborne radar tracks such as TN 4131 populated onto operator

screens, the combat information center crew expected the anti-air warfare TAO to perform his assigned task of independent information and kinematics verification for the track (Dotterway, 1992). Captain Rogers certainly expected this role to be filled, and depended upon the outputs of these analyses in his own decision making (Dotterway, 1992).

In the DON, the roles and responsibilities of each crew position are formally codified, and members' performance in these roles is evaluated prior to deployment. Having similar ideas of each members' responsibilities enables all members of the team to visualize how each person contributes to the overarching goal of the team (Marks et al., 2002). These visualizations are a component of team mental models, and dictate the task dependencies as well as synchronization requirements (Marks et al., 2002). This understanding includes knowledge of team weaknesses, and is essential when relying on backup behaviors (Marks et al., 2002). More specifically, it helps team members predict what their fellow teammates will do in a situation, and what type of help they may need (Marks et al., 2002). In ideal circumstance, such back up actions improve the team's performance (Marks et al., 2002). As demonstrated, the USS Vincennes was not in ideal conditions.

When the anti-air warfare TAO began to delegate his tasks to the anti-air warfare coordinator, the team was demonstrating backup behaviors. These behaviors are essential to team performance, especially in time critical and high-stakes situations (Smith-Jentsch et al., 2009). While it is acknowledged that backup is critical to a team's effectiveness, it works best when members have been working together for an extended period, possess a similar mental model of the task, understand team coordination processes, and are familiar with the person that they request backup from (Marks et al., 2002; Smith-Jentsch et al., 2009). The performance and qualifications of the person with excess capacity play a role in these requests. Teammates are less likely to request assistance when they know other members of the team are busy (Smith-Jentsch et al., 2009). Members are also less likely to request backup from someone whom they perceive as ineffective (Marks et al., 2002).

The relationship between backup behaviors and team performance explains why the USS Vincennes combat information center crew was unable to dynamically work through the problem presented by Iran Air flight 655. The anti-air warfare coordinator was

perceived as a weak leader, did not possess the requisite watch stander qualifications to function in their role, and did not have the excess capacity to handle the anti-air warfare TAO's task overflow. As a result of this overload, the anti-air warfare coordinator failed to perform his duties. This breakdown resulted in communications problems, contributing to increased loads on single channel radio nets 15 and 16, eroding the team's ability to communicate necessary information to the rest of the team (Dotterway, 1992). Additionally, members of the air team began to issue ambiguous communications, which further eroded the ability of the combat information center to make sense of the environment (Dotterway, 1992). In some cases, altitudes and airspeeds were shouted across the compartment with no intended recipient or accompanying track number (Dotterway, 1992). With these communications processes disrupted, the team was unable to form an effective mental model of their environment.

Compounding communications problems, the air team failed to proficiently operate the Aegis system. As the team processes began to erode, the actions of certain crew members disrupted sense-making processes that were essential for developing accurate mental models. While the Aegis system was notorious for requiring significant amounts of operator intervention and interpretation, the air team was familiar with the limitations of the system (Dotterway, 1992). Members of the team hooked different radar tracks, and failed to move these hooks throughout their target interrogation process (Dotterway, 1992). This caused the Aegis to present misleading information about the targets that the combat information center crew thought they were interrogating (Dotterway, 1992). For example, the communication from the identification supervisor identifying TN 4131 was based on inaccurate IFF data that identified Mode II readouts from a hook stuck at the Bandar Abbas airfield (Dotterway, 1992).

Among the most glaring breakdowns in crew interactions with the Aegis system was the failure of the crew to catch the link violations that led to the changing target numbers correlated in Aegis (Dotterway, 1992). The backup behaviors exercised between the anti-air warfare TAO and the anti-air warfare coordinator, both of whom were overloaded, prevented the information verification task from occurring. Thus, the many

interaction processes failed to provide the information needed to develop an accurate mental model, and impeded the team's goal.

Coordination dynamics are essential to a high-performing team. Coordination accounts for the interdependencies between tasks, and ensures that teams are able to synchronize their actions and accomplish goals (Marks et al., 2002). Generally speaking, the more interdependent the tasks, the more that teams will come to rely on coordination mechanisms to ensure that they are functioning as intended (Marks et al., 2002). When the communications between crew positions in the combat information center broke down, it no longer became possible to share team interaction information. This was further compounded by the aforementioned Aegis operations problems. The team in the combat information center was not able to process the actions of those around them, leaders failed to regulate actions, and the air team became unsynchronized as a result. This was evidenced by the captain of the ship asking about incorrect track numbers, miscorrelated information communicated into the medium, and the failure of the crew to perform their assigned roles. As the shift in workload occurred due to the backup behaviors, it is possible to detect minute by minute breakdowns in the combat information center crew's coordination processes and a subsequent inability to develop and accurate mental models (Dotterway, 1992).

The failings of the team to optimize communication, coordination, and interaction dynamics contributed to a flawed understanding of the external environment, and directly contributed to the shutdown of Iran Air flight 655. As more work was offloaded to different roles within the team, the team itself was unable to adapt, and failed to correct deficiencies. The root cause of the breakdown in team processes stems from the ad hoc team structure utilized by the combat information center crew. Modifications to the doctrinal structure ingrained into the team disrupted its ability to reinforce critical crew positions. As such, when the team needed to implement adaptive behaviors, it was unable to effectively do so, which reinforced the adverse conditions present in the combat information center. This reinforcing loop placed more work onto crew positions that were already overloaded, and as a result led to the development of a poor mental model. This

inaccurate mental model resulted in the prosecution of what the crew believed was an Iranian F-14, but was actually an Airbus laden with 290 souls (Dotterway, 1992).

(2) Modeling the “To-Be” System

In the fixes that fail archetype presented in Figure 16, rapid fixes exacerbate the problem within a system (Kim, 2000). Solving this dilemma requires one to set the symptoms aside and study the system. This task may be difficult to perform, as the unintended consequences of previous actions and the problem’s symptoms may be more easily recognizable (Kim, 2000; Senge, 2006). In some cases, the effects of these symptoms and side effects become pronounced and overshadow the original problem (Kim, 2000). Ultimately, the “as-is” model depicted in Figure 16 may be corrected by shifting focus away from the symptoms of the problem and attempting to solve the fundamental challenges that underly the system (Kim, 2000).

Within the “as-is” system depicted in Figure 16, the ad-hoc team structure carried into the engagement set conditions for task overload within critical combat information center positions that were essential to anti-air warfare functions (Dotterway, 1992). While team structure and backup behavior challenges were certainly present within the combat information center, the fundamental problem present in the combat information center’s air team lies in its interaction dynamics. Ultimately, the challenges that the combat information center crew faced stemmed from a situation and an environment where training conditions did not match the operational environment. This meant that previously rehearsed procedures were unable to adapt to the dynamic tasks required by the environment. Interestingly, this problem was discussed by Gorman et al. (2010), whose findings suggest that in such novel circumstances, teams should be adaptive and possess the ability to vary their interactions to meet the demands of their environment.

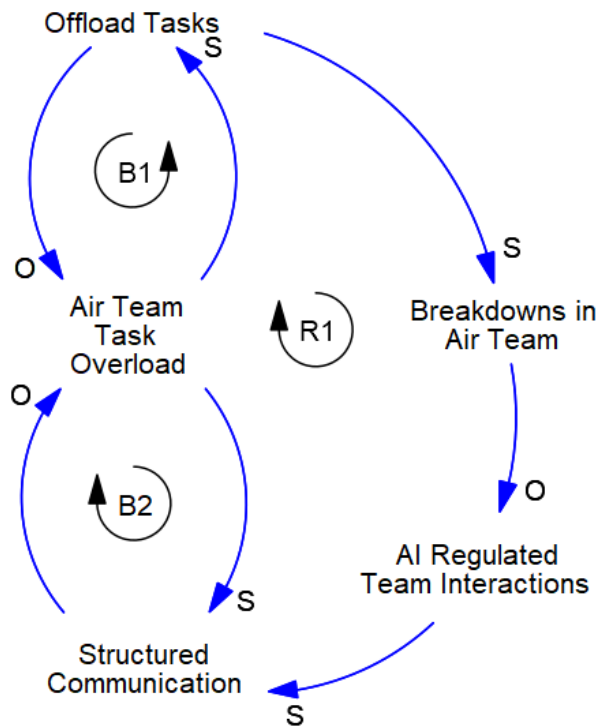


Figure 17. AI Regulates Team Interactions

Balancing loop one (B1) of Figure 17 depicts the combat information center air team’s struggles to keep up with the requirements of their situation. As the team became overloaded, they tried to offload tasks to other crew positions (Dotterway, 1992). When this task overwhelm persisted and the need for communication and coordination increased, the team was increasingly unable to perform and some additional offloading of tasks occurred (Dotterway, 1992). This led to a breakdown in the air team. Specifically, the key leaders responsible for ensuring the accuracy of information reporting and fusing a mental model for further action became overwhelmed and failed to keep up with task demands (Dotterway, 1992).

Further, unregulated interactions injected erroneous information into the team, causing a poor orientation. This was evident when Captain Rogers inquired about the wrong track number and an unnamed crew member provided information without adequate context (Dotterway, 1992). Absent in the communication was the range, bearing, and other information that may have cued the team into the existence of a second aircraft. Thus, the crew failed to communicate and coordinate effectively with each other, preventing

constructive, situational awareness forming interaction processes from occurring (Canan & Demir, 2021).

Reinforcing loop one (R1) depicts a way to prevent team process breakdowns from detracting from the team's performance. Given the task of regulating team interactions, the machine partner may track task allocation, completion statuses, and monitor for task saturation. Monitoring interactions, and then disseminating information to those who need it ultimately supports the ability of team members to forecast how teammates will coordinate future tasks (Marks et al., 2002). If certain positions within the team are saturated, the machine partner may notify other team members and coordinate support for overloaded team members. Team members may be less likely to ask overloaded teammates for assistance when they are able to determine workloads in advance of backup requests (Smith-Jentsch et al., 2009). Thus, regulating and monitoring interactions and then sharing interaction information may correct performance deficiencies and aide teams in functioning as an integrated unit (Cooke et al., 2004). Through optimizing the allocation of tasks within the team, crew members are able to continue to send the outputs of completed processes and exchange information with those on the team who need those outputs for their tasks. These continuing interactions may assist the team in developing accurate mental models and improve team situational awareness (Mohammed et al., 2010; NASEM, 2021).

One of the key challenges in regulating interactions is predicting the series of inputs, processes, and outputs needed by the team to advance towards their goal. Although not every event or task can be foreseen, common occurrences and requirements may be trained into the system. These common events may be considered a "playbook" where the machine may be empowered to execute specific tasks, which may reduce the workload on human team members (NASEM, 2021). When a "play" is activated within the "playbook," the team may be able to determine the task, structure, and interaction processes needed to accomplish the goal (NASEM, 2021). Thus, a "playbook" may streamline task performance and communications while presenting the team with a series of expected outcomes (NASEM, 2021).

“Playbooks” may enable the machine partner to assist in the regulation of interaction behaviors, because it will be able to place tasks and interdependent roles into a sequence. Mapping progress against the forecasted sequence, the machine partner may backup the team when it cannot interact effectively. Backup behaviors could be in the form of regulating the team’s interactions, or completing suitable tasks present in the team’s task backlog. Although modeling the intentions, expectations, needs, and preferences of human teammates is assessed as a challenge for machine partners, these may be implicitly communicated to the machine partner through the designation of a “play” within the team’s “playbook” (Chakraborti and Kambhampati, 2018; Russell, 2019). Such “plays” may be learned by the machine teammate through rehearsing standard operating procedures during exercises, pre-deployment training, as well as daily operations.

Within human teams, the planning, problem-solving, decision, and actions taken are based upon a team’s interaction dynamics (Cooke et al., 2004). Maintaining information about team tasks, equipment, structure, and interactions can stimulate mental model similarity within the team (Marks et al., 2002). With this information, the team may be able to coordinate more effectively, and ultimately improve their performance (Marks et al., 2002). Adaptive teams, such as those found in the DOD, require the ability to dynamically adjust their coordination and interaction mechanisms to respond to novel conditions in their environment (Gorman et al., 2010). The addition of a machine teammate into a team’s structure may aide teams in maintaining an accurate model of their structure, tasks, and goals, ultimately enabling the team to perform the right series of interactions mandated by their operational environment.

d. Poor Communications Behaviors Prevented Effective Situational Awareness

(1) Modeling the “As-Is” System

Lacking situational awareness, the USS Vincennes shot down Iran Air flight 655. Given all of the sophistication of the U.S. Navy’s most recently upgraded “robo-cruiser,” it begs the question of how a combat-ready team failed to make sense of their environment (Dotterway, 1992). While the Aegis system onboard the USS Vincennes correctly assessed

the IFF modes and codes from TN 4131, the crew operating the system did not (Dotterway, 1992). Multiple link violations led to the injection of new radar and track data, which was then correlated by systems onboard the USS Vincennes (Dotterway, 1992). It is worth noting that a limitation of the Aegis system was that it would not provide notifications to operators that autocorrelations and reclassification of tracks had occurred (Dotterway, 1992). Thus, it was the crew's responsibility to manipulate the Aegis system to provide the information that they needed to derive an understanding of the air common operational picture (Dotterway, 1992).

Interestingly, the Aegis system also did not provide all of the information needed on one screen or in one location (Dotterway, 1992). This means that operators would be forced to interact with different menus, flip through different screens, and manipulate a broad range of systems' control mechanisms to execute the tasks that they were assigned (Dotterway, 1992). After the series of interactions with the Aegis system, operators would be forced to piece together all of the data elements and cues in the environment to derive some semblance of situational awareness (Dotterway, 1992; Kennedy, 2021). As Kennedy (2021) notes, sense-making practices such as these, which force operators to work with discrete data elements, do make more information available to a user, however it may contribute to information overload. A human's cognitive capacity is consumed when this information overload occurs. Akin to a saturated buffer, new information is lost for want of processing. Supporting the assertion of information overload, Dotterway (1992) noted that the combat information center crew of the USS Vincennes adopted a non-doctrinal structure to attempt to mitigate the massive volume of air traffic present in the Persian Gulf. The failings of this new team structure will be discussed in a following section.

In the case of the USS Vincennes, the combat information center crew perceived information from its environment primarily from its sensors, such as the SPY-1 radar. These sensors extracted and formatted information from the environment for the human operators to make sense of (Kennedy, 2021). Such an incomplete representation of the operational environment contributed to situational awareness challenges because the ship's sensors only presented the team with a subset of the data set present in the environment (Kennedy, 2021). This data was then subject to processing, filtering, and other operations

before it was displayed for human operators (Kennedy, 2021). Human operators will only be able to perceive what is on their display screens, and will only be able to comprehend a smaller subset of what is available on their system (Kennedy, 2021; Shattuck & Miller, 2006). This comprehension ultimately influences how humans interpret the data presented to them and enable them to form mental models about current and future situations (Kennedy, 2021; Shattuck & Miller, 2006).

This section of the study will demonstrate that the USS Vincennes combat information center crew was unable to communicate effectively, derived an incorrect mental model, and performed actions that did not build a satisfactory level of situational awareness. As a result of this failing, their ability to observe, orient, decide, and act was compromised. Figure 18 specifically models the team's situational awareness challenges, which ultimately contributed to the undesired outcome.

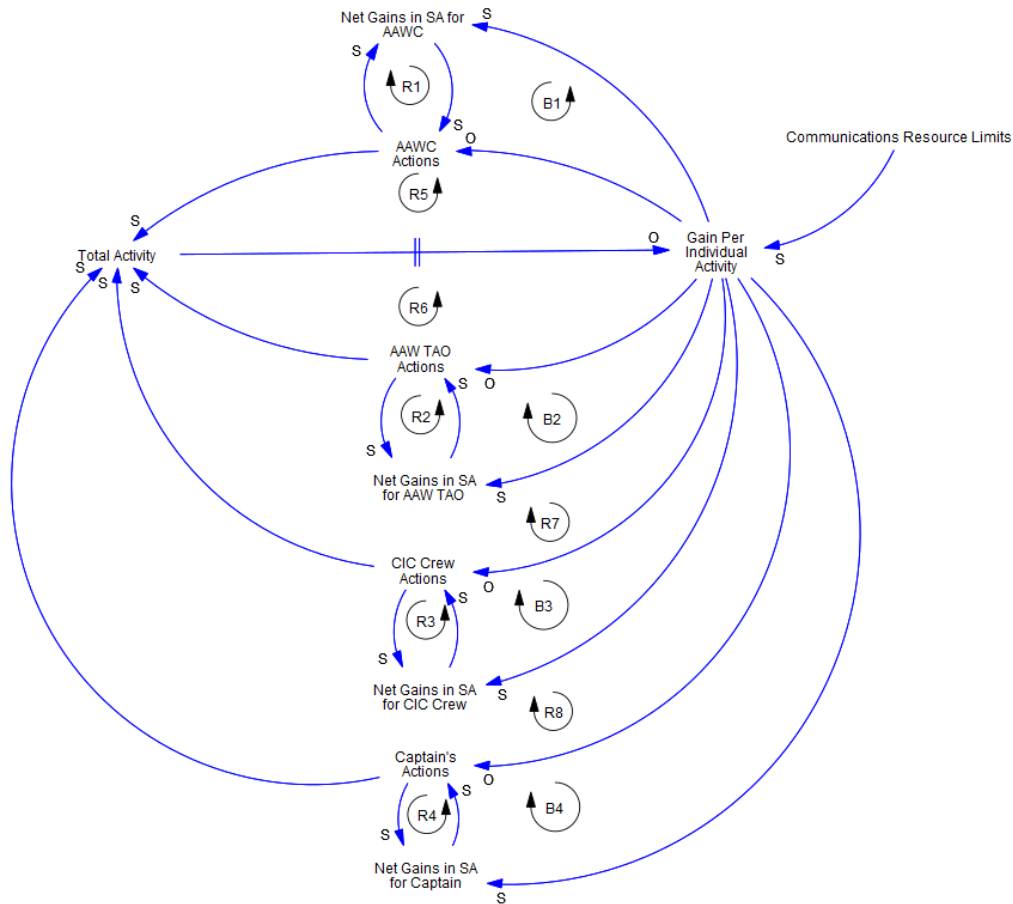


Figure 18. Poor Team Behaviors Erode Communications Patterns and Situational Awareness

Figure 18 is based upon the tragedy of the commons archetype. Within this archetype, each human pursues actions that are individually beneficial, but do not return the anticipated benefits (Kim, 2000). The primary reason that the activities do not create the desired benefits is that the system grows too large, consuming the common resource that supports the activities (Kim, 2000). The system begins to break down when the resource pool reaches saturation, and the humans who are doing the work cannot continue in their tasks because of resource limitations (Kim, 2000).

As will be discussed in the remainder of this section, this model explains how the time competitive nature of the operating environment and the team's ad-hoc use of communications means disrupted its ability to maintain an accurate understanding their environment. This disruption had significant impacts on combat information center crew's

coordination mechanisms, contributing to poor team performance. Situational awareness is based upon a perception of environmental elements, described in terms of time, space, and context of actions (Cooke et al., 2004). Based on the interpretations of these elements, humans extrapolate future actions through forecasting (Cooke et al., 2004). Team situational awareness is based upon a shared mental model, generally containing information about the nature of the task, team status, as well as its equipment (Cooke et al., 2004). Higher levels of team situational awareness are attributed to high performing teams (Cooke et al., 2004). This makes intuitive sense; teams that maintain an accurate mental model of the environment, task, team, and goal are able to interact more efficiently and effectively. The coordination mechanisms of these high performing teams were found to be both implicit as well as explicit, with implicit mechanisms being beneficial in stressful conditions (Mohammed, Rico, & Alipour, n.d.; Smith-Jentsch et al., 2009).

Aboard the USS Vincennes, the combat information center team's goal was to defend themselves from the approaching target, TN 4131. TN 4131's continual progress towards the USS Vincennes ensured that the combat information center team perceived this goal to be in jeopardy. The combat information center crew relied heavily on the Aegis system to provide information about their environment since visual identification and confirmation was not possible (Dotterway, 1992). However, a series of operator errors injected false information into the system, contributing to the development of a flawed mental model. Chappell (2020) noted that if a machine is processing more information than is available to the human partner, the human begins to lose situational awareness. This condition was present on the USS Vincennes. The Aegis outputs and displays only presented crew members with a limited, processed subset of the environment, however the ways in which the crew needed to interact with the system to collect information were not easily performed (Dotterway, 1992).

Further compounding the situational awareness challenge, the combat information center of the USS Vincennes was not arranged in such a way as to promote non-verbal and other implicit communications mechanisms. Instead, most crew members had to use internal and external single channel radio nets to communicate. During the investigation into the incident, it became evident that a non-doctrinal communications structure was

deployed within the USS Vincennes. The number of users on communications nets 15 and 16 saturated the nets, making information exchange difficult to conduct (Dotterway, 1992). More specifically, internal noise, static, and garbled transmissions flooded the medium (Dotterway, 1992). It was noted that the amplifiers powering the nets could not physically handle the load placed on the system, resulting in watch standers being unable to hear each other (Dotterway, 1992). As the technological communications systems degraded, internal communications were handled by shouting across the compartment (Dotterway, 1992).

Further limiting information and communication flows was the poor adherence to procedures. Operators issued ambiguous communications that lacked accompanying information needed to build context. The sender, receiver, track numbers, and kinematics were not shouted or radioed (Dotterway, 1992). When aircraft kinematics, IFF modes and codes, or other information were conveyed into the medium without track numbers, it became increasingly difficult for the combat information center crew to make sense of their environment.

As the combat information center crew interacted with each other, the Aegis system, and their environment, they generated information. This information served as inputs to processes at other stations, allowing the team to coordinate team activities, and adapt to task demands (Mathieu et al., 2010). In ideal conditions, the activities taken by crew positions would contribute to the total amount of activity required to complete a task. This process is reflected in reinforcing loops 5, 6, 7, and 8 in Figure 18. In the case of deriving situational awareness inside the combat information center aboard the USS Vincennes, the number of actions taken by the combat information center crew had the opposite effect.

Individual actions contributed to increasing situational awareness gains for crew members, but placed an intolerable load on the communications system. Reinforcing loops 1, 2, 3, and 4 of Figure 18 illustrate that as different members of the team act, they generate situational awareness. However, after a certain number of actions, the actions begin to deliver a decreasing amount of situational awareness per activity. This process is reflected in balancing loops 1, 2, 3, and 4 of Figure 18. The previously discussed overload on and subsequent failure of internal radio nets 15 and 16 contributed to this sense-making

challenge. Each activity on the net subtracted vital capacity from the system. In addition to tying up the half-duplex medium, so that others may not use it, the total power of the system was reduced. As an increasing number of users signed on to the net and communicated, resources depleted at a similar rate. Once the radio infrastructure was saturated and collapsed, team members resulted to shouting across the compartment. The poor adherence to communication procedures similarly ensured that as each individual communicated to the team, they contributed incomplete information that could not develop a rich understanding of the environment.

Threat-based intelligence drove Captain Rogers' need to decide as TN 4131 crossed the 20 nautical mile mark, and the information collected by the team influenced the decision that they made (Dotterway, 1992; Kennedy, 2021). The team's interactions with Aegis as well as each other informed their situational awareness. Time pressures, an ad-hoc team structure, as well as poor regulation of communications and interaction dynamics prevented combat information center crew members from effectively assembling team situational awareness. The cognitive knowledge of disparate crew members began to conflict with knowledge held by sailors in critical gatekeeping positions such as the captain of the ship, anti-air warfare TAO, and the anti-air warfare coordinator (Dotterway, 1992). This was evidenced by the divergent understanding of track numbers and aircraft kinematics in the minutes leading to the engagement (Dotterway, 1992). The interactions that needed to occur to validate information contained within the system were not possible for the aforementioned reasons, indicating that coordination between the right people at the right times failed to occur. As a result of these shortcomings, the team was unable to develop a requisite level of team situational awareness needed to prevent the engagement of Iran Air flight 655 (McNeese, Demir, Cooke, & She, 2021).

Gorman et al. (2020) noted that experienced teams are able to reorient themselves and dynamically adapt communications patterns to achieve their goal despite the presence of perturbations. When there are too many roadblocks, or perturbations are insurmountable Cooke et al. (2004) noted that coordination dynamics are disrupted, and that time was required to overcome the obstacle. As described by Kennedy (2021), gaining situational awareness is a time competitive activity, and the literature demonstrates that significant

amounts of coordination are required to maintain it, especially in team contexts (Brannick, Prince, Prince, & Salas, 1995; Cooke et al., 2004; Fout & Ploski, 2018; McNeese et al., 2021). Put simply, there is only so much time, and only so many activities that can happen in this time. Poorly developed mental models occur when situational awareness is limited by sensor perceptions and the occurrence of communications and coordination breakdowns (Chappell, 2020; Gorman et al. 2020).

(2) Modeling the “To-Be” System

Kim (2000) noted that diagnosing and then repairing the underlying issue within a tragedy of the commons system archetype requires solutions to address the behaviors of the entire group interacting within the system. The “as-is” system depicted in Figure 18 illustrated that breakdowns in team communications processes contributed to difficulties with team interactions. The team’s interaction challenges ultimately resulted in an erosion of situational awareness within the combat information center. Correcting the interaction and situational awareness difficulties becomes possible by addressing the team’s communications.

Optimizing the combat information center team’s communications to fit within the constraints of their single channel radio network capacity may be done through the growth and underinvestment archetype. Within this systems archetype, performance standards drive the development of capacity to ensure that actual performance meets requirements (Kim, 2000). This archetype encourages teams to anticipate the amount of capacity required to meet their needs before future growth constrains their performance (Kim, 2000). Utilizing this archetype, Figure 19 offers one solution to the USS Vincennes’ combat information center team’s communications challenges.

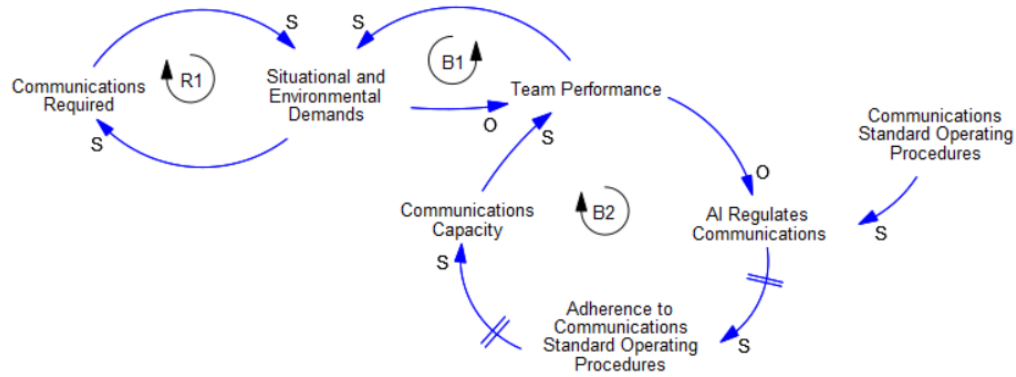


Figure 19. AI Regulates Team Communications

Considering that the combat information center team’s communications were ineffective because of uncorrected lapses in radio net discipline as well as the number of team members saturating two key information channels, any proposed corrective action must also address these problems (Dotterway, 1992). Within Figure 19, reinforcing loop one (R1) describes how ever-changing situational and environmental demands prompt some form of communications to exchange information with the rest of the team. Onboard the ship, situational and environmental information are captured through the USS Vincennes’ sensors, and interpretation of sensor outputs is performed by different roles (Dotterway, 1992; Kennedy, 2021). Relaying this information within the team’s structure such that aggregation occurs in the correct crew positions requires some form of team communication to occur (Canan & Demir, 2021).

Balancing loops one (B1) links the performance of the team to its ability meet the demands of their situation and environment. As novel conditions present themselves, successful teams adapt their resource usage and allocation behaviors alongside communication and coordination strategies (van den Oever & Schraagen, 2021). In the case of the USS Vincennes’ combat information center, the team failed to avoid saturation and overallocation of their communications resources (Dotterway, 1992). Task overwhelm in critical crew positions caused the team’s practiced communications dynamics to fail because overloaded personnel failed to ensure that the team followed its communications standard operating procedure, and maintained radio discipline (Dotterway, 1992).

Documented in balancing loop two (B2), Figure 19 illustrates that a machine teammate may regulate the team's communications by ensuring the adherence to team standard operating procedures. Adhering to the team's standard operating procedures may improve radio network capacity by ensuring that transmitting stations conduct themselves on the appropriate radio nets. Capacity may also be improved by reducing the amount of ambiguous information interjected into the medium. Members of the air team issued track information without identifying sender, intended recipient, track identification number, and other information essential to coordination and the construction of situational awareness (Dotterway, 1992).

Regulating communications, a machine partner may listen to the information pushed onto the medium and validate that its contents are aligned with the intended purpose of the channel. By further analyzing the contents of the message, the machine partner can help to regulate communications by ensuring that communications are purposeful and non-ambiguous. If this capability is outside the bounds of current technology, the machine may still ensure that predetermined communications patterns are adhered to. For example, as the team within the combat information center was increasingly focused on TN 4131, Captain Rogers asked for TN 4474's kinematics information (Dotterway, 1992). This request for information coincided with a breakdown in the team's shared mental model of the environment, and occurred when critical crew positions were overwhelmed (Dotterway, 1992). Recognizing that the captain asked a question, the machine partner could have informed the crew of TN 4474's kinematics, and based on recent message history, it may have additionally issued kinematic information for the team's track of interest, TN 4131.

When members of the team omit certain parts of messages that do not adhere to standard operating procedures, the machine partner may notify the teammate that they may not be communicating effectively. Another useful application of the machine partner could also be to notify the team that they are not communicating at the predetermined rate. During the minutes leading up to the engagement of Iran Air flight 655, the USS Vincennes combat information center crew failed to adhere to communications standard operating procedures (Dotterway, 1992). Frequently, messages were pushed into the medium without identifying

the sender, receiver, track number, or other relevant information needed to build context (Dotterway, 1992). These actions overloaded the communications medium and did not contribute to situational awareness (Dotterway, 1992). By analyzing the messages that are being sent into the medium, the machine partner could inform team members that their communications are ambiguous, and that additional information is needed to ensure clarity. Improving the effectiveness of communication could also be done by reminding team members that there are communications that need their attention, or that they have not replied to a teammate awaiting their attention. With these practices, a machine partner could treat team communications as an optimization problem, and enable the team to work towards an increased understanding of the operational environment (Fout & Ploski, 2018).

When crew members within the combat information center misidentify IFF modes and codes for tracks, fail to move radar read gates, and issue incorrect kinematics data, the machine partner may be used to correct these issues. Because the machine partner can make sense of utterances as well as text-based communications, it could improve the analyses of the team by validating communications against Aegis records and then explicate differences. Supporting the flow of correct information to the team would improve performance as well as ensure that the team was able to understand the status of their goals, tasks, team, and external environment (Fout & Ploski, 2018; Chappell, 2020; NASEM, 2021). Correcting the information as it flows between crew positions can ensure that the right actions occur, preventing the crew from wasting precious time.

In the case of channel or medium saturation, the machine teammate may perform automatic link establishment functions. With this function, it may optimize communications and protect against communication capacity challenges by identifying available frequencies and moving communications to those openings. This may enable teams to be able to dynamically adjust their communications patterns when facing novel events, and could avoid medium saturation challenges that occur as teams attempt to organize to meet new challenges (Gorman et al., 2020).

NASEM (2021) noted that teams must coordinate to aggregate their environmental and situational information; this includes the practice of exchanging information to align their goals. Given that teams use interdependent roles to accomplish tasks, such

coordination is essential to a team’s performance (NASEM, 2021). Underpinning coordination mechanisms are the team’s communications behaviors (NASEM, 2021). Communications enable the team to synchronize actions, construct an understanding of their environment, and develop an accurate shared mental model (Mohammed et al., 2010; Canan & Demir, 2021; NASEM, 2021).

e. Interactions Between USS Vincennes and the Environment

(1) Modeling the “As-Is” System

Clearly there were many challenges with the teaming behaviors within the USS Vincennes combat information center. The question of what led a billion dollar warship to engage an airliner was explained—the answer is simple, self-defense. That answer is insufficient. A better question is how did an expert crew controlling a one billion dollar warship purpose-built for anti-air warfare arrive at the conclusion that self-defense was necessary? By themselves, poor team behaviors fail to explicate “the how” in this case. The answer lies in the way that the crew of the USS Vincennes interacted with their environment and the other actors within the operational area. Dotterway (1992) presented a minute by minute guide of crew interactions, which allowed the authors to describe how the USS Vincennes conducted an escalating series of actions that resulted in the destruction of a civilian airliner. This escalating pattern of behavior will be depicted in Figure 20.

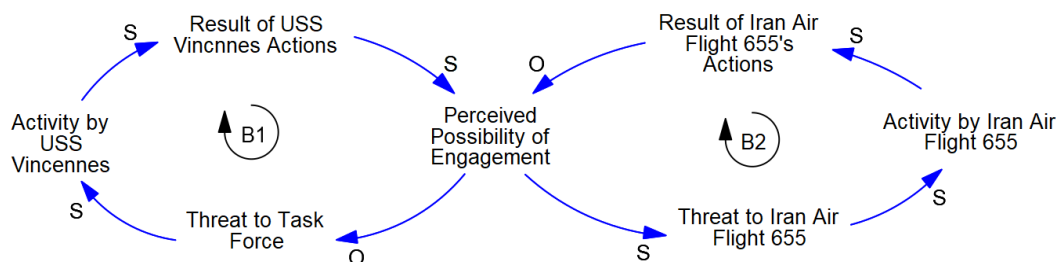


Figure 20. Interactions Between the USS Vincennes and Iran Air Flight 655

This causal loop diagram is based upon the escalation archetype (Kim, 2000). Within this archetype, one party takes actions that are perceived as a threat by a second party (Kim, 2000). The second party continues their actions, or acts threateningly to the

first party, continuing an insidious cycle of action and counter-action (Kim, 2000). Essentially, two balancing loops are clashing to bring the system into its desired state.

Already engaged in a sustained surface battle, the combat information center crew of the USS Vincennes observed TN 4474 depart from Bandar Abbas (Dotterway, 1992). Given the presence of an Iranian P-3 maritime surveillance aircraft, the crew believed that they were witnessing Iranian team backup behaviors, manifested as offensive air support (Dotterway, 1992). This belief was reinforced by recently received intelligence reports. Investigations into the incident revealed that a volume of intelligence reporting found its way into the hands of the U.S. Navy crews operating in the Persian Gulf ahead of the 4th of July weekend. Intelligence reports warned of an increased possibility of hostile action, kamikaze aircraft, as well as information that Iranian aircraft would not fly within their established patterns (Dotterway, 1992). Reports additionally indicated that anti-surface capable F-14s transferred to Bandar Abbas (Dotterway, 1992). These items were incorporated into the judgement of combat information center crew members, who characterized TN 4474 as an assumed enemy track (Dotterway, 1992). When the incorrect use of the Aegis system injected Mode II IFF information onto the identification supervisor's screen, he then conveyed this information onto radio nets 15 and 16 (Dotterway, 1992). Thus, the perceived actions of civilian and Iranian actors were characterized as threatening actions.

These information inputs sparked a flurry of action within the USS Vincennes combat information center. At 1017L, the flight would be classified as an Iranian F-14, leading others within the combat information center to scrutinize the track (Dotterway, 1992). By 1018L, the identification supervisor was searching the commercial air schedule to ensure that the track was not a scheduled commercial flight, and the USS Sides illuminated the track with fire control radar (Dotterway, 1992). During this time, Iran Air flight 655 maneuvered into its air corridor and began a long, slow ascent directly in line with the USS Vincennes (Dotterway, 1992). By 1019L, the crew of the USS Vincennes was challenging TN 4131 over the military air distress nets, asking the "unidentified Iranian aircraft on course 203, speed 303, altitude 4000... to state its intentions" (Dotterway, 1992, p. 38). This request for information would go unanswered because the

civilian air traffic control and Iran Air flight 655 were not monitoring military air distress communications nets (Dotterway, 1992).

By 1020L, the USS Vincennes was broadcasting warnings onto the international air distress nets (Dotterway, 1992). Unfortunately, the communications were ambiguous and did not refer to the aircraft based on information that a civilian aircrew could use to identify themselves (Dotterway, 1992). Thus, Iran Air flight 655 maintained a steady course, likely believing the transmission to be intended for a different recipient (Dotterway, 1992). During the remainder of this minute, the identification supervisor reported Mode II IFF on internal radio nets 15 and 16, and TN 4131 would be categorized as an F-14 on the combat information center's large screen display (Dotterway, 1992). By 1021L, the anti-air warfare TAO was relaying information updates to headquarters, and informing the commander, JTF Middle East of the USS Vincennes' intention to fire on the target at 20 nautical miles (Dotterway, 1992). This request prompted higher headquarters to require the USS Vincennes to issue additional warnings on both the international and military air distress nets (Dotterway, 1992).

The added communications task caused the anti-air warfare coordinator to direct continuous warnings across both nets (Dotterway, 1992). Interestingly, during this minute, the combat information center officer informed Captain Rogers that TN 4131 was "possibly a commercial airliner," (Dotterway, 1992, p. 40). It is also worth noting that the Bandar Abbas air traffic control tower did not relay any of the warnings conveyed by the USS Vincennes (Dotterway, 1992). By 1022L, the members of the USS Vincennes combat information center perceived TN 4131 as descending as well as accelerating (Dotterway, 1992). Unaware that the track designation changed over the previous minutes, and demonstrating the collapse of a shared mental model, the captain of the ship inquired into the status of TN 4474 (Dotterway, 1992). As a result, the anti-air warfare coordinator began efforts to illuminate TN 4474 with their fire control radar (Dotterway, 1992).

Interestingly, this interpretation of the TN 4474's kinematics data conflicted with that of the USS Sides. The TAO aboard the USS Sides interpreted the track as ascending, while critical positions within the USS Sides combat information center began shouting within the compartment that the USS Vincennes was targeting a commercial aircraft

(Dotterway, 1992). By 1022L, the commanding officer of the USS Sides declared that TN 4131 was not a threat to his ship and directed his team's attention to the hostile Iranian P-3 oriented to their west (Dotterway, 1992).

Iran Air flight 655 reached the 20 nautical mile decision point by 1023L. Coincidentally, this is when a U.S. Navy A-6, tracked as TN 4474 began its descent into a surface combat air patrol (Dotterway, 1992). At this decision point, the crew of the combat information center aboard the USS Vincennes searched for radio-frequency and other electromagnetic spectrum indicators from TN 4474 indicating that they were being targeted for attack, but did not detect any signals (Dotterway, 1992). The crew continued to issue continuous warnings, and the tactical information coordinator began to update all stations on internal nets 15 and 16 during every instance where the net was open (Dotterway, 1992). TN 4474 continued to descend and accelerate (Dotterway, 1992). This was problematic for the team in the USS Vincennes combat information because they had confused the attack profile kinematics of TN 4474 with the flightpath of track TN 4131, and observed TN 4131 closing on their position (Dotterway, 1992).

Acting in self-defense, the commanding officer turned the firing key at 1024L, and 14 seconds later the missile systems began to initiate a launch sequence, ultimately firing a pair of SM-2 anti-aircraft missiles into the sky at TN 4131 (Dotterway, 1992). Interestingly, the international air distress talker was in the process of issuing another warning when the missiles were launched from the ship (Dotterway, 1992). Iran Air flight 655 was destroyed approximately 8 nautical miles from the USS Vincennes at an altitude of 13,500 feet and an airspeed of 383 knots (Dotterway, 1992).

As the authors explored the two balancing loops, it became clear that the relatively low altitude and heading of Iran Air flight 655 triggered the USS Vincennes' sense of danger. Placing this in context with observed adversary tactics, techniques, and procedures, it is understandable why the crew could have believed TN 4474 or TN 4131 were hostile actors. Because Iranian military aircraft were observed emitting both Mode II and Mode III IFF signals, the ship was already engaged in combat, as well as the fact that there were Iranian maritime surveillance aircraft in the skies, it could be argued that this was a reasonable assumption (Dotterway, 1992). This assumption was continually reinforced

when the crew confused TN 4474's kinematics with those of TN 4131, and they believed the track to be descending and accelerating—a common profile for surface attack aircraft beginning an attack run.

As the warnings went unanswered, and Iran Air flight 655 continued its flight toward the USS Vincennes, the crew perceived themselves to be in danger at an increasing rate. This triggered still more activity, which did not change the behavior of the track, leading the crew to believe that hostilities were going to increase against them. To avoid suffering the same fate as the USS Stark, the USS Vincennes escalated until the system reached a desired state—the hostile track was removed from the sky, and the ship was safe from airborne threats.

From the perspective of Iran Air flight 655, it had no idea that the actions it was taking were leading to a deteriorating spiral of escalation. It continued to adhere to its flight plan, emitted the appropriate IFF signals, and followed the flight path issued by Bandar Abbas tower. They did not return any communications to the USS Vincennes, due to two factors. The aircraft did not monitor the military air distress net, and those transmissions on the international air distress net were ambiguous enough that the aircrew likely believed those transmissions were not intended for them (Dotterway, 1992). Essentially the aircraft flew straight into hostilities, unaware that its actions were threatening or would trigger an engagement.

(2) Modeling the “To-Be” System

When attempting to diagnose the underlying issues present within an escalation archetype, akin to Figure 20, Kim (2000) asserted that it is necessary to identify how one agent is forced to compete against the other. Kim (2000) also noted that it is important to model the significant delays in the system, as they may misrepresent the nature of perceived threats. Within the “as-is” system depicted in Figure 20, the combat information center crew was not aware that they were targeting an airliner. This lack of situational awareness was attributed to the time pressures facing the crew (C-Span, 1992). Similarly, Iran Air flight 655 was not aware that it was perceived as hostile, and therefore did not adjust its actions (Dotterway, 1992).

Although the time pressures felt by the crew of the USS Vincennes were notable, they were more pronounced because of a combination of other factors previously discussed. Figure 12 depicted how U.S. Navy ships were essentially ordered to avoid taking the first hits in an engagement. Figure 14 demonstrated how time pressures eroded the combat information center crew's productivity and mental capacity. Figure 16 discussed how the ad-hoc team structure and novel task environment prevented the crew from effectively responding to their environment. Figure 18 demonstrated how the team was unable to effectively communicate, coordinate, or interact because of challenges with the team's dynamics. When attempting to diagnose the problems facing the crew of the USS Vincennes, time pressures again emerged as a root cause for team's situational awareness deficiencies.

Based upon the limits to success archetype, Figure 20, depicts that the efforts of the team to improve their situational awareness and perform the correct series of actions were limited by time. From takeoff to shootdown, Iran Air flight 655 was transiting an air corridor towards the USS Vincennes, meaning that there was a fixed period of time before the aircraft reached the 20 nautical mile decision point for engagement (Dotterway, 1992). Correcting the issue of time pressures requires a way to combat the effects of the constraint (Kim, 2000). Unfortunately, time pressures are ever-present in situational awareness and decision-making scenarios such as these (Fout & Ploski, 2018; Kennedy, 2021). Kennedy (2021) noted that efforts to gather an adequate amount of situational awareness must always contend with time pressures that bound sense-making activities. Thus, one way to mitigate effects of time constraints is to ensure that crew members have accurate information from which to drive interactions and coordination mechanisms. This practice is depicted in Figure 21.

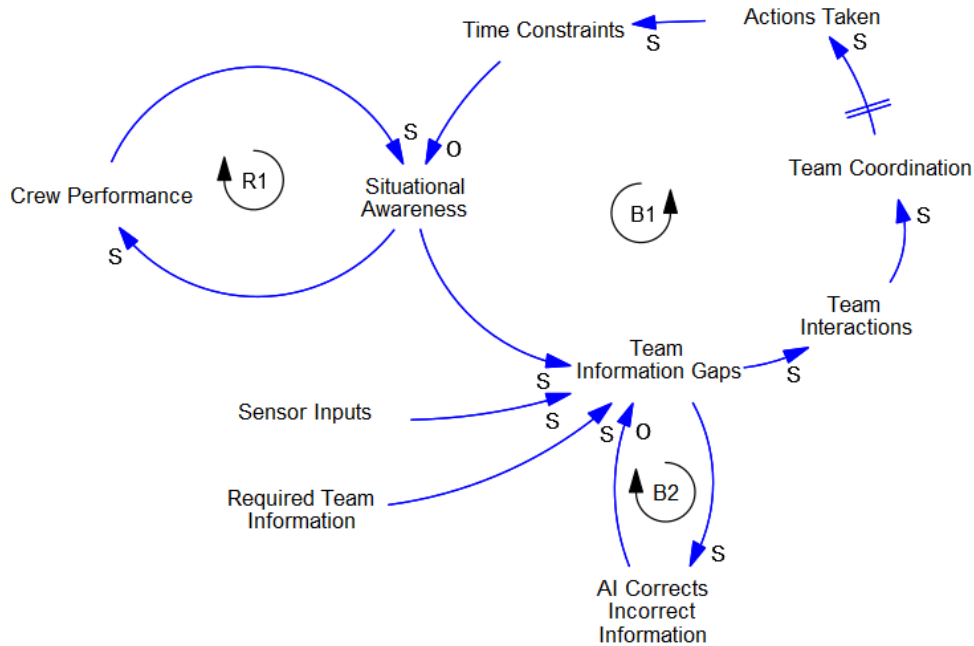


Figure 21. AI Enables Higher Levels of Situational Awareness

Reinforcing loop one (R1) demonstrates that the crew’s performance is tied to its situational awareness, with higher levels of team situational awareness implying improved performance (Cooke et al., 2004). Situational awareness is tied to balancing loop one (B1), where information gaps spark interaction and coordination mechanisms to address these gaps, resulting in actions taken (Kennedy, 2021). All of the events occurring in balancing loop two take time to execute, which ultimately influences the level of situational awareness that the team may achieve (Kennedy, 2021).

Balancing loop two (B2) illustrates how information gaps may be mitigated through the use of a machine teammate. Given an understanding of the team’s information requirements from sources such as playbooks or through learning opportunities presented during training exercises, the machine partner may use collected data from sensor inputs to implicitly “fill in” information gaps. Such action is akin to forecasting the requirements and actions of teammates so that the team may quickly alter interaction patterns to contend with their current situation (Gorman et al., 2020).

Team members may share a mental model of their situation, yet be ultimately incorrect about the conditions of their environment (Lim & Klein, 2006). Such mismatches

may result in choosing strategies and executing tasks that are incompatible with the environment (Lim & Klein, 2006). Interestingly, given similar situational awareness inputs, the commanding officer of the USS Sides did not classify TN 4131 as a threat (Dotterway, 1992). In fact, crew members within the USS Sides' combat information center understood TN 4131 to be Iran Air flight 655 (Dotterway, 1992). It is worth noting that the USS Sides was an older ship and did not possess capabilities on par with Aegis, however this still presents an excellent opportunity for the machine teammate. As the USS Sides was a vital component of the JTF Middle East air defense team, it may be of interest to the commander of the USS Vincennes that the USS Sides' leadership arrived at a different interpretation and conclusion.

Leveraging the concept of cognitive consensus, the machine partner may be able to notify decision makers when teams arrive at different interpretations of information, so that mental models and fundamental assumptions may be evaluated (Mohammed & Dumville, 2001). Although one could certainly argue that the time constraints facing the USS Vincennes did not support a lengthy discourse about the fundamental assumptions and subsequent assessments of the USS Sides' captain, there may be circumstances that support the development and validation of team mental models. As required, the machine partner may provide teams with the estimates of other team members to facilitate mental model convergence. Mental model convergence may ultimately improve team processes such as back-up behaviors, communication dynamics, as well as strategy selection and implementation (Mohammed et al., 2010).

Akin to the notion of "plays" within a "playbook," that were discussed in Section C, the machine partner could present recommendations to the human crew, and explain the information used in its analyses accompanied by a confidence interval (Chappell, 2020; NASEM, 2021). In addition to offering a variety of potential responses to facilitate coordination, such information could improve the team's situational awareness of their operational environment as well as friendly and enemy forces (Fout & Ploski, 2018; NASEM, 2021). The addition of this explanatory information may serve to provide the human partner with important information that may improve the team's ability to build an understanding of their environment and therefore successfully adapt to novel

circumstances (Gorman et al., 2010; Chappell, 2020). Use of a playbook may also improve the human partners' ability to observe, predict, and direct the machine partner's actions, enabling the machine partner to be better integrated into the team (Fout & Ploski, 2018).

Within teams of humans, it is acknowledged that gathering information to develop situational awareness and support decision making is limited by the amount of time available to decision makers (Kennedy, 2021). Teams must possess accurate information to ensure that they develop a reasonably correct model of their environment, structure, and tasks from which they will direct and control actions (Chappell, 2020). Situational awareness is essential to teams as they interact with their environment, especially if perturbations or other impediments occur (Cooke et al., 2004). Ultimately, higher levels of team situational awareness have been demonstrated to be effective predictors of a team's performance (NASEM, 2021).

D. CHAPTER SUMMARY

Using two well-studied vignettes, this chapter operationalized some of the attributed behaviors and capabilities that would enable a machine partner to improve the performance of teams operating in military contexts. Specifically, the conceptual models described how artificial intelligence may improve team performance as well as team communication, coordination, and interaction dynamics. Also modeled was the ability of a machine partner to support team backup behaviors, improve team situational awareness, and assist in the progress towards team goals. As a part of these vignettes, the machine partner was assigned specific roles, responsibilities, and was modeled as performing interdependently with the remainder of the team. These factors illustrate that an artificial intelligence may be a member of a team (NASEM, 2021).

IX. SUMMARY AND FUTURE WORK

A. CONCLUSIONS

Human machine teaming may offer a novel means to outcompete great power competitors in future conflict (Tangredi & Galdorisi, 2021). Recent research indicates that the relative strengths and weaknesses of humans and machines are complementary, and placing them together in team constructs may produce superior results compared to deploying humans or machines in isolation (Stumborg et al., 2019). Based upon these assertions, the authors studied how humans and machines may partner in team constructs.

Initial review of the DON's AI and HMT strategies reveals that to fully understand the deeper potential of AI as a component of HMT, an exploration of pertinent topics within several literature fields was necessary. Although teaming theory in humans has been well researched, very little research explores military C2 concepts and associated doctrine. This study applies these aspects to HMT and AI theories while accounting for the capabilities and limitations of AI technologies. This lens allowed for the creation of a baseline to develop HMT components that may optimize the assets available to a commander.

This study introduced a conceptual systems dynamical model to synthesize the desired behaviors of AI in human machine team constructs. Modeling was used to make the various interdependencies of a complex system, such as a team of military personnel, more readily apparent. Dotterway (1992) noted that it becomes possible to understand how even subtle changes propagate through all elements in a system. Two vignettes served as the foundation for the models, and provided context for how HMT may improve upon the well-studied teaming challenges present in Snook's (2011) study of Operation Provide Comfort and Dotterway's (1992) examination of the USS Vincennes incident.

The models in which these vignettes are explored are built around essential team behaviors such as the ability to communicate, coordinate, and interact towards the completion of a goal. Additionally, the models were designed to capture the ability of a machine teammate to participate in the construction of team situational awareness, the development of team mental models, as well as the ability to provide support, such as team

backup behaviors. Based upon the current capabilities of narrow AI, the authors assert that machines can partner with humans in team constructs, and that such integrations may improve the team's agility. Finally, the authors modelled how machine partners may be more observable, predictable, and directable through the use of interactive team cognition-based approaches (Fout & Ploski, 2018).

1. Research Question 1: What does it mean to be a teammate?

A team is a group of members with unique, specialized roles and responsibilities that works interdependently to accomplish a specific goal (NASEM, 2021). As a part of the inputs, processes, and outputs performed by a team, a mental model of the environment, team structure, allocation of tasks, and progress towards the team's goal are modelled by its participants (NASEM, 2021). Information to populate such models is collected through team interactions. These interaction patterns emerge as teams communicate with each other and coordinate actions to complete assigned tasks. Thus, a team is a group possessing interaction patterns, a common understanding of its tasks as well as the team-related factors utilized to achieve its goals (Cannon-Bowers & Salas, 2001; Mathieu et al., 2010; Canan & Demir, 2021; NASEM, 2021).

2. Research Question 2: How can machines be partners with humans in teams?

The authors assert that machines may partner with humans in team constructs. The "to-be" models developed in this study demonstrate how a machine partner may perform specific roles on the team that improve team performance. More specifically, they demonstrated how a machine teammate may perform tasks that regulate team communications, coordination, and interaction behaviors, ultimately improving team situational awareness, agility, and performance.

From an interactive team cognition perspective and relying on a machine partners ability to perform listening functions, the machine demonstrated the ability to improve a team's communications processes. Leveraging the strengths of the machine partner to mitigate ambiguous and ineffective communications may improve a team's ability to coordinate when presented with novel circumstances (Gorman et al., 2020). This study also

demonstrates how the regulation of team's communications patterns may improve a team's agility and performance.

Monitoring the structure of a team, identifying gapped roles, and facilitating task allocation may help teams continue to perform if their structures are somehow disrupted. For example, a machine partner may be able to monitor team structures in ways that improve agility and performance. Modifying the structures of the team may support adaptive teaming; an essential behavior for military teams that are almost certainly guaranteed to encounter novel conditions on the battlefield (HQ USMC, 1997; Gorman et al., 2010; Bray & Moore, 2021).

Interaction dynamics are essential to the agility of a team. Teams primarily navigate around obstacles and perturbations through patterns of interaction, which gives them the information and situational awareness needed to adjust to unanticipated conditions (Gorman et al., 2020; Canan & Demir, 2021; NASEM, 2021). This study demonstrated how machine partners may improve inter and intra team interaction dynamics.

Adjusting to unforeseen tasks and conditions in the environment may require team backup behaviors. The machine partner may support human partners in a limited capacity by acting to fill in for certain tasks when humans are removed from a team's structure. A machine partner may execute task sharing behaviors for certain tasks, ultimately improving a team's ability to handle increased task loads. By performing task monitoring and allocation, humans and machine partners may be able to identify when task overload conditions exist. Identifying conditions of overload may then trigger backup responses to prevent lapses in performance (Marks et al., 2002; Smith-Jentsch et al., 2009).

This study shows the integration of a machine partner with sensors may enable a flow of validated sensor information to be presented to the team. By conjoining the machine partner with sensor collection capabilities and the team's information networks, the machine partner may scrutinize team communications for mismatches between sensor data and team communications. In addition to assisting humans in identifying incorrect observations, such actions may enable the machine partner to provide needed information updates to the team to support the construction of situational awareness and accurate mental

models. Possessing accurate information about environmental cues, activities, and goals supports the development of situational awareness (Kennedy, 2021). Situational awareness is an important factor in decision making as well as performance (Cooke et al., 2004; Kennedy, 2021)

3. Research Question 3: How can the integration of machine partners into previously human-human teams improve the team's agility?

Team agility is a critical component to organizational development and transformation based on the iterative nature of change (Song et al., 2022). Based on AI behaviors that were modeled in chapter eight and further explicated in research question two, the authors noted that AI may improve team agility by driving or otherwise supporting adaptive team behaviors. These behaviors manifest as responses to perturbations in established teamwork and taskwork patterns. The “to-be” models in chapter eight demonstrate that AI integration improves team effectiveness and behavioral elasticity. With a machine teammate, routine or otherwise information-heavy tasks may be delegated to a machine partner that is more capable of quickly processing large amounts of data more precisely than a human partner (Stumborg et al., 2019). HMT provides unique opportunities in solving team problems, and as technology progresses the integration of machine partners into human teams may also improve team communication and coordination. The end state will result in better performance of teams and more agile adaptations to team disruptions (Song et al., 2022). In sum, team agility may be improved through the integration of machine partners.

4. Research Question 4: Can machine partners be better understood with shared team mental models or interactive team cognition approaches?

Currently, humans hold varied feelings towards AI teammates but do maintain a positive outlook towards future teamwork (Grosz et al., 2016; Russell, 2019; Zhang et al., 2021). Understanding a machine partner will take a willingness to team with a machine, shared understanding between AI and humans, clear communication abilities, and proven performance (Bansal et al., 2019; NASEM, 2021; Zhang et al., 2021). Trust is one of the many factors that will ultimately determine the success of HMT (Russell, 2019; NASEM,

2021). The models used in this study demonstrate design ideas that may give rise to team subsequent team interaction models for HMTs. These models do not inflate or annihilate trust in human-machine relationships. Instead, the suggested design ideas described in this study's models support interactive team cognition. With further study, team cognition and shared mental models may help mature HMT processes.

It is worth noting that team and shared mental models are an important aspect of human teams, and as identified during this study, will be an important factor in HMT (Mohammed et al., 2010; NASEM, 2021). NASEM (2021) identifies that developing mental models with a machine partner represented a significant challenge at the time of their study. Given the evolving nature of certain machine systems, such as those that employ machine learning algorithms, it may become difficult for human partners to construct an accurate mental model given systems' flexible error boundaries and difficulties in explicating conclusions (Bansal et al., 2019; NASEM, 2021).

Shared and team mental models reside within the umbrella of team cognition (Mohammed et al., 2010). During the course of this study, it became apparent that an interactive team cognition approach, which accounts for shared mental models, may prove more beneficial to exploring how a machine may support human partners in team structures. While operating within an HMT construct, humans and machines must coordinate and interact to achieve goals (NASEM, 2021). These interaction mechanisms between teammates exist to move information about interdependent tasks to those who need it. The elements that underpin these interaction mechanisms are highly inter-related and may be performed by AI systems (NASEM, 2021). As demonstrated herein, the interactive team cognition framework appears to offer a promising lens for future HMT research.

B. ADDITIONAL FINDINGS AND LIMITATIONS

1. Additional Findings

Despite the evidence that human attributes have a considerable impact on machine performance, the human elements of human-machine teams have not been modelled with the same level of fidelity as machine partners (Wang & Kosowski, 2019). During the

course of this study, the authors also discovered machine partners in some studies were not true AI. For example, in Chakraborti and Kambhampati's (2018) analysis, both human and machine were modeled with the aid of Amazon's Mechanical Turk. Similarly, Chan, Doyle, McElfresh, Conitzer, Dickerson, Borg, and Sinnott-Armstrong (2020) used Amazon's Mechanical Turk to study trust dynamics between humans and artificial AI given certain conditions. In Chan et al.'s (2020) study, both humans and AI were represented by the Mechanical Turk. While these studies are exceptional, future research into AI and HMT may require the use of established human teams and authentic AI systems to suitably evaluate HMTs.

2. Limitations

The most notable limitation to this study is that there are few, if any, prior cases of HMT to study. Additionally, current market-comparable AI systems require narrow use cases that do not generalize to military environments. Battlefields are infamous for their complexity and will require capabilities far beyond those represented by current narrow AI systems to contend with the ambiguities and challenges of warfare.

Additionally, the research process for this study was limited to unclassified materials contained within the public record. The primary sources driving the analyses included the comprehensive Dotterway (1992) thesis on the USS Vincennes incident, as well as Snook's (2011) definitive work on the Operation Provide Comfort incident. These cases were selected because the process failures and outcomes of both incidents have been extensively studied. The available information enabled the authors to develop conceptual models of team process failures, diagnose the root causes of these failures, and identify fundamental fix actions. As such, these incidents represented militarily useful vignettes for modeling how team dynamics may be improved with the addition of a machine teammate.

C. RECOMMENDATIONS FOR FUTURE WORK

It may be pertinent to examine the effects of multiple AIs filling roles within a team. Given the limitations of current AI and its narrow focus, multiple systems may be required to support teams.

It is recommended that future research explore how machine partners can learn the concept of commander's intent and take actions within these bounds. The use of selected "plays" from a well-rehearsed and trained "playbook" may be linked to the concept of commander's intent, and thus may be a starting point for future exploration of the concept. Leveraging a machine partner within the maneuver warfare construct implies that some understanding of commander's intent may be required (HQ USMC, 1997)

Humans and animals have been successfully working together in team constructs for countless years. Akin to AI, animals are not given equal status to humans within teams, and are believed to be limited in their ability to communicate, coordinate, and interact with human partners as a team member (NASEM, 2021). Given that human-animal teaming has been successfully applied in military contexts, such as those demonstrated by military working dogs, it may prove beneficial to HMT to explore human-animal teaming. Human-animal teaming may offer additional insights into creating behavioral cues and signals that could be utilized by an AI working alongside human partners (NASEM, 2021).

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- Ahmed, A., Agarwal, S., Kurniawan, I., Anantadjaya, S., & Krishnan, C. (2022). Business boosting through sentiment analysis using Artificial Intelligence approach. *International Journal of System Assurance Engineering and Management*, 13(1), 699–709.
- Allard, C. (1996). *Command, control, and the common defense*. National Defense University.
- Allspaw, J. (2015). *Trade-offs under pressure: Heuristics and observations of teams resolving internet service outages*. [Master's thesis, Lund University].
https://www.academia.edu/attachments/41309101/download_file?st=MTY1MDA0NDc4NSw3MS4yMDluMTA3LjE1&s=swp-splash-paper-cover
- Andrus, J. (1994). *AFR 110–114 aircraft accident investigation board report of investigation: U.S. Army Black Hawk helicopters 87-26000 and 88-26060*. U.S. Air Force.
- Apple Jr., R. (1988, 5 July). The downing of flight 655. *New York Times*.
<https://www.nytimes.com/1988/07/05/world/the-downing-of-flight-655-military-errors-the-snafu-as-history.html>
- Argyris, C. (1973). The tendency is to analyze and study phenomena in their full-blown wholeness, complexity. *Readings in the Study of Leadership*, 379.
- Aycock, A., & Glenney IV, W. (n.d.). Trying to put Mahan in a box: Insights from attempting to develop a decision aid for the operational commander. In S.J. Tangredi and G. Galdorisi (Eds.), *AI at war: How big data, artificial intelligence, and machine learning are changing naval warfare* (1st ed., pp. 265–286). Naval Institute Press.
- Bansal, G., Nushi, B., Kamar, E., Lasecki, W., Weld, D., & Horvitz, E. (2019). Beyond mental models: The role of mental models in human-AI team performance. *The Seventh AAAI Conference on Human Computation and Crowdsourcing (HCOMP-19)*, 2–11.
- Beck, R., Dibbern, J., & Wiener, M. (2022). A Multi-Perspective Framework for Research on (Sustainable) Autonomous Systems. *Business & Information Systems Engineering*, 1–9.

- Bergin, R. (2021). *Information systems strategy and policy: Course recap* [Class notes for IS4182: Information Systems Strategy and Policy]. Department of Information Sciences, Naval Postgraduate School. https://cle.nps.edu/access/lessonbuilder/item/170695/group/0a38d2fb-13ff-44e7-ae98-db091dcdd3f/Lessons/_Week%2010_%20The%20B_l_ysis%20%20_%20Recap/IS_4182_Fall%2018%20Course%20Recap.pptx
- Birkinshaw, J., & Ridderstråle, J. (2010). Adhocracy for an agile age. *Organization Science*, 22(5), 1286–1296.
- Black, N. (1987, July 27). Two officers accept blame in frigate attack, will leave service. *Associated Press*. <https://apnews.com/article/4dd26be262aab7b7f9f5404dc46d1111>
- Bode, I., & Watts, T. (2021). Meaningless human control: Lessons from air defence systems on meaningful human control for the debate on AWS. Centre for War Studies & Drone Wars UK.
- Borges et al. (2021). Decision-Making Framework for Implementing Safer Human-Robot Collaboration Workstations: System Dynamics Modeling. *Safety* (7)75. <https://doi.org/10.3390/safety7040075>
- Bowron, E. (1998). *Operation Provide Comfort review of U.S. Air Force investigation of Black Hawk fratricide incident*. General Accounting Office.
- Bray, W., & Moore, D. (2021). The Department of the Navy's commitment to big data, AI, and machine learning. In S. J. Tangredi and G. Galdorisi (Eds.), *AI at war: How big data, artificial intelligence, and machine learning are changing naval warfare* (1st ed., pp 62–80). Naval Institute Press.
- Brannick, M., Prince, A., Prince, C., & Salas, E. (1995). The measurement of team process. *Human Factors*, 37(3), 641–651.
- Brose, C. (2020). *The kill chain: Defending America in the future of high-tech warfare*. Hachette Books.
- Builder, C. H., Bankes, S. C., & Nordin, R. (1999). *Command concepts: A theory derived from the practice of command and control*. Rand.
- Burke, M. (2022, January 21). Tesla driver charged with manslaughter in deadly autopilot crash raises new legal questions about automated driving tech. *NBC News*. <https://www.nbcnews.com/news/us-news/tesla-driver-charged-manslaughter-deadly-autopilot-crash-raises-new-le-rcna12987>
- C-Span. (1987, May 18). *USS Stark* [Video]. <https://www.c-span.org/video/?172074-1/uss-stark>

- C-Span. (1992, July, 21). *USS Vincennes attack on an Iranian airliner* [Video].
<https://www.c-span.org/video/?27276-1/uss-vincennes-attack-iranian-airliner>
- Canan, M., & Demir, M. (2021). Addressing two central issues of team interaction dynamics: The whole is greater than the sum of its parts. *Proceedings of the AHFE 2021 Virtual Conferences on Neuroergonomics and Cognitive Engineering, Industrial Cognitive Ergonomics and Engineering Psychology, and Cognitive Computing and Internet of Things*, 61–69. https://doi.org/10.1007/978-3-030-80285-1_8
- Cannon-Bowers, J., & Salas, E. (2001). Reflections on shared cognition. *Journal of Organizational Behavior*, 22, 195–202.
- Cella, C. E. (2017). Machine listening intelligence. *arXiv preprint arXiv:1706.09557*.
<https://doi.org/10.48550/arXiv.1706.09557>
- Center for Naval Analyses [CNA]. (2021a, November 8). *AI and autonomy in Russia* (Report No. DOP-2021-U-031157-Final). CNA.
- Center for Naval Analyses [CNA]. (2021b, November 9). *The China AI and autonomy report* (Report No. DIM-2021-U-030977-Final2). CNA. <https://www.cna.org/centers/cna/cip/china/china-ai-newsletter>
- Chakraborti, T., & Kambhampati, S. (2018). Algorithms for the greater good! on mental modeling and acceptable symbiosis in human-ai collaboration. *arXiv:1801.09854*. <https://doi.org/10.48550/arXiv.1801.09854>
- Chan, L., Doyle, K., McElfresh, D., Conitzer, V., Dickerson, J., Borg, J., & Sinnott-Armstrong, W. (2020). Artificial intelligence: Measuring influence of AI ‘assessments’ on moral decision-making. *2020 Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 214–220. <https://doi.org/10.1145/3375627.3375870>
- Chappell, J. (2020). *Command and control with human-machine teams* [Master’s thesis, Naval Postgraduate School]. NPS Archive: Calhoun. <http://hdl.handle.net/10945/66057>
- Choi, C. (2021, September 21). 7 revealing ways AIs fail: Neural networks can be disastrously brittle, forgetful, and surprisingly bad at math. *IEEE Spectrum*.
<https://spectrum.ieee.org/ai-failures>
- Christensen, H., Amato, N., Yanco, H., Mataric, M., Choset, H., Drobnis, A. ... Trinkle, J. (2020). A roadmap for U.S. robotics—From internet to robotics, 2020 edition. *Foundations and Trends in Robotics*, 8(4), 307–424.

- Clarke, A. & Knudson III, D. (2018). Examination of cognitive load in the human-machine teaming context [Master's thesis, Naval Postgraduate School]. NPS Archive: Calhoun. <http://hdl.handle.net/10945/59638>
- Coakley, T. (1992). *Command and control for war and peace*. National Defense University Press: U.S. G.P.O.
- Cooke, N., Gorman, J., Myers, C., & Duran, J. (2013). Interactive Team Cognition. *Cognitive Science*, 37, 255–285.
- Cooke, N., Salas, E., Kiekel, P. & Bell, B. (2004). Advances in measuring team cognition. In E. Salas and S. Fiore (eds.), *Team Cognition: Understanding the factors that drive and performance* (1st ed., pp. 3–8). American Psychological Association.
- Costa, A. (2003). Work team trust and effectiveness. *Personnel Review*, 32(5), 605–622. <https://doi.org/10.1108/00483480310488360>
- Damacharla, P., Javaid, A. Y., & Devabhaktuni, V. K. (2018, July). Human error prediction using eye tracking to improvise team cohesion in human-machine teams. In *International Conference on Applied Human Factors and Ergonomics* (pp. 47–57). Springer, Cham.
- Dane, E. (2010). Reconsidering the trade-off between expertise and flexibility: A cognitive entrenchment perspective. *Academy of Management Review*, 35(4), 579–603.
- Denning, P. & Dunham, R. (2010). *The innovator's way: Essential practices for successful innovation*. The MIT Press.
- Department of Defense. (2017, May 8). *Autonomy in Weapon Systems* (DOD Directive 3000.09 w/ ch 1). <http://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/300009p.pdf>
- Department of Defense. (2018). *Department of Defense [DOD] artificial intelligence strategy: Harnessing AI to advance our security and prosperity*. <https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMARY-OF-DOD-AI-STRATEGY.PDF>
- Department of Defense. (2020). Implementing the national defense strategy: A year of successes. Department of Defense.
- Department of the Navy [DON]. (2021a). *Department of the Navy science and technology strategy for intelligent autonomous systems*. <https://s3.documentcloud.org/documents/21019491/2021-dist-a-don-st-strategy-for-intelligent-autonomous-systems-2-jul-2021.pdf>

- Department of the Navy [DON]. (2021b). *Unmanned campaign plan framework*.
https://www.navy.mil/Portals/1/Strategic/20210315%20Unmanned%20Campaign_Final_LowRes.pdf?ver=LtCZ-BPIWki6vCBTdgtdMA%3D%3D
- Diukman, A. (2012). *Swarm Observations Implementing Integration Theory to Understand an Opponent Swarm* [Master's thesis, Naval Postgraduate School]. NPS Archive: Calhoun. <http://hdl.handle.net/10945/17358>
- Domingos, P. (2015). *The master algorithm: How the quest for the ultimate learning machine will remake our world*. Basic Books.
- Dotterway, K. (1992). *Systemic analysis of complex dynamic systems: The case of the USS Vincennes* [Master's thesis, Naval Postgraduate School]. NPS Archive: Calhoun. <http://hdl.handle.net/10945/30594>
- Epstein, D. (2020). *Range: How generalists triumph in a specialized world*. Pan Books.
- Espinosa, J., Lerch, F., & Kraut, R. (2004). Explicit versus implicit coordination mechanisms and task dependencies: One size does not fit all. In E. Salas and S. Fiore (eds.), *Team Cognition: Understanding the factors that drive and performance* (1st ed., pp. 107–129). American Psychological Association.
- Fiore, S. & Schooler, J. (2004). Process mapping and shared cognition: Teamwork and the development of shared problem models. In E. Salas and S. Fiore (eds.), *Team Cognition: Understanding the factors that drive and performance* (1st ed., pp. 133–152). American Psychological Association.
- Fout, J., & Ploski, J. (2018) *Unmanned tactical autonomous control and collaboration (UTACC) human machine communication and situational awareness development*. [Master's Thesis, Naval Postgraduate School]. NPS Archive: <http://hdl.handle.net/10945/59661>
- Fridman, L. (2018). *Self-driving cars: State of the art (2019)*[Video]. <https://cle.nps.edu/access/lessonbuilder/item/167371/group/4b62cc67-51c2-4c3e-93e2-5eb50cfac8cc/Su21/Videos/MIT%20SofttheA%202019.mp4>
- Fusco, R., Grassi, R., Granata, V., Setola, S., Grassi, F., Cozzi, D., Pecori, B., Izzo, F., & Petrillo, A. (2021). Artificial intelligence and COVID-19 using chest CT scan and chest X-ray images: Machine learning and deep learning approaches for diagnosis and treatment. *Journal of Personalized Medicine*, 11(10), 993–112. <https://www.mdpi.com/2075-4426/11/10/993#>
- Ginnett. (2010). Chapter 3 – Crews as Groups: Their Formation and their Leadership. In *Crew Resource Management* (Second Edition, pp. 79–110). Elsevier Inc. <https://doi.org/10.1016/B978-0-12-374946-8.10003-2>

- Gombolay, M., Yang, X., Hayes, B., Seo, N., Liu, Z., Wadhwanian, S., Yu, T., Shah, N., Golen, T., & Shah, J. (2016). Robotic assistance in the coordination of patient care. *The International Journal of Robotics Research*, 37(10), 1300–1316. http://people.csail.mit.edu/gombolay/Publications/Gombolay_RSS_2016.pdf
- Gorman, J., Cooke, N., & Amazeen, P. (2010). Training adaptive teams. *Human Factors*, 52(2), 295–307. <https://doi.org/DOI: 10.1177/0018720810371689>
- Gorman, J., Stevens, R., Galloway, T., Willemsen-Dunlap, A., & Halpin, D. (2020). Measuring real-time team cognition during team training. *Human Factors*, 62(5), 825–860. <https://doi.org/10.1177/001872081985279>
- Grosz, B., Altman, R., Horvitz, E., Mackworth, A., Mitchell, T., Mulligan, D., Shoham, Y., Stone, P., Brookes, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., Kraus, S., Leyton-Brown, K., Parkes, D., ... Teller, A. (2016). *Artificial Intelligence and Life in 2030* (p. 52). Stanford University.
- Gutwin, C. & Greenberg, S. (2004). The importance of awareness for team cognition in distributed collaboration. In E. Salas and S. Fiore (eds.), *Team Cognition: Understanding the factors that drive and performance* (1st ed., pp. 177–201). American Psychological Association.
- Han, S., & Choi, H. (2022, January). Checklist for Validating Trustworthy AI. In *2022 IEEE International Conference on Big Data and Smart Computing (BigComp)* (pp. 391–394).
- Hartwick, J. and Barki, H. (1994) Explaining the role of user participation in information system use. *Management Science*, 40, 440–465. doi:10.1287/mnsc.40.4.440
- Headquarters 8th Air Force. (1995). *Record of trial, general court martial, U.S. v. Captain Jim Wang*. Maureen A. Nation, ed., Tinker Air Force Base, Oklahoma.
- Headquarters [HQ] United States Marine Corps [USMC]. (1997). *Warfighting* (MCDP 1). Washington, D.C.: Headquarters U.S. Marine Corps.
- Headquarters [HQ] United States Marine Corps [USMC]. (1996). *Command and control* (MCDP 6). Washington, D.C.: Headquarters U.S. Marine Corps.
- Headquarters [HQ] United States Marine Corps [USMC]. (2017). *Marine Corps operations* (MCDP 1-0). Washington, D.C.: Headquarters U.S. Marine Corps.
- Hirschberg, J., & Manning, C. (2015). Advances in natural language processing. *Science*, 349(6245), 261–266. <https://doi.org/10.1126/science.aaa8685>

- Hoadley, & Lucas, N. J. (2020). *Artificial intelligence and National Security* (CRS Report No. R45178). Congressional Research Service. <https://a51.nl/sites/default/files/pdf/R45178.pdf>
- Hoehn, J. R. (2021). *Joint all domain command and control: background and issues for congress* (CRS Report No. R46725). Congressional Research Service. <https://apps.dtic.mil/sti/pdfs/AD1126249.pdf>
- Howard, A. (2020). Are We Trusting AI Too Much?: Examining Human-Robot Interactions in the Real World. *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction, 1–1*. <https://doi.org/10.1145/3319502.3374842>
- Hrnjica, B., & Softic, S. (2020). Explainable AI in manufacturing: A predictive maintenance case study. *2020 IFIP International Conference on Advances in Production Management Systems, 66–73*. https://link.springer.com/chapter/10.1007/978-3-030-57997-5_8ex
- Ingle, S., & Phute, M. (2016). Tesla autopilot: Semi-autonomous driving, an uptick for future autonomy. *International Research Journal of Engineering and Technology, 3(9)*, 369–372. <https://www.academia.edu/download/54330579/IRJET-V3I969.pdf>
- Jena, A. B., Prasad, V., Goldman, D. P., & Romley, J. (2015). Mortality and treatment patterns among patients hospitalized with acute cardiovascular conditions during dates of national cardiology meetings. *JAMA internal medicine, 175(2)*, 237–244. <https://jamanetwork.com/journals/jamainternalmedicine/article-abstract/2038979>
- Johnson, M. (2014). Coactive design: Designing support for interdependence in human-robot teamwork. *Journal of Human-Robot Interaction, (3)1*. <https://doi.org/10.5898/JHRI.3.1.Johnson>
- Johnston, B. (2011). The collection of physical knowledge and its application in intelligent systems. In *International Conference on Artificial General Intelligence* (pp. 163–173). Springer.
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Strauss, and Giroux.
- Katzenbach, J., & Smith, D. (1993). *The wisdom of teams: creating the high-performance organization*. Harvard Business School Press.
- Kennedy, M. (2021). *OODA LOOP 2.0: A model for completion and expeditionary advanced base operations* [Master's Thesis, Naval Postgraduate School]. <https://calhoun.nps.edu/handle/10945/67756>

- Khan, S., Shiwakoti, N., & Stasinopoulos, R. (2021). A conceptual system dynamics model for cybersecurity assessment of connected and autonomous vehicles. *Accident Analysis and Prevention*, 165. <https://doi.org/10.1016/j.aap.2021.106515>
- Kim, D. (2000). *Systems thinking tools*. Pegasus Communications, Inc.
- Klimoski, R., & Mohammed, S. (1994). Team mental model: Construct or metaphor? *Journal of management*, 20(2), 403–437.
- Kline, J. (2017). Impacts of the robotics age on naval force design, effectiveness, and acquisition. *Naval War College Review*, 70(3).
- Kotter, J. (1978). *Organizational dynamics: Diagnosis and intervention*. Addison-Wesley.
- Kreisher, O. (2016, June 22) *Triton UAV passes full-motion video to P-8 during flight test*. USNI News. <https://news.usni.org/2016/06/22/triton-uav-passes-full-motion-video-p-8-flight-test>
- LaGrone, S. (2017, May 17). *The attack on USS Stark at 30*. USNI. <https://news.usni.org/2017/05/17/the-attack-uss-stark-at-30>
- Lange, D., & Carreno, J. (2021). How AI is shaping navy command and control. In S.J. Tangredi and G. Galdorisi (Eds.), *AI at war: How big data, artificial intelligence, and machine learning are changing naval warfare* (1st ed., pp. 202–219). Naval Institute Press.
- Lawless, W. F. (2020). Quantum-Like Interdependence Theory Advances Autonomous Human–Machine Teams (A-HMTs). *Entropy*, 22(11), 1227.
- Lawrence, P. & Lorsch, J. (1967). *Organization and environment; managing differentiation and integration*. Division of Research, Graduate School of Business Administration, Harvard University.
- Legaspi, A., Mah, J., & Hszieh, S. (2021). Communicating at the speed of war: The future of naval communications and the rise of artificial intelligence. In S.J. Tangredi and G. Galdorisi (Eds.), *AI at war: How big data, artificial intelligence, and machine learning are changing naval warfare* (1st ed., pp. 183–202). Naval Institute Press.
- Levine, J. & Choi, H. (2004). Impact of personnel turnover on team performance and cognition. In E. Salas and S. Fiore (eds.), *Team Cognition: Understanding the factors that drive and performance* (1st ed., pp. 153–176). American Psychological Association.
- Lewis, L., & Vavrichek, D. (2019). *An AI Framework for the Department of the Navy* (DQR-2019-U-021795-Final). Center for Naval Analyses (CNA).

- Lim, B., & Klein, K. (2006). Team mental models and team performance: A field study of the effects of team mental model similarity and accuracy. *Journal of Organizational Behavior*, 27. <https://doi.org/DOI: 10.1002/job.387>
- Lipman-Blumen, J., & Leavitt, H. J. (1999). *Hot groups seeding them, feeding them, and using them to ignite your organization*. Oxford University Press.
- Luhrsen, S. (2007). *Understanding systems theory for U.S. Marines*. [Master's thesis, United States Marine Corps Command and Staff College]. NPS Archive: Calhoun. <https://apps.dtic.mil/sti/pdfs/ADA505411.pdf>
- Macmillan, J., Entin, E., & Serfaty, D. (2004). Communication overhead: the hidden cost of team cognition. In E. Salas and S. Fiore (eds.), *Team Cognition: Understanding the factors that drive and performance* (1st ed., pp. 61–82). American Psychological Association.
- Marks, M., Burke, C., Sabella, M., & Zaccaro, S. (2002). The impact of cross-training on team effectiveness. *Journal of Applied Psychology*, 87(1), 3–13. <https://doi.org/10.1037//0021-9010.87.1.3>
- Massachusetts Institute of Technology [MIT] Horizon. (2022a). *Benefits: The advantages and opportunities of artificial intelligence* [Class Notes for AI 101]. Joint Artificial Intelligence Center. <https://horizonapp.mit.edu/article/benefits-ai>
- Massachusetts Institute of Technology [MIT] Horizon. (2022b). *Common misconceptions: Misunderstood and misrepresented claims about artificial intelligence and the truths underlying them* [Class Notes for AI 101]. Joint Artificial Intelligence Center. <https://horizonapp.mit.edu/article/common-misconceptions-ai>
- Massachusetts Institute of Technology [MIT] Horizon. (2022c). *How AI works: An introduction to artificial intelligence—what it is and how it works* [Class Notes for AI 101]. Joint Artificial Intelligence Center. <https://horizonapp.mit.edu/article/how-ai-works>
- Massachusetts Institute of Technology [MIT] Horizon. (2022d). *Limitations: The challenges and disadvantages of artificial intelligence* [Class Notes for AI 101]. Joint Artificial Intelligence Center. <https://horizonapp.mit.edu/article/limitations-ai>
- Mathieu, J., Rapp, T., Maynard, M., & Mangos, P. (2010). Interactive effects of team and task shared mental models as related to air traffic controllers' collective efficacy and effectiveness. *Human Performance*, 23, 22–40. <https://doi.org/DOI: 10.1080/08959280903400150>

- McNeese, N., Demir, M., Cooke, N., & She, M. (2021). Team Situation Awareness and Conflict A Study of Human–Machine Teaming. *Journal of Cognitive Engineering and Decision Making*, 15(2), 83–96.
- Miller, J. (2021). *Disruptive by design: Gamechanger: There's a new game in town!* Signal. <https://www.afcea.org/content/disruptive-design-gamechanger-theres-new-game-town>
- Mintzberg, H. (1981). Organization design: Fashion or fit? *Harvard Business Review*, 59(1), 103–119.
- Mohammed, S., & Dumville, B. (2001). Team mental models in a team knowledge framework: Expanding theory and measurement across disciplinary boundaries. *Journal of Organizational Behavior*, 22, 89–106.
- Mohammed, S., Rico, R., & Alipour, K. (n.d.). Team cognition at a crossroad: Toward conceptual integration and network configurations. *Academy of Management Annals*. <https://doi.org/10.5465/annals.2018.0159>
- Mohammed, S., Ferzandi, L., & Hamilton, K. (2010). Metaphor no more: A 15-year review of the team mental model construct. *Journal of Management*, 36(4), 876–910. <https://doi.org/10.1177/0149206309356804>
- Mondal, A., Giraldo, J. H., Bouwmans, T., & Chowdhury, A. S. (2021). Moving Object Detection for Event-based Vision using Graph Spectral Clustering. *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 876–884). https://openaccess.thecvf.com/content/ICCV2021W/GSP-CV/html/Mondal_Moving_Object_Detection_for_Event-Based_Vision_Using_Graph_Spectral_Clustering_ICCVW_2021_paper.html
- Moore, M. (1990, 23 April). 2 Vincennes officers get medals. *The Washington Post*. <https://www.washingtonpost.com/archive/politics/1990/04/23/2-vincennes-officers-get-medals/cf383f02-05ce-435b-9086-5d61de569ed8/>
- Morando, A., Gershon, P., Mehler, B., & Reimer, B. (2021). A model for naturalistic glance behavior around Tesla Autopilot disengagements. *Accident Analysis and Prevention*, 161. <https://doi.org/10.1016/j.aap.2021.106348>
- Nachmany, Y., & Alemohammad; H. (2019). Detecting roads from satellite imagery in the developing world. *2019 Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 83–89. http://openaccess.thecvf.com/content_CVPRW_2019/html/cv4gc/Nachmany_Detecting_Roads_from_Satellite_Imagery_in_the_Developing_World_CVPRW_2019_paper.html

- National Academies of Sciences, Engineering, and Medicine [NASEM]. (2021). *Human-AI Teaming: State of the Art and Research Needs*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/26355>.
- National Institute of Standards and Technology [NIST]. (2019, August 9). *U.S. Leadership in AI: A Plan for Federal Engagement in Developing Technical Standards and Related Tools*. https://www.nist.gov/system/files/documents/2019/08/10/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf.
- Nissen. (2014). *Harnessing dynamic knowledge principles in the technology-driven world*. Information Science Reference.
- Oneto, L. (2022). Advances in artificial neural networks, machine learning and computational intelligence. *Neurocomputing*, 470, 300–303.
- Office of the Chief of Naval Operations. (2020, October 1). *Project overmatch* [Memorandum]. Department of Defense. <https://s3.documentcloud.org/documents/7276835/pages/Project-Overmatch-Memos-p2-normal.gif?ts=1604347169298>
- Osinga, F. (2007). *Science, Strategy and War: The Strategic Theory of John Boyd* (Vol. 18). Routledge. <https://doi.org/10.4324/9780203088869>
- Oxford. (2021). Agility. In *Oxford English dictionary*. Retrieved 30 July 2021, from <https://www-oed-com.libproxy.nps.edu/view/Entry/3983?redirectedFrom=agility#eid>
- Papernot, N., McDaniel, P., Goodfellow, I., Jha, S., Celik, Z., Swami, A. (2016). *Practical black-box attacks against deep learning systems using adversarial examples*. arXiv:1602.02697v2
- Peters, G. (2009). The philosophy of improvisation. *The Philosophy of Improvisation*. University of Chicago Press.
- Perry, W. (1994). Memorandum for the Secretaries of the Army, Navy, Air Force, Chairman of the Joint Chiefs of Staff, and Chiefs of Staff of the Army, Air Force, and Chief of Naval Operations: Subject: Aircraft Accident and Corrective Action. Office of the Secretary of Defense.
- Posada, S. (2021). *Autopilot & First Responder Scenes* (Report No. PE 21–020). Department of Transportation. <https://www.nhtsa.gov/sites/nhtsa/files/2021/09/NOA-PE21020-1893.PDF&clen=939438&chunk=true>.
- Rahman, M., Hossain, S., Showail, A., Alrajeh, N., & Ghoneim, A. (2021). AI-Enabled IIoT for Live Smart City Event Monitoring. *IEEE Internet of Things Journal*. <https://doi.org/10.1109/JIOT.2021.3109435>

- Rentsch, J. R., & Klimoski, R. J. (2001). Why do 'great minds' think alike?: Antecedents of team member schema agreement. *Journal of Organizational Behavior: The International Journal of Industrial, Occupational and Organizational Psychology and Behavior*, 22(2), 107–120.
- Rentsch, J., & Woehr, D. (2004). Quantifying congruence in cognition: Social relations modeling and team member schema similarity. In E. Salas and S. Fiore (eds.), *Team Cognition: Understanding the factors that drive and performance* (1st ed., pp. 11–31). American Psychological Association.
- Richards, C. (2020). *Boyd's OODA Loop*. <https://hdl.handle.net/11250/2683228>
- Robertson, G. (2020, July 25). 'Without early warning you can't have early response': How Canada's world class pandemic alert system failed. *The Globe and Mail*. <https://www.theglobeandmail.com/canada/article-without-early-warning-you-cant-have-early-response-how-canadas/>
- Rudd, G.(2004). *Humanitarian intervention : assisting the Iraqi Kurds in Operation Provide Comfort, 1991*. Dept. of the Army.
- Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Penguin Books.
- Saenz, M. J., Revilla, E., & Simón, C. (2020). Designing AI systems with human-machine teams. *MIT Sloan Management Review*, 61(3), 1–5.
- Salas, E. & Fiore, S. (2004). Why team cognition? An overview. In E. Salas and S. Fiore (eds.), *Team Cognition: Understanding the factors that drive and performance* (1st ed., pp. 3–8). American Psychological Association.
- Scales, R. (1993). *Certain victory : United States Army in the Gulf War*. Office of the Chief of Staff, U.S. Army.
- Scharre, P. (2018). *Army of none: Autonomous weapons and the future of war*. W. W. Norton and Company.
- Scharre, P. (2021). The Navy at a crossroads: The uneven adoption of autonomous systems in the Navy. In S.J. Tangredi and G. Galdorisi (Eds.), *AI at war: How big data, artificial intelligence, and machine learning are changing naval warfare* (1st ed., pp. 80–92).
- Scholkopf, B., Locatello, F., Bauer, S., Ke, N., Kalchbrenner, N., Goyal, A., Bengio, Y. (2021). *Towards causal representation learning*. arXiv:2102.11107v1

- Schramm, H., & Clark, B. (2021). Artificial intelligence and future force design. In S.J. Tangredi and G. Galdorisi (Eds.), *AI at war: How big data, artificial intelligence, and machine learning are changing naval warfare* (1st ed., pp. 236–252). Naval Institute Press.
- Senge, P. (2006). *The fifth discipline: The art and practice of the learning organization*. Doubleday.
- Shannon, & Weaver, W. (1963). *The mathematical theory of communication*. University of Illinois Press.
- Shattuck, L., & Miller, N. (2006). Extending Naturalistic Decision Making to Complex Organizations: A Dynamic Model of Situated Cognition. *Organization Studies*, 27(7), 989–1009.
- Shultz, R., & Clarke, G. (2020, August 25). *Big Data At War: Special Operations Forces, Project Maven, and Twenty-First-Century Warfare*. Modern Warfare Institute. <https://mwi.usma.edu/big-data-at-war-special-operations-forces-project-maven-and-twenty-first-century-warfare/>
- Singer, P. W. (2009). *Wired for war: The robotics revolution and conflict in the 21st century*. Penguin.
- Sinval, Aragão e Pina, J., Sinval, J., Marôco, J., Marques Santos, C., Uitdewilligen, S., Maynard, M. T., & Passos, A. M. (2020). Development of the Referee Shared Mental Models Measure (RSMMM). *Frontiers in Psychology*, 11, 550271–550271. <https://doi.org/10.3389/fpsyg.2020.550271>
- Sloan, E. C. (2002). *Revolution in Military Affairs* (Vol. 5). McGill-Queen's Press-MQUP.
- Smith-Jentsch, K., Kraiger, K., Cannon-Bowers, J., & Salas, E. (2009). Do familiar teammates request and accept more backup? Transactive memory in Air Traffic Control. *Human Factors*, 51(2), 181–192. <https://doi.org/10.1177/0018720809335367>
- Snook, S. (2011). *Friendly fire: The accidental shootdown of U.S. Black Hawks over Northern Iraq*. Princeton University Press.
- Song, B., Zurita, N., Gyory, J., Zhang, G., McComb, C., Cagan, J., & Stump, G. (2022). Decoding the agility of artificial intelligence-assisted human design teams. *Design Studies*, 79. <https://doi.org/10.1016/j.destud.2022.101094>
- Stanton, N. & Baber, C. (2017). *Modelling Command and Control : Event Analysis of Systemic Teamwork*. CRC Press.

- Stanton, N., Jenkins, D., Salmon, P., Walker, G., Revell, K., & Rafferty, L. (2017). *Digitizing Command and Control : a Human Factors and Ergonomics Analysis of Mission Planning and Battlespace Management* (First edition.). Taylor and Francis.
- Sterman, J. (2000). *Business dynamics: Systems thinking and modeling for a complex world*. McGraw-Hill.
- Stumborg, M., Brauner, S., Hughes, C., Kneapler, C., Leaver, J., Patel, R., Shields, C., & Collins, E. (2019). *Research and development implications for human-machine teaming in the U.S. Navy* (Report No. DRM-2019-U-019330-1REV). Center for Naval Analysis.
- Sycara, K., & Lewis, M. Integrating intelligent agents into human teams. In E. Salas and S. Fiore (eds.), *Team Cognition: Understanding the factors that drive and performance* (1st ed., pp. 203–231). American Psychological Association.
- Tangredi, S. (2021). AI programs of potential military opponents: Propositions and recommendations. In S.J. Tangredi and G. Galdorisi (Eds.), *AI at war: How big data, artificial intelligence, and machine learning are changing naval warfare* (1st ed., pp. 92–117). Naval Institute Press.
- Tangredi, S. & Galdorisi, G. (2021). *AI at war: How big data, artificial intelligence, and machine learning are changing naval warfare*. Naval Institute Press.
- Tawalbeh, M., Quwaider, M., & Lo'ai, A. (2020, April). Authorization model for IoT healthcare systems: case study. *11th International Conference on Information and Communication Systems (ICICS)* (pp. 337–342). IEEE.
- Thilmany, J. (2007). The emotional robot: Cognitive computing and the quest for artificial intelligence. *EMBO Reports*, 8(11), 992–994. <https://doi.org/10.1038/sj.embor.7401106>
- Tirpak, J. (2021, 7 October). *From the October issue: Unmanned flying teammates*. Air Force Magazine. <https://www.airforcemag.com/article/unmanned-flying-teammates/>
- United States Marine Corps. (2018). *Unmanned aircraft system operations* (MCRP 3–20.5). [https://www.marines.mil/portals/1/Publications/MCWP%203-20.5%20\(Formerly%20MCWP%203-42.1\).pdf?ver=2016-06-06-123323-993](https://www.marines.mil/portals/1/Publications/MCWP%203-20.5%20(Formerly%20MCWP%203-42.1).pdf?ver=2016-06-06-123323-993)
- U.S. Air Force. (1993). Multi-Command Regulation 55–33, *E-3 Operating Procedures Aircrew*. Washington: Headquarters, Department of the Air Force.
- van den Oever, F., & Schraagen, J. (2021). Team communication patterns in critical situations. *Journal of Cognitive Engineering and Decision Making*, 15(1), 28–51.

- van Ments, L., & Treur, J. (2021). Reflections on dynamics, adaptation and control: a cognitive architecture for mental models. *Cognitive Systems Research*, 70, 1–9.
- Verhovek, S. (1995). Air Force Officer Is Acquitted In Downing of Army Aircraft: Officer Cleared in Downing of Copters. *New York Times*.
- Wallach, W. & Allen, C. (2009). *Moral machines: Teaching robots right from wrong*. Oxford University Press.
- Wang, J., & Kosowski, L. (2019). Modelling human attributes in autonomous systems operations (DST-Group-GD-1056). Australian Government Department of Defense Science and Technology Group.
<https://www.dst.defence.gov.au/sites/default/files/publications/documents/DST-Group-GD-1056.pdf&clen=757151&chunk=true>
- Waterman, S. (2021). The misplaced fear of a military technology gap. *Signal*, 75(5), 25–27.
- Weaver, W. (1953). Recent contributions to the mathematical theory of communication. *ETC: A review of general semantics*, 261–281.
- The White House. (2019, February 14). *Maintaining American leadership in artificial intelligence* (Executive order 13859). <https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence>
- Wills, S. (2016). The effect of the Goldwater-Nichols act of 1986 on naval strategy, 1987–1994. *Naval War College Review*, 69(2), 20–40.
- Work, R. (2021). AI, autonomy, and the third offset strategy: Fostering military innovation during a period of great change. In S. J. Tangredi and G. Galdorisi (Eds.), *AI at war: How big data, artificial intelligence, and machine learning are changing naval warfare* (1st ed., pp. 45–62). Naval Institute Press.
- Zhang, R., McNeese, N., Freeman, G., & Musick, G. (2021). “An Ideal Human:” Expectations of AI Teammates in Human-AI Teaming. *Proceedings of the ACM on Human-Computer Interaction*, 4 (CSCW3), 1–25.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

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