Artificial Intelligence for Command and Control of Air Power

Matthew R. Voke, Major, USAF

Wright Flyer Papers

Air Command and Staff College

Evan L. Pettus, Brigadier General, Commandant
James Forsyth, PhD, Dean of Resident Programs
Paul Springer, PhD, Director of Research
John P. Geis, PhD, Essay Advisor

Please send inquiries or comments to

Editor
The Wright Flyer Papers
Department of Research and Publications (ACSC/DER)
Air Command and Staff College
225 Chennault Circle, Bldg. 1402
Maxwell AFB AL 36112-6426

Tel: (334) 953-3558
Fax: (334) 953-2269

Email: acsc.der.researchorgmailbox@us.af.mil

AIR UNIVERSITY

AIR COMMAND AND STAFF COLLEGE

Artificial Intelligence for Command and Control of Air Power

Matthew R. Voke, major, USAF

Wright Flyer Paper No. 72

Air University Press
Muir S. Fairchild Research Information Center
Maxwell Air Force Base, Alabama

Commandant, Air Command and Staff College
Brig Gen Evan L. Pettus

Director, Air University Press
Lt Col Darin M. Gregg

Project Editor
Dr. Stephanie Havron Rollins

Copy Editor
Carolyn B. Underwood

Illustrator
L. Susan Fair

Print Specialist
Megan N. Hoehn

Distribution
Diane Clark

Air University Press
600 Chennault Circle, Building 1405
Maxwell AFB, AL 36112-6010
https://www.airuniversity.af.edu/AUPress/

Facebook:
https://www.facebook.com/AirUnivPress

and

Twitter: https://twitter.com/aupress

Accepted by Air University Press April 2018 and published November 2019.

Disclaimer

AIR UNIVERSITY PRESS

Contents

List of Illustrations

Figures

Foreword

It is my great pleasure to present another issue of The *Wright Flyer Papers*. Through this series, Air Command and Staff College presents a sampling of exemplary research produced by our residence and distance-learning students. This series has long showcased the kind of visionary thinking that drove the aspirations and activities of the earliest aviation pioneers. This year's selection of essays admirably extends that tradition. As the series title indicates, these papers aim to present cutting-edge, actionable knowledge—research that addresses some of the most complex security and defense challenges facing us today.

Recently, The *Wright Flyer Papers* transitioned to an exclusively electronic publication format. It is our hope that our migration from print editions to an electronic-only format will fire even greater intellectual debate among Airmen and fellow members of the profession of arms as the series reaches a growing global audience. By publishing these papers via the Air University Press website, ACSC hopes not only to reach more readers, but also to support Air Force–wide efforts to conserve resources. In this spirit, we invite you to peruse past and current issues of The *Wright Flyer Papers* at https://www.airuniversity.af.edu/AUPress/Wright-Flyers/.

Thank you for supporting The *Wright Flyer Papers* and our efforts to disseminate outstanding ACSC student research for the benefit of our Air Force and war fighters everywhere. We trust that what follows will stimulate thinking, invite debate, and further encourage today's air, space, and cyber war fighters in their continuing search for innovative and improved ways to defend our nation and way of life.

BRIAN HASTINGS
Colonel, USAF
Commandant

Abstract

Computational power, data collection, and algorithm capabilities are increasing at an exponential rate. Artificial Intelligence (AI) advances demonstrate the ability to augment human thoughts and actions in countless areas, among which include the Command and Control (C2) of joint airpower. To triumph in future wars, the United States requires the capability to create multiple dilemmas across multiple domains at an overwhelming speed while preventing the enemy from doing the same. AI will provide the cognitive agility required to C2 forces in providing this capability overmatch. The side with an information advantage and ability to react with high-velocity decision-making will decide the outcome of future wars. This paper attempts to familiarize the reader with some common types and functions of AI, explores specific application areas, and recommends solutions assisting joint targeting using airpower. The development of a weapon to target a pairing system reveals specifics using an example AI creation process. Along with explaining the construction of AI models, this paper also proposes a process for preparing and validating AI for operational use and discusses essential implementation considerations. The desired end state for AI employment in the C2 of joint airpower is efficient human-machine teaming and increased cognitive agility.

About the Author

Major Matthew "Jerry" Voke penned this paper while assigned to ACSC, Air University, Maxwell AFB, AL. Before his assignment at ACSC, Major Voke flew the U-28 while stationed at Hurlburt Field, FL, as a copilot, aircraft commander, instructor, and evaluator pilot. He has deployed ten times around the globe supporting US counterterrorism efforts. Most recently, he instructed United States Air Force Weapons School students as the U-28 Assistant Director of Operations in the 14th Weapons Squadron at Hurlburt Field, FL. His formal military education includes attending the Air and Space Basic Course, Squadron Officer School, United States Air Force Weapons School, and ACSC. He has a Bachelor of Science in Physics from the United States Air Force Academy and a Master of Business Administration from Trident University International. He has helped write several Air Force and multi-service publications and has published a classified Weapons School paper for the vectoring of naval assault forces using aircraft.

Introduction

General "Hap" Arnold once stated that World War I was won by brawn, and World War II was won by logistics.[1] The side with an information advantage and the ability to control the conflict in time will win "World War III." Victory in future conflicts will favor the side that can create multiple concurrent challenges for their adversary and adapt to environmental changes at overwhelming speeds while preventing the enemy from doing the same. These future overmatches are possible by leveraging Artificial Intelligence (AI) in automation and augmented human decision-making. AI provides strength in combining human-level intuitive problem-solving with the speed, accuracy, and persistence of machines. Additionally, rapid advances in AI are forging novel applications in man-machine collaboration.[2] The world is in an AI race, and the United States must leverage AI to dominate future wars and deter potential adversaries.[3]

The potential uses of AI in revolutionizing airpower operations range from tactical to the strategic. China and Russia are among the nations pursuing an AI future, with Vladimir Putin remarking in 2017, "Whoever becomes the leader in [artificial intelligence and cyberspace] will become the ruler of the world."[4] Beijing's leaders similarly state, "The rapid development of AI will profoundly change human society and life, and change the world . . . China must firmly seize the strategic initiative in the new stage of international competition in AI development."[5] Realizing this revelatory potential, the United States must strive to maintain dominance and initiate a culture shift toward developing and implementing AI automation and machine-man teaming.

To develop the ensuing asymmetric advantage, the United States must pursue novel applications of AI in autonomy and augmenting decision-making. By ensuring the lead in AI, the United States can restore the diminishing overmatch against any potential adversaries and hence strengthen deterrence.[6] AI augmentation and autonomy would enable US forces to create multiple dilemmas across multiple domains at overwhelming speeds, rendering human adversaries incapable of reacting if the United States chose to do so. This overmatch requires cognitive advantages only possible with AI that learns, enhances human performance, better prepares human decision-makers, fosters networked solutions, leverages smart devices, coordinates operations, and optimizes the selection and delivery of effects.[7] This paper will specifically look at providing assistance in joint targeting and assessment inside the command and control (C2) of airpower, as an initial implementation area. Although the scope of this paper is limited, the United States should seek a higher level of AI autonomy and augmentation across many areas in C2 to enable the fastest and most practical decision cycles.

The discovery of novel and powerful AI solutions is increasing exponentially, and the US military may be struggling to keep pace with the corporate world. The Department of Defense (DoD) budgeted $195 million in 2017 for machine learning, $238 million applied to deep learning, and $82 million toward language processing.[8] At first glance, these numbers may seem like a significant amount of money devoted to AI research; however, DoD investments and research are growing significantly slower than the corporate sector. Between 2012 and 2017, overall DoD investments in AI, big data, and cloud research grew at 5.7 percent compound annual growth rate (CAGR).[9] Corporate investments and AI research are skyrocketing at a CAGR of approximately 35 percent, with AI driving change in how businesses interpret data and interact.[10] Corporate investments in AI are doubling roughly every two years. DoD investments in AI are not keeping pace with corporate investments or AI's exponential growth.

C2 of joint airpower has historically been effective at deliberately planning operations against less capable adversaries, although it is arguably not efficient. Efficiency is likely to be a substantial requirement against a near-peer enemy or in a highly dynamic war. The *National Defense Strategy* states the DoD is "over-optimized for exceptional performance at the expense of providing timely decisions, policies, and capabilities to the warfighter" and must "deliver performance at the speed of relevance."[11] Specifically, joint targeting has proven effective but not optimally efficient at directing a large number of assets against preplanned targets, as exampled in Desert Storm, Operation Allied Force, and Operation Enduring Freedom.[12] The United States conducted these deliberate operations in a preplanned manner, in relatively small wars, against less capable adversaries. Deliberate planning typically begins 24–to–72–plus hours from operations, while allowing for some exceptions in dynamic targeting.[13] US C2 has reinforced lessons learned from past success against less capable adversaries, growing more effective; however, efficiency and agility required for combating a near-peer adversary are yet to be developed.

The United States must continuously strive to keep airpower employment efficient and agile. As Colonel John Boyd said in *Patterns of Conflict*, "in order to win, we should operate at a faster tempo or rhythm than our adversaries—or, better yet, get inside the adversary's observe-orient-decide-action time cycle or loop."[14] AI can supplement human efforts in many areas to significant effect and speed. US forces must initiate the next airpower evolution and leverage AI to augment human decisions and actions. AI solutions are necessary to facilitate the speed, strength, balance, flexibility, and coordination necessary to create multiple dilemmas across multiple domains with speed in a future conflict. The United States must remain the leader in air, space, and cyberspace power, especially as Russian and Chinese investments in these areas grow.

AI is showing exponential growth in ability and speed in the corporate world, disrupting areas from natural language processing to a cancer diagnosis. Some experts liken the surge in AI to the invention of electricity in the late 19th century, stating it will likely start an equally significant transformation of industries.[15] Advances in AI are making it possible to cede to machines many tasks long regarded as impossible for machines to perform. As machine-to-machine communications and man-to-machine teaming solutions continue to mature, AI will become increasingly prevalent in human-intensive processes.[16]

One of the most significant areas of AI contribution lies in its ability to draw correlations from data, potentially invisible to humans. Sensors and perceivers collect or create data—which then must then be stored, cleaned, and structured. Processing turns collected raw data into usable AI fuel. Algorithms can then create models, develop and test insights, draw correlations, and detect anomalies for human decision-makers or actors.

AI can simultaneously make decisions on a time scale incomprehensible to humans. AI can assist humans with target discovery, intelligence fusion, target prioritization, commander's analysis, assessment, force assignment, mission planning, mission monitoring, and execution. AI will enable rapid decisions across multiple domains and multiple levels of war that humans cannot outpace. AI will aggregate, integrate, distill, and present common operational pictures and aid in decision-making cycles, achieving effects that humans today cannot grasp. AI may prove capable of shortening the lengthy joint targeting cycle into a loop that rapidly updates at varying speeds, quickly reacts to environmental changes, and is agile and aggressive. Every area of human thought or action is ripe for AI disruption, including airpower employment.[17]

America must posture itself to respond to relevant information with high-velocity decision-making.[18] Automation and machine-man teaming will improve and expedite decision-making and tasking cycles in the C2 of airpower. Victory in future conflict will go to leaders who can command, control, and direct their forces at a pace that overwhelms their adversaries.[19] Future wars will favor the belligerent leader that can rapidly process information and make decisions at each level of war. The warfighting goal is not only to react to the enemy quickly; but also it is to drive the fight in directions and at the pace of one's choosing. The belligerent leader that is three steps ahead of the enemy will set their adversary "back on their heels," continually struggling to react, and unable to gain the initiative

This paper attempts to familiarize the reader with AI's potential and highlight critical considerations for employment in the C2 of airpower. This paper's primary focus is to demonstrate the latent AI potential that the United States must leverage in tomorrow's fights. The goal is that the reader will be galvanized through specific application examples. The author's hope is that readers will walk away with a fascination and foundational familiarity with AI. The required precursor to the development and fielding of the next asymmetric overmatch is an innovative culture and a desire to adapt, fostered through encouraging creative problem-solving and stoking innovative passions throughout the joint force.

This paper briefly introduces some essential C2 principles and processes—setting the stage—and highlights areas for potential AI applications. Second, it describes the balance of centralization and some principles of US C2 joint airpower. Third, it surveys US joint targeting and target selection processes. Fourth, a survey of AI basics and examples of areas for using AI will provide a foundation of concepts and common AI types, followed by several examples of using AI in military applications. This allows for a discussion of a weapon-to-target pairing model that illustrates how to handle data, algorithm selection, and decision-making with model outputs. The paper then discusses AI-aided assessment, including dominant indicators and feedback loops. Finally, essential implementation considerations will cover required steps toward fielding and changes for successfully building trust and aiding decision-makers in the employment of AI solutions in airpower.

Principles of Command and Control (C2) of Airpower

This section will briefly acquaint the reader with some C2 basics in the employment of airpower. Succinctly, C2 is "The exercise of authority and direction by a properly designated commander over assigned and attached forces in the accomplishment of the mission."[20] Joint air operations are operations performed by forces made available by the services for joint air tasking.[21] Joint air operations can vary wildly across operational environments and the range of military operations; however, it is crucial that the reader have a foundational understanding of the C2 of airpower.

Centralization Versus Decentralized Airpower

US doctrine prescribes the conduct of joint air operations using the principle of centralized control and decentralized execution.[22] An oversimplified synopsis driving this tenet of airpower is that centralization of control enables the senior echelon commanders to control, mass, and lead forces effectively; decentralization of execution allows forces to seize the initiative, respond to uncertain and changing environments, and fosters flexibility in lower echelons.

Technological developments frequently shift the equilibrium of this tenet of airpower. Robust communications connectivity has increased the shared operations picture at all levels, but it has also enabled senior leadership involvement in the finest details of employment.[23] This duality has created inherent tension between the imperatives of political control and those of efficient mission accomplishment that leaders must understand.[24] Although centralized control and execution are possible in many situations, a conscious effort to delegate execution authority appropriately will ensure the maintenance of US airpower agility.

The balance of centralization and decentralization can shift between and during conflicts, and leaders must strive to increase what some call "agility" in airpower employment. The *Air Force Future Operating Concept* defines agility as the ability to react rapidly to situations. Agility is a combination of one's flexibility, speed, coordination, balance, and strength.[25] Decentralization generally favors flexibility and speed, while centralization generally favors coordination and strength. One's understanding of agility might imply physical capability, but agility also includes a cognitive capability to react to a dynamic opponent, moving target, or shifting environment.[26] The power of AI employment in airpower C2 lies in the cognitive speed and strength it can bring synergistically toward the goal of agility.

C2 Structure

C2 systems control joint air operations, typically built around the C2 system of the service component commander, that has the preponderance of air assets and the most exceptional ability to control them.[27] The air operations center (AOC) is the senior element of the theater air control system, which ensures the effective planning and conduct of air, space, and cyberspace operations. The AOC construct may also apply when fighting with joint or coalition partners as a joint air operations center or a coalition air operations center. The size of an AOC can vary wildly between staff in the single digits to more than a thousand officers, enlisted, and civilian members. Each AOC's organization differs, but their common goal is to match available means toward tasked military objectives. The responsibilities of the AOC typically include planning and controlling joint air operations, recommending priorities in air apportionment, airspace coordination, air defense coordination, space coordination, and cyberspace coordination.

Technological improvements in computing, communications, and information sharing have disrupted the tiered C2 structure creating disruptions from operational planning to tactical execution. AOC leadership is capable of making decisions historically conducted at lower operational or tactical echelons. Today, for example, the highest levels of operational C2 may not delegate target identification and weapons release authority, which was once only possible at the tactical edge. The relative overmatch wielded when facing less capable adversaries, and the aversion to civilian casualties and loss of life may have lulled the United States into complacency and a false assumption that future wars will take place in the same benign environments.[28] Increases in efficiency and agility needed to defeat near-peer adversaries are possible through distributing control outside the AOC, decentralizing the execution of air assets, and leveraging AI's speed and cognitive strength across multiple levels of C2.

The C2 of airpower has excellent potential to leverage AI augmentation and automation for increased cognitive agility. The next section discusses one example of AI use area in airpower C2—joint targeting. There are currently data gaps and inefficiencies within the targeting cycle that complicate the efficient transfer of critical information when planning and executing joint air operations.[29] The next sections will cover how AI can prepare decision-makers with a better understanding of their operating environment, filter and fuse the fire hose of battlefield data into relevant information, and increase decision speeds.

Joint Targeting

Joint targeting's purpose is to match available means provided to the joint force commander (JFC) with ways of employing air, space, and cyberspace power to meet military and political ends. It is the process of selecting and prioritizing targets and matching the appropriate response to them, considering operational requirements and capabilities.[30] Joint targeting can occur at whatever level of war necessary to achieve the JFC's objectives. Joint targeting selects targets, matches desired effects to those targets, and finally selects the means capable of delivering the required effects.

Complicating the effort of matching ways and means to targets is the fog and friction of war, target duplication, unknowns, integration requirements, second and third-order effects, and lack of available resources. Joint Publication (JP) 3-30, *Command and Control of Joint Air Operations*, highlights the importance of efficiency, stating, "An effective and efficient target development process coupled with the joint air tasking cycle is essential for the Joint

Force Air Component Commander (JFACC) to plan and execute joint air operations. The joint targeting process should integrate the intelligence database, analytical capabilities, and data collection efforts of national agencies, combatant commands, subordinate joint forces, and component commands."[31] This section will discuss how AI can aid in bringing the required efficiency to the joint targeting cycle.

To gain efficiency in the joint targeting cycle, the individual steps and automated connections of steps in the process must gain efficiency. The six phases of joint targeting, as shown in Figure 1 include;

1. End state and commander's objectives development,
2. Target development and prioritization,
3. Capabilities analysis,
4. Commander's decision and force assignment,
5. Mission planning and force execution, and
6. Assessment.[32]

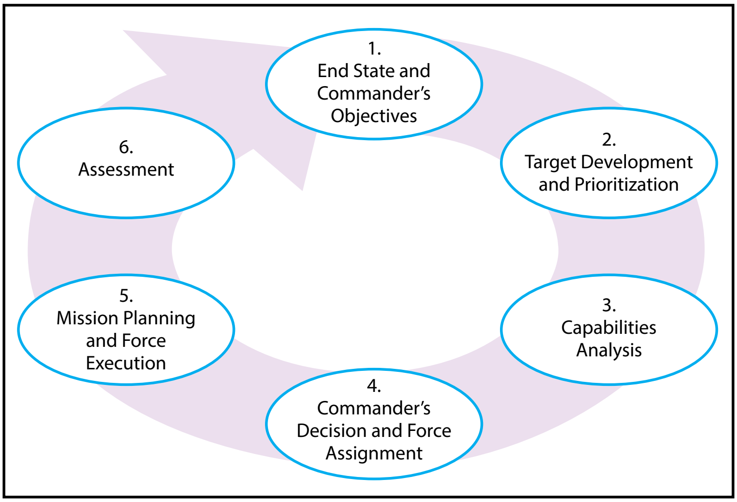

Figure 1: Joint Targeting Cycle (JP 3-60)

AI will likely hasten and optimize target development and prioritization, provide near-instantaneous capabilities analysis, rapidly and iteratively plan against targets, and automate the development of attack plans.

C2 staffs conduct joint targeting in the joint air tasking cycle, which is deliberate by design but allows for some execution of dynamic targets outside of the cycle. The JFACC can change the process to match the environment; however, the airpower tasking cycle itself remains deliberate and generally fixed in duration. Claims of expeditious or efficient dynamic targeting or time-sensitive targeting inside the air tasking cycle are relative to past capabilities and timeliness. Any "pop-up target" generally requires at least 12 hours for action within the air tasking order (ATO) process. Exceptions to this rigidity are flexibly tasked assets, planned or launched without a specific target, and able to react to real-time opportunities. The agility to reroll deliberately tasked assets using the ATO process within the 12-hour window is generally minimal, and the current operations division will likely provide tasking outside the normal airpower tasking process. The flexibility and speed required in time-sensitive or highly dynamic targeting occurs despite the rigid joint air tasking cycle, not with the assistance of it.

One of the most significant difficulties in targeting is the efficient and comprehensive melding of targets on varying time scales. The joint forces continue to discover, prioritize, and pair effects against target sets that may be different at the time of execution. Moreover, target development cycles occur at different speeds, and sensors may discover high priority targets late in the rigid targeting process. The United States requires the ability to adjust to the environment and enemy actions as instantaneously as possible. Joint targeting cycles assisted by AI could prosecute targets both deliberately and dynamically, and be able to match effects and prioritize targets using continually updated cycles. When C2 pairs an asset to a high-value "pop-up" target, their original target should instantaneously flow back into the targeting solution, be re-prioritized, and potentially handed to another strike asset. The scale and speed of conducting such a targeting cycle is only possible with the cognitive speed and strength of AI. A joint air targeting cycle fueled by AI could be deliberately dynamic and continually iterative.

Targets

The definition of targets is any person, place, or thing considered for possible action to alter, degrade, or neutralize the function it performs for the adversary.[33] The JFACC prioritizes targets by contribution level toward the joint force objectives, probability of achieving the desired effects, the cost of engagement, and many other factors.[34] Targets can be people, facilities, equipment, or infrastructure; conversely, targets may also have digital capability or be an entity in cyberspace. Target analysis in joint targeting defines not only targets but also their strengths, weaknesses, and interconnections.

Planners must document target characteristics in detail so that they can develop, correlate, and fuse targets. Characteristics are specific to the target type, but can generally include specifics, such as location, size, detailed appearance, target makeup, dispersion, hardening, electromagnetic signature, emitters, and mobility. Humans currently manually input target characteristics, compile targets, correlate targets, eliminate redundancies, and prioritize targets through disconnected databases incapable of robust machine-to-machine communication.

The JFACC's staff then prioritizes targets, matches them with effects, and tasks their execution with available assets, input, and coordination from component commanders. To close the metaphorical feedback loop in targeting, assessment teams analyze effects from operations in meeting the military objectives. An assessment provides measurements of success that feed back into training human decision-makers and could have applications for teaching AI systems. This iterative loop reinforces optimal behaviors and attenuates poor behavior.

Why AI?

Implementing AI-aided intelligence fusion and target prioritization could dynamically recommend re-tasking assets faster than any human process at scale. The current deliberate and dynamic targeting cycles employed in the AOC contain minimal automation, are rife with redundant human efforts, and lack effective cross-communication. Automating human-intensive tasks inside the targeting cycle could tighten the observation, orientation, and decision loops. Even basic automation facilitating the effective flow of objectives, tasks, targets, characteristics, and weapons between the steps of the cycle would shave hours off the processes. More comprehensive AI implementation would facilitate near-instantaneous fusion of target characteristics, prioritization to match objectives, matching of ordnance, and adjustment of air tasking. One goal of AI implementation is to reduce the joint air targeting process from the current three-to-five days to one without a fixed duration that also iteratively adapts to meet the changing operational environment and threat in real-time.

AI Background

As Pablo Picasso once said of computers: "but they are useless. They can only give you answers."[35] The witty insightfulness of his comment implies the inability of computers to ask important questions or express creativity. Picasso would be awestruck, a half-century later, at the apparent creativity expressed by AI and the deeply perceptive questions they are now capable of asking.[36] Definitions and the perceptions AI vary, but for this paper, AI is *an unnatural agent with the ability to learn and adapt to changes, on par or better than humans.*[37] The next section will explore AI as a foundation for later discussion, and it will differentiate between narrow and general AI with examples, and then will explore the concept of exponential growth. It will conclude with a discussion of when decision-makers should seek AI implementation in solving problems.

Types of AI

Narrow AI is AI designed and trained for a specific purpose of limited scope[38]. AlphaGo is one example of narrow AI. AlphaGo was the first machine to beat world champions in the game of Go in 2015 and 2016, but it cannot fold a towel or predict weather conditions.[39] Narrow AI has shown immense ability to defeat human experts in select areas, but adapting to unforeseen environments or different applications proves challenging.

Artificial general intelligence (AGI), in contrast, is AI that can work and learn across a broad spectrum of areas.[40] An example of an attempt at broadening toward AGI is Q-network, an AI developed by DeepMind.[41] Q-network is a deep reinforcement-learning agent that can achieve greater than human-level performance on 49 Atari games.[42] Q-network's only inputs are pixels on the screen, and it uses the same broad algorithm and model structure across the spectrum of games. Q-network is an example of the progress made in high-dimensional sensory inputs, "resulting in the first artificial agent that is capable of learning to excel at a diverse array of challenging tasks."[43] Although Q-network is proficient across multiple types of Atari games, it is not proficient at all Atari games. Experts continue to make progress toward AGI, but "general" intelligence remains elusive and is more difficult to achieve than narrow AI.

Although AGI currently remains out of reach, engineers are attempting to broaden the application of AI in certain areas. By taking a narrow AI and increasing its peripheral application to similar or joined areas, an AI can develop breadth beyond its particular narrow focus. An example of the broad application of AI is traffic management. An early large-scale implementation of this is New York's "next-generation traffic control system" which controls 12,400 traffic signals, creating the most extensive traffic control system.[44] The system incorporates time-of-day, construction, motorcades, accidents, unique event data, real-time field sensors, radio-frequency identification readers, and cameras to improve the flow of traffic. By increasing the diverse inputs, learning, outputs, and communication means, the system has broadened from its initial narrow purpose to overall traffic control. Although the scope has broadened, the system does not possess general intelligence applicable outside of its specific use area. Narrow and broad applications of AI have demonstrated AI's capabilities and have paved the way for future use in the C2 of airpower.

Although AI becomes more and more capable of problem-solving and continually increases in breadth and depth, the metaphorical goalposts defining intelligence keep moving further away. The requirements that define intelligence continually increase with scientific advances. According to McCarthy, "as soon as it works, it no one calls it AI anymore."[45] The Turing test—coined in 1950 by British mathematician Alan Turing—is an attempted objective test of a machine's ability to demonstrate human-level intelligence.[46] The Turing test posits that a machine is intelligent if a human judge cannot distinguish a machine from a man in a moderate length text-based conversation. Several applications of AI have passed the Turing test since it was envisioned, including one named Eugene Goostman in 2014.[47] Some experts are more skeptical, for instance, Ray Kurzweil claims the victory as "premature," rife with restrictions, and able to fool naïve judges.[48] Simply defining machine and human intelligence is not a straightforward matter.[49] Although human perceptions of what defines intelligence constantly shifts, AI's proficiency and use areas continues to grow at an exponential rate. When assessing AI's intelligence for specific applications, simple metrics can be speed and accuracy when compared to human equivalent operations. Leaders should implement AI solutions capable of performing tasks quicker, cheaper, and more accurately than human performance.[50]

Exponential Growth

An understanding of exponential growth is required for the understanding of AI's potential. Humanity's technological growth is producing outsized gains yearly. One example of exponential growth is Moore's Law—the number of calculations that one can purchase for $1,000 will double roughly every 18 months. See Figure 2 for a graphical depiction of the exponential growth of transistor density.[51]

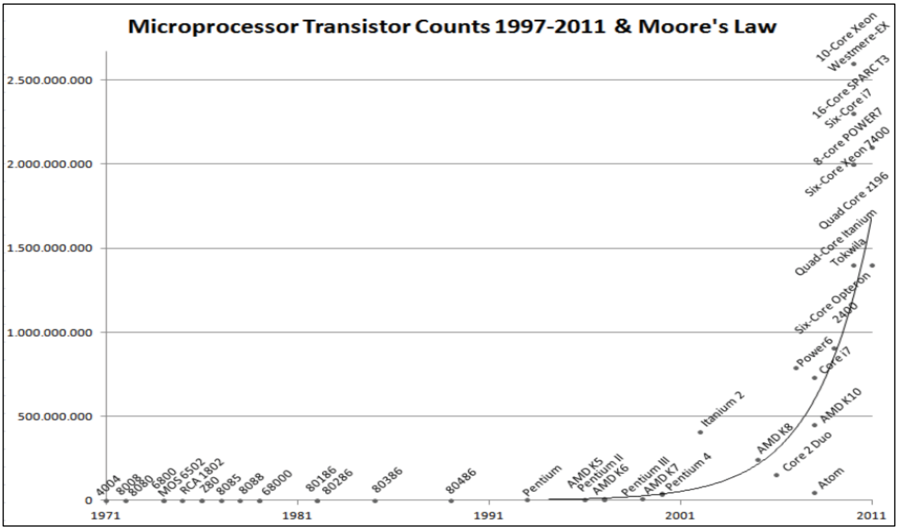

Figure 2: Moore's Law Transistor Counts

Exponential growth is the growth of a quantity for which the rate of growth is directly proportional to the amount present.[52] As the amount increases, the rate of growth increases. Because humans are comfortable thinking linearly, it is sometimes difficult to understand the importance of exponential growth upon the future.

There are many examples that explain the importance of exponential growth; perhaps the easiest is the invention of the chessboard. Ancient Chinese lore tells of a farmer who invented a game of strategy that was a gameboard— with 64 squares and 32 wooden figures—that he presented to the provincial king. The king was overjoyed after the farmer explained the strategy and gameplay that resulted from such a simple invention. The king promised to reward the farmer with anything the farmer could name for this genius idea. The farmer stated that he is a simple peasant farmer, and he does not require much in return. All the farmer requested was some grains of rice out of the massive provincial stores. Since he invented the 64-square board, he proposed payment using the board. On the first square, one grain of rice is the reward, and then on each successive square, the payments are doubled—2, 4, 8, and 16. The king realized the farmer could have asked for gold or any other valuable commodity and is overjoyed to render payment in such a humble manner. However, less than halfway through the board, the king realizes the power of exponential growth. The first square on the fourth row would require a metric ton of rice delivered to the farmer. The outcome of the story differs from this point. Half of the stories speak of the peasant taking the king's throne; other versions express the king's displeasure at being deceived by executing the peasant. Perhaps Earnest Hemmingway explained exponential growth the best: "it happens first slowly, then all at once."[53]

Humans have a difficult time understanding the implications of exponential growth. If a human traveled back in time 50 years with a smartphone, people would not believe it possible. Today's smartphones not only have more capability than the computers that took the Apollo astronauts to the moon; today's smart phones have more computational power than the National Aeronautics and Space Administration (NASA) at that time.[54] Some examples of areas leveraging exponential growth include data production, data storage costs, and computational power. A failure to understand exponential growth prevents some from seeing the importance of technology's impact and its potential. Because of the impact of exponential and combinatorial technological change, the future is quite unlikely to be an extension of the present.[55] Future assumptions, frameworks, and underlying logic will have changed. One cannot extrapolate forward today's requirements and situations in predicting tomorrow's wars. If the United States does not develop and leverage technology to meet tomorrow's challenges, someone else will.

Data

Humans produce data at exponentially increasing rates; for example, 2017 produced more data than in the previous 5,000 years combined.[56] In addition to the advances in data production, the AI algorithms and modeling used to interpret the data have also significantly advanced. AI can recognize patterns in massive amounts of data, highlighting correlations. AI can predict future states or outcomes from model correlations or classifications, sometimes better than humans can. The ability to predict future events hinges on the information of today's circumstances mixed with corollary and causal relationships. In theory, the current situational context and all relationships can be known or approximated. The challenge with predicting the future lies in the small deviations from reality, causing significantly outsized errors, and as humans, AI is not immune to this pitfall.[57]

AI cannot exist without data from which to train and learn. AI learning is only possible with sufficient amounts of data to draw meaningful correlations. AI not only needs large amounts of data to learn, but it also requires structured and applicable data. As an example, think about a house-price prediction model. Some data collected on home sales are more usable for predicting sale prices than others are. Square footage and the number of bedrooms will likely correlate more directly to a sales price than the color of the house, for example. Having massive amounts of data is critical to learning; however, it must contain the correct data.[58]

The power of data is in its ability to explain the world, often in more ways than humans can observe. Computer algorithms previously required programmers to direct every action the computer took. AI is currently a mixture of programmer direction and learning from data. Recent advances in AI show possibilities in letting AI design tailored AI for specific applications. Humans should leverage AI's strengths in understanding environments and AI's ability to predict events when they are superior to humans.

Machine Learning

One category of AI is machine learning. Machine learning is a subset of AI where machines learn from data without explicitly programmed rules.[59] The object of machine learning is to find an appropriate model and fit the model to the data, instead of using parameters defined by a human. Machine learning uses training, or "learning," to adjust the parameters of the model through continuous optimization of performance in matching data. In some ways, machine learning is more about probabilities and statistics than it is about programming and robotics. Some machine learning models are incredibly intricate and complicated structures, while others can be straightforward and easily understood. To paraphrase Einstein, machine learning models should be as simple as possible but no simpler.[60] Additionally, data scientists and machine learning experts must avoid overfitting models to explain data too specifically.[61] The three subcategories of machine learning include supervised, unsupervised, and reinforcement-learning as shown in Figure 3.

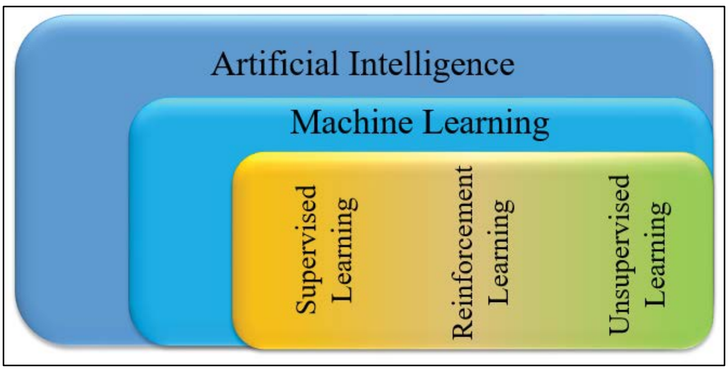

Figure 3: Types of Machine Learning

In supervised learning it is typically faster to optimize a model because known labels or classifications accompany the input data. For x input, there is a known y output. Primary example areas for supervised learning include classification and regression.[62] A known output allows for the comparison of a model's prediction to the true correlation, classification, or prediction. By comparing the model's prediction to known truth, errors can be iteratively calculated and then minimized. The prediction error drives small changes in the model. Like a drop of water, attempting to descend an incline, the fastest way down is to find the largest descending gradient. Both the water and machine learning models improve by repetitively running through adjustments downward (decreasing error by the steepest gradient). Models trained through machine learning run quickly, using minor resources, and can scale as much as needed with minor oversight. An additional benefit of supervised learning is the ability to update a model's training to unexpected or unique situations using new 'known truths,' such as those one might see in the opening days of major combat operations.

Unsupervised learning differs from supervised learning in that the output or classification of the data may be unknown. Unsupervised learning can use unlabeled data. Nobody can teach or supervise the learning of the algorithm and model because humans do not necessarily know what to instruct. For example, humans can very quickly determine whether a picture contains a cat or not, but it would be extremely difficult or impossible for a human to code a computer on how to find cat features in pictures without machine learning. An example use of unsupervised learning occurred in 2012 when Google programmers were able to train a neural network on unlabeled data, teaching itself to recognize the facial features of humans and cats.[63] Some unsupervised learning areas include clustering, anomaly detection, adversarial networks, and blind signal separation. Unsupervised learning explores data with unknown or hidden patterns.

The third type of machine learning is reinforcement-learning. Reinforcement-learning was inspired partly by theories of animal cognition and reinforcement techniques.[64] It attempts to learn an optimal policy, through trial and error, for sequential decision-making problems, similar to giving a dog a treat when performing a trick.[65] There are many types of reinforcement-learning, but the goal is to achieve the ideal behavior over time by maximizing reward. One example of reinforcement-learning is the Q-network example mentioned earlier in the broad intelligence discussion. Q-network learned by attempting to maximize Atari game scores. The reinforcement-learning rewarded higher scores, reinforcing specific learned behaviors. Reinforcement-learning is useful because of its broad application; however, it can be very memory intensive and costly. Reinforcement-learning was one of the Massachusetts Institute of Technology's (MIT) *Technology Review* "10 Breakthrough Technologies of 2017," suggesting it will play a crucial role in achieving AGI.[66]

Machine learning and prediction is possible because the world is full of regularities.[67] As Mark Twain said, history does not repeat itself, but it often rhymes.[68] One of the advantages of machine learning is that it can explain unknown or significantly complicated relationships with simpler models by discerning underlying correlations. Models can be adjusted for simplicity while explaining as much of the data as possible. In machine learning, like human learning, the simpler the model, the easier and quicker the learning. The downside to simplicity emerges when the assumptions are too broad, creating errors in explanations or predictions. The world is full of complexity, and model optimization requires confidence and accuracy after the simplification of complexity. In machine learning, balances occur during selection, development, and training, and there is not a cookie-cutter option for machine learning development. Every application of machine learning is unique because every problem and data set is unique.

Neural Networks

Neural networks are one approach to machine learning that many people know and their inspiration comes from the human brain's network of neurons. Some early efforts to create AI developed "neurons" that would "fire" when they reached a sufficient threshold from previous neuron inputs.[69] This coupling of neurons creates a network of nodes that individually participate only a minuscule amount, however together they create capability to recognize or react to patterns. See Figure 4 for a graphical depiction of a simple neural network.[70] In this simple, example neural network, there is one input layer, one hidden layer, and one output layer.

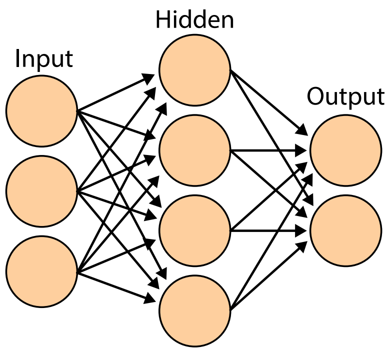

Figure 4: Neural Network

The input layer perceives a tiny bit of data from the environment, and then acts on its learned perception, sending output to the next layer. The next layer is fed by the outputs of the previous layer and acts according to its learned behavior to provide outputs to the next layer. The behavior of the input and output layers are easy to see and understand. The inputs may be numbers or a pixel from a picture. The outputs are typically classifications or regressions explaining data. However, it is sometimes difficult to understand the hidden layer's learned behavior because it may have learned behaviors incomprehensible to human observers. The hidden layers might be discovering the edge of a picture or a unique feature of a class of objects that human observers may not understand. Because of a human's lack of interpretability from the middle layers, they are termed "hidden." Additionally, if a neural network has multiple hidden layers, it is deemed a "deep neural network."

Neural networks are efficient as classifiers, providing confidence in correlating an output from input data. They are also capable of performing maximum likelihood estimation.[71] For example, a neural net could provide an imagery analyst the likelihood that an image contains a subsonic cruise unarmed decoy (SCUD) missile or other specified targets. Each neuron in the neural network interprets a small piece of data, with the first row taking in the pixels of the image. Subsequent layers of neurons may make more abstract or more profound conclusions from there (e.g., an edge or specific image feature) that could lead to overall understanding. The output of a neural network could either be as simple as providing the likelihood the picture contains a SCUD or as complex as trying to classify every object the image contains.

Where to Use AI

AI works well when underlying patterns or correlations exist in data. Therefore, one should seek the use of machine learning where correlations likely exist, and there is a sufficient amount of training data upon which to draw these correlations. Second, the model that attempts to explain the underlying correlations should be simple enough to solve future problems quickly. There are exceptions to these generalizations, but these determinations are a good starting point for exploration. AI needs data that contain correlations to make conclusions about the environment and predictions. Humans are currently more effective at solving novel or unstructured problems; AI thrives in environments with high correlation.

AI currently shows superiority over humans in several areas. First, machine learning excels in situations where humans cannot determine the rules governing the data or relationships. This determination is possible where a large number of variables exist, there are complicated interconnections of the variables, or humans are not able to "label" the data. An example of this would be in predicting diseases based on DNA. Unaided, humans would not be able to process all the data in one human genome, much less tie correlations from thousands of examples. The other niche where machine learning excels is when humans may understand correlations, but rules cannot explain them through brute force coding. It is easy for humans to recognize handwritten characters, but it is prohibitively difficult to program a computer with all the rules required to recognize handwriting. AI superiority in this situation emerges by the ability to scale the application relatively cheaply. AI reads and automatically routes mail, reducing postal service staffing requirements and decreasing costs to the consumer. Another area of machine learning potential is in performing tasks that humans can accomplish; however, the situation regularly changes, and it would be cost-prohibitive for humans to reprogram the machines constantly. AI can update perceptions almost instantaneously based on billions of data points. An example of this area is Amazon's recommendation system. Using AI, Amazon updates recommendations for products in near-real-time, tailored explicitly to hundreds of millions of people. AI iteratively adapts and learns with changing data where there are too many instances to understand or program individually.[72] AI is useful in areas where humans do not understand the rules governing data, they understand the rules but cannot accurately code the solution, or where it would be cost-prohibitive or impractical to scale or iteratively adapt the solutions.

Humans maintain relative superiority in areas surrounding creativity, originality, responsibility, and empathy.[73] Although AI can pen original sonnets and paint distinctive masterpieces that judges are unable to differentiate from human artists, their creativity and originality are always matters of debate. These AI are essentially emulating creativity and originality through learning to mimic human examples. These areas of superiority may be eroding, however. One example of this erosion is the transition from DeepMind's AlphaGo, which learned to expertly play Go through analyzing a human play, to AlphaGo Zero, which learned solely through reinforcement-learning and self-play (not observing human-played games).[74] Machines are forcing humans to reevaluate the definitions of originality and creativity. Responsibility is perhaps the most "black and white" area currently dominated by humans. In the business of war, humans should not delegate responsibility to machines because they cannot be held accountable for their decisions and actions. Finally, AI is incapable of understanding the intricacies of human empathy, morals, beliefs, values, or implicit purposes. AI may learn to simulate ethics or follow rules of engagement; however, AI is unlikely to internalize human empathy, or the ability to think critically about human morals, values, and some actions anytime soon.[75]

For these reasons, the nature of war will likely remain a predominantly human endeavor for the predictable future.[76] Humans will retain superiority in operational design and areas characterized more closely with the "art of war." AI will see a surge in data handling, information fusion, prioritization, analysis, processing, and procedures commonly referred to as the "science of war."' AI will aid humans across the levels of war—ranging from assistance to full autonomy—each complementing the other's strengths and asymmetrically applied against an adversary's weaknesses.

Creating AI Models

The previous sections created a foundation for the reader in joint targeting inside airpower C2 and AI basics; the next step is to create an AI framework built on top of that foundation. There is an infinite number of ways to approach each AI application, and this section will explain one generic framework for the nontechnical reader—fear not, there is not any math or programming involved! This generic framework will step chronologically through;

1. How to identify use areas for AI,
2. How to find and prepare data,
3. How to select the algorithms and generate models, and then,
4. How to create appropriate decision rules and use the outputs.[77]

A weapon-target pairing model is included as an example in each step to add some specificity and understanding of the concepts.

Identify Use Area

The obvious first step in creating a machine learning solution is to identify use areas, such as staff-hour intensive or highly repetitive tasks. AI is good at solving problems that contain strong correlations or example cases.[78] Additionally, it is essential to remember not to force AI solutions onto problems just for the sake of having AI solutions. AI should provide accuracy, speed, and/or increased efficiency when addressing a problem.

When analyzing use areas, it is best to start implementation with relatively simple problems and work toward complex problems. What problems are relatively easy to solve by humans but require massive amounts of time or repeat frequently? Joint targeting is an area where humans are easily able to link specific targets to desired effects, but the process changes frequently and is staff-hour intensive. Inside joint targeting, several areas would benefit from AI: weapon-target pairing, mission-aircraft pairing, target prioritization, battlefield situational awareness, blue and red force tracking, system node analysis, master air attack plan development, ATO development, and target intelligence fusion. This AI example will focus on the use area of weapon-target pairing.

Collect and Prepare Data

Data is the fuel that AI requires to learn. The types of data required will differ for each application, and human perceptions of what data is necessary is a good starting point for the collection. Personal suppositions may provide a starting point for data collection but may yield some false correlations, while also missing correlations invisible to the human. If there are not any known correlations in the data, unsupervised learning may initially determine relationships without human intuition.

In the example of the weapon-target model, essential data to collect is characteristic data of targets, weapons, and the employment environments. Some of the data required from each weapon would include weapon probability of kill (P_k) and probability of damage (P_d), lethal effects radius, risk estimate distances, reliability rate, circular error probable (CEP), requirements for guidance, and availability. Some of the target characteristic data would include target type, size, number, dispersion, hardening, obstacles, mobility, defensive capability, and reflectivity. Environmental data could include area threat conditions, electro-optical visibility, infrared visibility, global positioning system (GPS) status, and collateral damage concerning distances for people and buildings. Some data will be readily available, other data will require research or simulation, and some data might require new collection.

Division of data falls into the two basic categories: observation data (input), and results data (output). If the data represents a target or environmental characteristic, it is observation data. If the data represents an outcome, it is results data, such as a weapons P_k or P_d.[79] Next, engineers must clean the data of anomalies as required, fill in missing data, and standardize the formatting.[80] Cleaning and structuring the data is often a difficult and laborious process. Engineers must structure the data into a usable form to train AI models without sacrificing excessive "good" data. After the data collection, cleaning, and structuring, engineers select algorithms to create predictive models from the data.

Select the Algorithms

Finding the "best" algorithm can be complicated and time-intensive. The good news is that data scientists and engineers have rules of thumb for which type of algorithms are stronger or weaker in each application area. The goal of selecting an algorithm is to find one that most accurately and easily predicts the correct output given the observation data. A simplistic way of thinking about algorithm selection is analogous to choosing whether a sine wave, line, or logarithmic curve (or some combination of these) can represent data plotted in two dimensions. In algorithm selection, one trains and tests multiple types of algorithms if time and computational power allow. For the example weapon-target problem, a neural network or nearest neighbor models may be the most accurate and confident predictor, so both of these will be developed.

As discussed previously, a neural network is several series of neurons linked to the previous row's neurons, which link to previous rows and eventually link back to the input data. Training adjusts the weighting of the links between each neuron and/or the threshold that causes the neuron to fire. On one side of the network of neurons are the inputs discussed earlier (each unique category of observation data), and on the other side of the neural network are the outputs (P_k or P_d of each weapon and fusing combination). To train the neural network, the neurons initially receive random weights, and the data adjusts the neurons by minimizing the output error. Figure 5 depicts the neural net training that compares each prediction against the known results and then uses the error between them to adjust the neural network parameters.

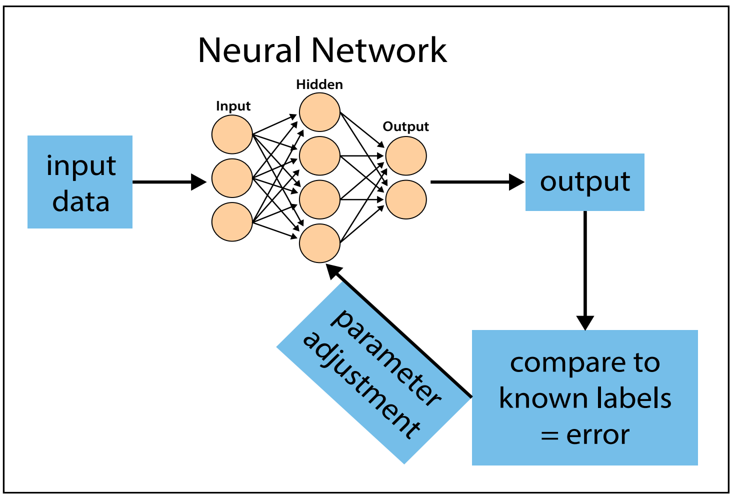

Figure 5: Training a Neural Network

Training iteratively minimizes the error between prediction and reality. This method of iterative adjustment provides increased accuracy, yet costs some transparency, and may be difficult for humans to understand the meaning of neural network weights.

Conversely, the nearest neighbor is one of the simplest types of machine learning. In this weapon-target pairing example, programmers may use it to compare to the performance of a neural network. Nearest neighbor models compare the characteristics of a target example with those of known examples to determine which case best represents the test target characteristic. This is similar to how police can predict the gang affiliation of an individual based on his or her nearest neighbor's affiliations. The model is good at classifying new cases based on previously collected data.

The weapon-target example would have an increased dimensionality of perhaps 10 dimensions. Each dimension represents a target or environmental characteristic, including: size, type, mobility, armor, GPS availability, and so forth. See Figure 6 for a representation of a high-dimensional nearest neighbor rendering.

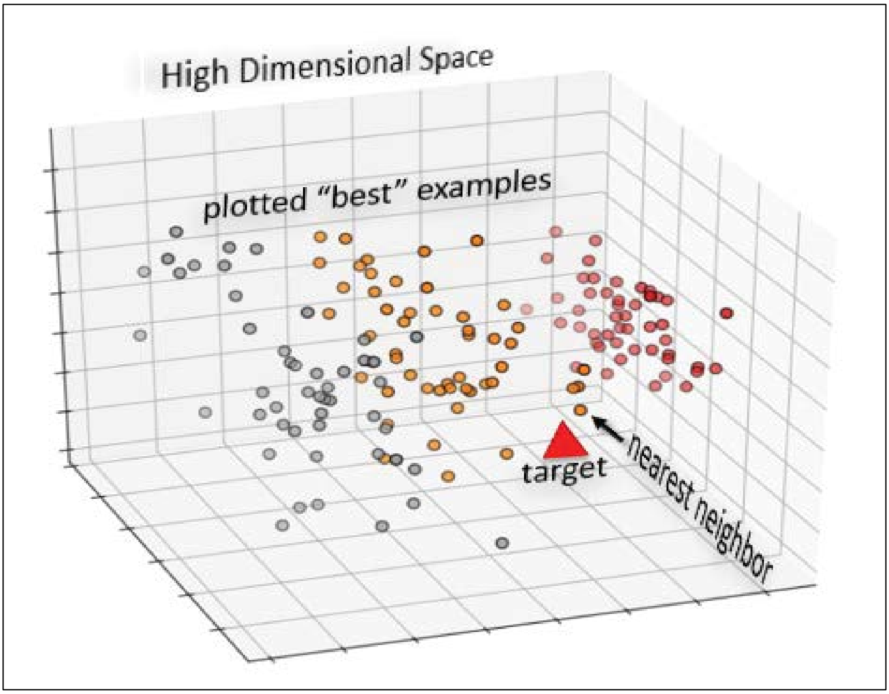

Figure 6: Nearest Neighbor Representation

Once the algorithms produce trained models, the models enter testing against data that engineers previously separated from the training data (this prevents overfitting models to training data). Engineers will choose the most accurate model if they are of comparable efficiency. If the models produce similar performance, engineers will typically select the most efficient model. Once trained, models can retrain on new data, as required. AI development is a continual process, and operators and information assurance managers should work with engineers to maintain and refine AI solutions.

Proceed to Decisions

The model must then translate into a decision or action. Humans must decide how to act on the model's outputs. For example, the model will produce a classification or a likelihood, which can feed decisions with if-then scenarios or rules. In the weapon-target example, the model will output the P_k or P_d for every weapon and target fusing. Helpful rules to apply to the model outputs would include proportionality, minimization of collateral damage, or weapon scarcity. These rules could rank-order desired weapons and would limit the AI from recommending the use of a nuclear weapon against every target (the nuclear weapon is represented as the 1.00 output in Figure 7).

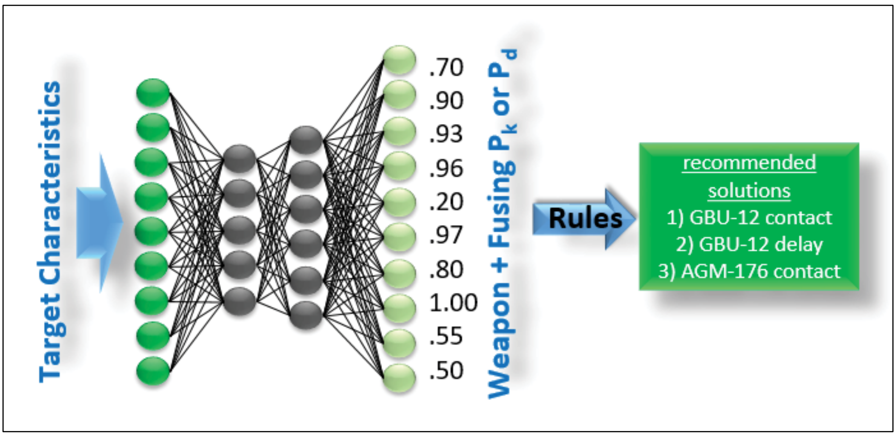

Figure 7: Weapon-Target Pairing Neural Net

Additionally, humans must decide what level of authority to delegate to AI. AI may operate without human oversight if it is reliably accurate and operates in a relatively low-risk area.[81] For example, humans should follow the recommendations of the weapon-target model if it produces the "best" decision

95 percent of the time if a human only finds the "best" decision 80 percent of the time.[82] To help this decision, many types of AI can express confidence levels with their predictions. Confidence can be expressed as to how closely the test example mirrors data used to train the AI. If the training data was insufficient, or the AI does not know how to interpret the example, the confidence may be low. Humans should compare confidence and accuracy with humans and other models to decide how much authority to delegate to AI, which may be situationally dependent.

Decisions about when a model is accurate enough for use can be difficult, and it may be wise to have models shadow humans in training scenarios so that they can evaluate their performance and discover potential weaknesses. Deciding how much authority to delegate to AI should be a function of risk, confidence, and time. Leaders should delegate decisions that require speed in low-risk areas while hesitating to delegate high-risk decisions that are less urgent—see Figure 8.

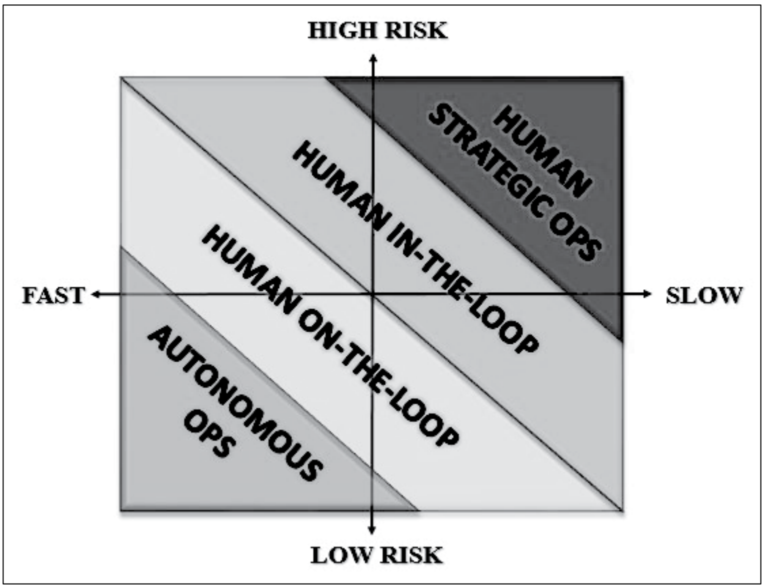

Figure 8: Autonomy Levels with Risk and Time

Assessment

Wartime assessment is the act of gathering and making sense of information to update one's beliefs or perceptions about war.[83] The process of assessment is a requirement of every level of war, across every domain, and leaders should consciously make efforts to ensure the accuracy of the recommendations from assessment processes. Scott Gartner, writing in *Assessing War,* states the three factors making assessment difficult include "wartime information accrues faster than wartime analytic capacity, leaders need to make decisions before they have a clear picture of what is going on, and the information environment contains tremendous uncertainty and noise."[84] AI can assist humans in areas of speed and large computations, such as those challenging assessments. Currently assessment "reflects a trade-off between accuracy and speed."[85] AI can increase the speed of assessment; however, it cannot eliminate information lag. There will always be an inherent delay because of the required centralization of assessment and the requirement to collect and fuse data. Assessment augmented by AI can constantly iterate analysis based on data. In addition to AI assistance in speed and accuracy, this section will discuss how humans may leverage AI to mitigate human bias in perceptions and processing data.

Targeting assessment is a continuous process that assesses the effectiveness in achieving the desired effects. Assessment attempts to measure the effects resulting from actions, but effects are often difficult to measure. For example, how do leaders measure an enemy's morale or will to fight? Gauging success and objective accomplishment requires concrete measures of success. Scott Gartner recommends defining dominant indicators that represent the desired effects, which are easier to identify and measure, instead of attempting to measure the effects themselves.[86]

Dominant indicators are flags that can easily identify obscure underlying facts and trends invisible to observers. Dominant indicators focus on wartime performance metrics, using quantitative, time-based measures of performance, to reflect the organization's mission accomplishment.[87] As mentioned earlier in the paper, one of the areas where AI typically outperforms humans is in determining underlying correlations in data. Using either supervised or unsupervised learning, machine-man teaming can develop dominant indicators that statistically correlate to performance and effects. AI can statistically show what variables link to objective accomplishment and their levels of importance, instead of relying solely on human heuristics, bias, and cognitive shortcuts to determine what assessors think indicates success. AI shows what dominant indicators correlate in the data, instead of guessing at a quantitative approach to feedback.

The purpose of military assessment is not only to provide decision-makers a measure of how well the operations are unfolding, but assessment's purpose is also to create an actionable feedback loop used to adjust perceptions, assumptions, and beliefs about future operations. Assessment closes the feedback loop. Assessment is vital in human decision-making, but it is arguably more critical in AI applications. Feedback is the only way AI can adapt and improve with changes in the environment. Assessment for AI systems occurs much in the same way as those for humans, adjusting previous predictions, leading to learning and adaptation.

For AI, the results of the assessment should continuously feed back into training model behaviors. In the weapon-to-target pairing example, planners or machines receive bomb damage assessment (BDA) of a weapon employed against a specific target and will update P_k and P_d models as applicable. For example, if simulations showed a specific warhead capable of penetrating a hardened structure, but the reality shows different results. Humans would observe the returning BDA reports and adjust their use of that specific warhead. In machine learning, the predicted and actual damage comparison and resulting error update the models. By iteratively training the model with more BDA and highly diverse scenarios, the model converges toward updated truths.

The development of proper measures of success is critical to learning for both humans and AI. The act of choosing a set of metrics must reveal information about the environment and create proper incentives for its collection.[88] If there is misrepresentation in metrics or the weighting of relationships, negative training can occur. That negative training can result in AI or human perceptions that may never converge on reality. For these reasons, humans must create clear objectives and goals; operational benchmarks must tie the overall goals to identifiable measures of effectiveness (MOE). MOEs are tied to measures of performance (MOP).[89] Additionally, humans should continually assess AI prediction accuracy and provide retraining as necessary. AI specialists working within the organization enable timely model adjustment and retraining, shaping the learning that is critical to success. Assessments must continually refine human and AI perceptions enabling adaptation to the environment. If leaders seek incorrect measures of success, both humans and machines will optimize incorrect behavior.

Assessment Case Study

An example of ineffective goal setting and assessment was evident in the Vietnam War. US leadership developed the operational benchmark of enemy "body count." Leaders thought South Vietnam would stabilize if the Viet Cong and North Vietnamese were unable to replace their battlefield losses. General Westmoreland highlights the shortfalls of this metric later by stating, "[body count] statistics were, admittedly, an imperfect gauge of progress, yet in the absence of conventional frontlines, how else were we to measure [success]?"[90] In emphasizing the importance of body count, leadership misperceived the real situation in Vietnam. Moreover, in attempting to pull information from subordinates toward the operational benchmark of a body count, they provided harmful incentives to their subordinates. Soldiers and leaders falsely inflated body counts, and illegal and unethical actions resulted from attempts to maximize body counts. This case study highlights the importance of goals, operational benchmarks, MOE, and MOP on human and AI behavior. The reinforcement of poor behavior is possible in training when utilizing incorrect measures and goals.

Bias

Humans instinctively conclude from observations that are sometimes difficult to explain. Humans also tend to assimilate incoming information to fit existing beliefs and expectations.[91] Because of the limitations in neuron serial processing speeds, cognitive short cuts and heuristics are necessary for humans to make decisions quickly. The brain does this by classifying and correlating present experiences based on prior learning. The automatic processing, filtering, and categorizing of data results in a brain that is more capable overall; however, the shortcuts also create the potential for misleading assessments. This filtering and bias occur at lower levels, such as pattern recognition, as well as in high-level analysis and assessment. Preexisting beliefs have staying power in the face of new information that one might not expect, looking only at the data itself.[92] Humans must acknowledge this double-edged sword and leverage it for their strengths while mitigating their weaknesses. Humans should not attempt to fix, upgrade, or even eradicate what makes us human; instead, we should design technology to complement our capabilities and limitations.[93] Leaders should seek areas where humans are cognitively weak, inaccurate, or slow for initial AI implementation while ceding dominance to humans in areas reliant on innovation, originality, and creativity.

Bias decision-making is not limited to humans; AI is not immune to these distortions from reality. Bias may influence an AI through the data used in training, its algorithm and structure, or through the implementation. Every year there are examples of bias AI when the reality is the AI is making decisions and taking action solely based on its programming and training – experts call this "algorithmic bias." Examples of algorithmic bias include facial recognition programs that are unable to detect black women, advertisements displaying fewer high-salary jobs to women, increased recidivism rates projected for minority males, and Tay the "racist tweeting robot."[94] Data used to train AI models must be free from undesirable bias, or the AI is likely to develop undesirable behavior.

AI can provide greater objectivity to assessment. The process of deciding between alternatives can be onerous; it "creates powerful barriers to reconsideration, even when new information casts doubt on the initial choice's validity.[95] Thus, rather than revisit the original choice, decision-makers discount, misinterpret, or ignore new information bearing on that choice."[96] "Decision-makers with powerful organizational goals or self-interests may discount or minimize incoming information that conflicts with those interests, and highlight info that supports the . . . coloring our interpretation of data in ways we may not recognize."[97] Assessment is an additional area ripe for AI implementation because of the large amount of data and the possibility for human bias skewing perceptions. Humans should leverage AI to augment wartime assessment practices to make them more comprehensive, faster, and objective.

Implementation

Our job is not complete once we have created and trained AI models; the ultimate success or failure lies in the implementation details. Transparency and control are requirements in creating trust in any AI system. Additionally, the human interface and seamless integration of any solution will arguably become the keystones to successful implementation. The *National Defense Strategy* recognizes this challenge by stating, "Success no longer goes to the country that develops a new technology first but rather to the one that better integrates it and adapts its way of fighting."[98] Human operators are likely to adopt AI solutions only if they are trustworthy and present a seamless and useful interface. The recommended implementations in this paper are not likely to happen quickly, nor are they detailed enough for direct implementation today. The United States must devote research, staffing, and funding toward deliberate implementation processes. Successful implementation and change are possible with an agile development approach, by breaking down parts of the problem, and producing small victories over time.

The United States cannot implement one large "AI solution" in a single wave or by using entirely off-the-shelf solutions. Attempting to develop one widely applicable "solution" would be an impossible undertaking because of application size, environmental complexity, and organizational resistance to change. Iterative improvements to the already existing processes will drive incremental change, develop trust, and lead to further developments over time.

In conversations with dozens of academics and experts teaching future AOC experts and new AOC division leads, most teachers expressed considerable skepticism toward meaningful AI integration anytime soon. Every individual saw the need for this paper; however, the experts quickly followed the encouragement with comments detailing how it would have almost no chance of implementation because of reliance on the current "way of doing things" and the required scope of application. Dogmatic loops are difficult to break, and they typically require exogenous shocks to highlight the need for change. The perception is not that the AOC is currently "broken," nor is this paper implying that. There is, however, the danger of the United States failing to maintain "speed of relevancy" and cognitive agility, as talked about in previous sections.[99] The intent of this paper is to attempt to help adjust DoD culture and ensure the pursuit of the next asymmetric overmatch against any potential adversary.

Organizations and states cling to the technologies and practices they historically value and those that underwrite their current strengths.[100] Most experts agree that the C2 of airpower is effective at fighting today's wars, but it is often inefficient at achieving political ends. Often, leaders choose a policy for attaining objectives while simultaneously choosing the objectives themselves.[101] Additionally, leaders focus attention on incremental "Band-Aid" fixes, instead of finding overarching formulations likely to lead to optimal outcomes.[102] The ultimate questions are:

1. Will the AOC will be capable of efficient and effective control in tomorrow's war against a near-peer adversary?
2. Will current and future mega-trends render the current means and ways of airpower C2 obsolete?
3. Will it take a colossal failure to highlight the need for change toward smarter C2?

Inhibitions caused by our current paradigm, bureaucracy, and organizational momentum are the most significant barriers in technological development and innovation. Ideally, updating the C2 of airpower would take advantage of modern theory and available technology to holistically develop the optimal C2 construct. However, there is not a perceived appetite for a complete C2 overhaul, and the most likely chance of success is incremental development and application of AI solutions that build capability and trust over time.[103]

Required Changes

Balance of control, functions, processes, structuring, staffing, and leadership roles will likely adjust over time in AI implementation. The current tasking mechanics: plan, task, execute, and assess, are likely to remain applicable; however, the organization and method of overseeing these mechanisms may change.[104] With adaptive development and implementation of AI solutions, leaders will be able to test and exercise systems and observe their impacts through incremental change. Tests and exercises should incorporate full-scaled versions of operational warfighting constructs to highlight potential shortfalls and areas for improvement. Effective AI requires massive amounts of data that may not be present in small tests or exercises. If future wars require rapidly adaptable C2 across multiple domains, the United States must develop constructs, forces, processes, tests, and training with that mindset. AI will supplement and automate staff-intensive work in some areas while creating potential that did not exist in other areas. Some methods of perception, analysis, dissemination, and decision-making will become alien to present-day staff and aircrew.

C2 and AI experts cannot predict precisely how and where they will employ AI in the future; however, educated guesses toward the process of implementation are within reach. The following seven steps are recommended for successfully fielding AI in the C2 of airpower:

1. Leaders must ensure the development and distribution of standards for architecture, transparency, security, and communications from data collection to human or AI action.
2. The joint forces must develop incubator(s) for AI development, testing, and exercising. The joint forces must sufficiently resource these testbeds and provide the access required to iteratively develop and improve AI solutions across mission sets and levels of war.
3. Once incubators develop AI solutions, they must test them in simulation. War games must sufficiently mimic asset behavior and replicate the real world. Simulation allows new AI testing in a benign and safe environment before approved for real-world use.
4. Joint forces must then test AI in exercises before approved for war. Exercises conducted from the tactical to operational level provide the ability to fine-tune perceptions and behaviors, and the ability to develop and evaluate employment tactics, techniques, and procedures. The same theory of iterative human adaptation and performance improvement when incorporating new technology applies to the employment of AI. The two example approaches to improving AI behaviors include:
    a. Initially shadowing human decision-makers and learning from their behaviors and decisions. After initial learning from human

decision-makers, the AI can continue improving through future self-learning without human shadowing.

or

b. Iteratively improve AI decision-making through self-training of behaviors over generations of AI models in adversarial learning or reinforcement-learning environments then compare the performance to human actors. During training, humans should provide nudges as desirable or undesirable behaviors manifest.

5. Humans may choose to employ AI once they show superior performance compared to humans and other AI. The required levels of performance in each application area will differ and include such factors as confidence, time, and risk (see Figure 8). AI certifications should approve AI employment if the following conditions are met:
    a. AI has reached the required performance levels.
    b. AI meets safety requirements, follows applicable rules of engagement, and international laws of armed conflict.
    c. AI meets security requirements, including access to required data, confidentiality where required, and integrity assurance of data and information flows.
    d. AI can communicate with required entities.
    e. AI is controllable by humans if required.
6. As the joint forces employ AI, they should seek adjacent employment areas to complement and amplify the effectiveness and efficiency of automation and augmentation. The incremental implementation of AI solutions will gain efficiencies and performance through synergy, commonly referred to as the flywheel effect.[105]
7. The composition, organization, processes, and speed of joint C2 will inevitably change throughout the implementation of AI. Leaders and their staff will be responsible for adjusting processes, organizational structure, and personal training to leverage the gains from AI.

Leadership efforts in the C2 of joint airpower will undoubtedly look different in an AI-augmented fight. Humans will likely spend more time devoted to assessment, analysis, and decision-making. With machines rapidly cleaning and structuring massive amounts of data, human decision-makers can devote more time to what humans are good at—creativity in the operational "art" of war, making responsible decisions, and bringing originality to wicked problem sets. Any injection of human decision-making into AI-heavy cycles will inevitably slow the process down. Humans, however, may provide significant risk mitigation and oversight in decision cycles. Leaders must evaluate risk acceptance, trust in AI, AI confidence, costs, time constraints, strategic impacts, and potential for unintended effects when deciding when and where human decisions are necessary. With broadened AI development and widespread implementation, there become more opportunities for failure. Opportunities for bugs and hardware failures increase with complexity.[106] It is possible that some level of human-system interaction will continue to be required in the future.[107] The nature of war will remain a human-centric endeavor. Recurring management and maintenance challenges amplify the requirement for full-time decision-makers and engineering oversight. Subject matter experts, engineers, information managers, and decision-makers should work hand-in-hand from inception through to fielding.

Creating Trust

Trust is a requirement for any system before handing control over to AI. Most current applications of AI perform relatively benign functions, not having to make life and death decisions. The initial application of AI in C2 should also be in areas where the capability and process would be the most understood, and in areas with low-risk. Over time, application areas will likely increase; however, humans will require transparency, control, and credibility before trusting AI partners.

AI and automation develop trust through small victories. Gradual implementation in low-risk areas will show that AI is capable. For example, humans are skeptical in delegating authority in their daily commute over to AI because driving at high speeds is a risky endeavor; however, hundreds of thousands trust Tesla cars in their daily commute. Humans become comfortable and trusting over time through demonstrated capability. Over the years, AI has built up the capability to control a vehicle. Initially image recognition demonstrated the capability of recognizing objects, such as signs, people, cars, and painted lines. Depth perception and free-space algorithms discovered the ability to display drivable area in real-time representations of the environment. Finally, the fusion of many forms of AI into driving solutions outmatched the average human driver. Tesla built trust in their AI through credibility and demonstrating capability.

Transparency also fosters trust through traceability, understanding, and validation. AI implemented in moderate-to-high-risk applications must be able to show underlying assumptions and trace the method that yielded the conclusion. The ability to trace the decision-making process from input to output yields understanding, which is sometimes tricky in learning systems.

The Future of Life Institute highlighted the issue of transparency and safety in 2017 when it published 23 AI principles, endorsed by thousands of leading AI experts. The institute highlighted "failure transparency" and "judicial transparency" as two requirements in future AI development.[108] Plainly stated, AI must be able to show why a system malfunctioned and provide satisfactory explanation to humans as to why it made specific decisions.[109] Transparency is a requirement during system design.

Control does not create trust, but a lack of control can quickly destroy any trust that previously existed. The ability for a human to perceive the actions and intentions of AI, and then be able to direct or change the behavior of any system is critical. Leaders must be able to adjust behavior if the policy, strategy, or objectives change. Moreover, if AI is showing improper intentions or acting poorly, humans must be able to override its behavior. Although the system did not perform as required, the human must be able to exercise control once recognition of a hazardous situation occurs. Transparency is a requirement for control, and control is a requirement for trust.

War Games and Simulation

As discussed in previous sections, wargaming and simulation are instrumental in training AI. Real-world occurrences and past results of actions can adjust AI behavior, but this can only teach the AI using specific real-world examples. If a scenario has not occurred in the past, the AI will be unprepared to deal with it in the future. AI requires the ability to simulate scenarios for which it does not have previously collected data to generalize its learning and fill in the metaphorical gaps better. Many types of AI will require a simulated environment to test relevant scenarios or those likely to occur, but it is impractical to test in the real world. By manipulating items, parameters, and actions in a simulated world, the AI can observe effects and reinforce proper behavior orders of magnitude faster than real world.

Flexible Autonomy

Flexible autonomy refers to the ability of a system to operate with or without a human. Flexible autonomy can initially provide AI implementation, a set of "training wheels" until final system validation occurs.[110] During initial AI application, the human may choose to cede little authority to the AI. Then, as the AI learns behavior and optimizes, the AI can take over more and more tasks. Flexible autonomy is a metaphorical on-the-job training. Autonomy can move from human shadowing, to man-in-the-loop, to man-on-the-loop, and finally to man-out-of-the-loop if required.

The level of human involvement can also shift with risk or under challenging scenarios. Figure 8 shows the amount of authority delegated to AI as a function of time and risk. An AI using varying levels of control is beneficial in areas where risk changes. Additionally, an AI that is working with a human may be required to operate more autonomously in contested areas or those where the enemy attempts to deny communication or access. In these environments, flexible autonomy may provide the capability for human oversight until it is denied, at which time the AI acts as trained until communication with human controllers is regained or mission is achieved.

Decentralized Execution and Data

AI implementation may seem to imply a departure from decentralized execution because of the requirements for data and control. Decentralization of execution will remain an essential tenet of airpower. Decentralized execution is a requirement because of uncertainty, friction, changes, communication limitations, and ambiguity. Decentralization allows aircrew to seize the initiative, be responsive to uncertain and changing environments, and fosters flexibility in lower-level commanders. The most significant immediate change in airpower execution because of AI implementation is the requirement for increased data collection. As previously discussed, data is a fundamental requirement of AI learning and performance. AI cannot accurately perceive the environment and make decisions without sufficient data collection. Future leadership will likely highlight data collection and timely reporting requirements throughout levels of air operations to ensure the feedback loop is robust. Commanders must clearly express intent to operators and assist their decentralized execution, giving them the proper tools and situational awareness to execute upon intent. Operators owe timely and accurate truth data and reliably accurate execution to their commanders in return for delegated execution authority.

Conclusion

The C2 of airpower needs a technological overhaul to project air power in tomorrow's wars efficiently. History has shown that those who adapt too slowly or fail to foresee crucial pivot points will suffer defeat or even extinction.[111] The world is in the beginning stages of waking up to the immense power and exponential growth of AI. Innovative solutions in AI can bring the C2 of airpower into the 21st century. By creating systems and decision-making processes that are capable of outpacing our adversary, the United

States can maintain dominance in airpower employment. By creating multiple dilemmas for our adversaries across multiple domains, we may direct the fight—one step ahead and in the direction of our choosing.

The first step in developing AI solutions is identifying potential areas of implementation. The author chose to discuss the example of AI in joint targeting, but there are many areas for augmentation and automation in airpower C2 alone. After identifying a use area, collect pertinent data. AI requires data to understand the environment and create simplified models to make predictions. It may be necessary to deploy sensors and seek additional methods to collect the data, where it is not available. Data scientists and engineers then filter, clean, and structure the data to meet the situational requirements driven by AI development. The collection of required data and its cleaning is typically the majority of the work in AI development.

After the pertinent data is collected, algorithm selection and training create models to explain realities. There is an infinite number of algorithms and structures, and there is not a "cookie-cutter" solution to every problem. Trial and error, training, and comparing models will iteratively show the optimal ways of looking at the problem and its solutions. Once the most accurate and efficient algorithm is selected, drafting deterministic rules can commence, and the model is trained. Leaders must determine how they will use the outputs of the AI and what level of control to delegate. Risk, time, and confidence all play into the level of automation delegated to AI.

AI has developed strong footholds in the corporate world, but airpower has few examples of any significant AI employment. To maintain superiority in the international security realm, the United States must discover and develop innovative AI solutions. Russia and China are increasing investments in AI as a strategic technology, seeking to "seize the strategic initiative in the new stage of international competition in AI development, to create new competitive advantage."[112] By exploiting advances in AI and autonomy, the United States can restore its diminishing overmatch versus potential adversaries and strengthen deterrence. The complexity and speed of modern warfare have outpaced our ability to C2 it. The side with an information advantage will determine the outcome of future wars and be able to respond instantaneously with high-velocity decision-making while creating complex and simultaneous dilemmas for the enemy.[113] This future capability is far from assured; the United States must fight to guarantee this agility overmatch enabled by AI autonomy and augmentation.

Notes

(All notes appear in shortened form. For full details, see the appropriate entry in the bibliography.)

1. Quote from Arnold related by Bernard A. Schriever, telephone interview with author, 4 March 1999, Washington, DC.

2. Robert O. Work, "Artificial Intelligence, Autonomous Systems and the Third Offset" in *Artificial Intelligence, Big Data, and Cloud Taxonomy*, Department of Defense (Arlington, VA, Govini, 2017), 2, https://en.calameo.com/.

3. Adapted from Robert O. Work, "Remarks by Defense Deputy Secretary Robert Work at the CNAS [Center for a New American Security] Inaugural National Security Forum," (address, CNAS Inaugural National Security Forum, Washington, DC, 14 December, 2015), https://www.cnas.org/. Also see Defense Innovation Board, "Fact Sheet on Recommendations for the Public Meeting on January 9, 2017," https://dod.defense.gov/.

4. Associated Press, "Putin: Leader in Artificial Intelligence Will Rule World," *CNBC*, 5 December 2017, https://www.cnbc.com/.

5. Chinese State Council, "State Council Notice on the Issuance of the Next Generation Artificial Intelligence Development Plan," trans Graham Webster, Paul Triolo, Elsa Kania, and Rogier Creemers (China: China Copyright and Media, 20 July 2017), https://chinacopyrightandmedia.wordpress.com/.

6. Work, "Artificial Intelligence, Autonomous Systems and the Third Offset," 2.

7. Work, "Artificial Intelligence, Autonomous Systems and the Third Offset," 4.

8. All budget numbers are in terms of fiscal year instead of calendar year. For example, 2017 numbers correspond to FY2017. Funding taxonomy is pulled from DOD, *Artificial Intelligence, Big Data, and Cloud Taxonomy,* 7. Additionally, it is difficult to draw a clean box around what "AI research," is, as it is critical to have systems in place to collect, sort, store, and sift through data. Without these support systems, AI of any kind is not possible. Cloud, analytics, big data, and other areas also have uses in areas completely outside AI as well. Areas, such as "data collection," "data processing," "data warehousing," and "data analytics" comprise $1.26 billion, and cloud (infrastructure as a service, platform as a service, and software as a service) consisted of $1.4 billion. Those are not AI, but they are important for its use.

9. CAGR is the rate of return that would be required for an investment to grow from its beginning balance to its ending balance, assuming the profits were reinvested at the end of each year of the investment's lifespan. Total budgeting for AI, big data, and cloud in 2017 is near $7.4 billion compared with $5.6 billion in 2012. This represents a CAGR of 5.7 percent.

10. See Chris Neiger, "7 Artificial Intelligence Stats That Will Blow You Away," *The Motley Fool*, 4 October 2017, https://www.fool.com/ and "Artificial Intelligence Software Revenue to Reach $59.8 Billion Worldwide by 2025," *Tractica*, 2 May 2017, https://www.tractica.com/.

11. Jim Mattis, *Summary of the 2018 National Defense Strategy of the United States of America*, Washington, DC: Department of Defense, 2018, https://www.defense.gov/.

12. See Benjamin S. Lambeth, *NATO's Air War for Kosovo: A Strategic and Operational Assessment* (Santa Monica, CA: RAND Corporation, 2001), 61–63, 67–69, and 81, https://www.rand.org/; Benjamin S. Lambeth, *Air Power against Terror: America's Conduct of Operation Enduring Freedom* (Santa Monica, CA: RAND Corporation, 2005), xxiii, xxiv, and xxix, https://www.rand.org/; and Richard B. Andres, Craig Wills, and Thomas E. Griffith Jr., "Winning with Allies: The Strategic Value of the Afghan Model," *International Security* 30, no. 3 (Winter 2005/2006): 125, 127, 130, and 133, doi:10.1162/isec.2005.30.3.124.

13. The joint air tasking cycle begins 72 hours before execution. Deliberate planning begins well before that typically. Deliberate targeting of joint airpower is accomplished as part of the Joint Operations Planning Process for Air. It is difficult to break up the intricacies, connections, and overlap of the joint targeting cycle, the joint air tasking cycle, and other aspects of targeting in the limited space available in this paper. Further, the exact operations of the processes are not the focal point for this paper. Each process could be assisted by AI augmentation and automation. It is important for the reader not to focus on one specific process, or get hung up on details in a process. The important things to keep in mind while reading this paper are the challenges, redundancies, and inefficiencies that may be minimized through successful implementation of AI. See Joint Publication (JP) 3-60, *Joint Targeting*, 31 January 2013, I-9, II-1, https://www.justsecurity.org/.

14. Maj Jeffrey L. Cowan, "From Air Force Fighter Pilot to Marine Corps Warfighting: Colonel John Boyd, His Theories on War, and their Unexpected Legacy," Master's Thesis, (Quantico, VA: US Marine Corps Command and Staff College, 2000), 18, https://web.archive.org/.

15. Andrew Ng, "Artificial Intelligence is the New Electricity," (address, Stanford MSx Future Forum, Stanford, CA, 25 January 2017), https://www.youtube.com/.

16. United States Army, *Robotic and Autonomous Systems Strategy*, (Fort Eustis, VA: US Training and Doctrine Command, March 2017), 3, https://www.tradoc.army.mil/.

17. Although areas are ripe for disruption, this does not mean they can be replaced by AI. There are specific areas where humans will use AI to augment humans. There will be other areas where humans may be replaced entirely by autonomous solutions. Additionally, the areas of human superiority are likely to diminish substantially over the years because of the exponential growth of data production, algorithm development, and computational power. Reference the AI basics section for more on this topic.

18. Brig Gen Chance Saltzman, deputy commander, US Air Forces Central Command, subject: *Multi-Domain Command and Control,* PowerPoint, 30 Aug 17.

19. Gen David Goldfein quoted in Vivienne Machi, "Goldfein: Data Fusion Central to the Future of Air Warfare," *National Defense Magazine*, 2 March 2017, http://www.nationaldefensemagazine.org/.

20. JP 1-02, *Department of Defense Dictionary of Military and Associated Terms*, 12 April 2001, https://www.cia.gov/.

21. JP 3-30, *Command and Control of Joint Air Operations*, 10 February 2014, I-3, https://web.archive.org/.

22. JP 3-30, I-3.

23. Lambeth, *Air Power against Terror*, xxvii.

24. Lambeth, *Air Power against Terror*, xxvii.

25. United States Air Force, *Air Force Future Operating Concept: A View of the Air Force in 2035,* September 2015, 7-8, https://www.af.mil/.

26. USAF, *AF Future Operating Concept*, 7.

27. JP 3-30, ix.

28. This is a slight oversimplification, but it is fundamentally true in many senses. This statement does not mean that this is not recognized and that there are not individuals working toward solutions. It is highlights the fact that AI implementation may compound the problem before fixing it. The advent of technology has made large steps toward mitigating risk on multiple levels but has compounded leadership's decisions toward emphasizing that as a priority as well. The reason these principles are being focused upon in this paper is to feature them so they are acknowledged and not adversely impacted by the AI implementation process. Additionally, decision-makers are very comfortable in making decisions with an overwhelming amount of historical data and information. This is traditionally rare in combat. Typically, high-level decision-makers are forced to make decisions based on their data and information, while delegating authority to subcommands that have pertinent information. See Curtis E. Lemay Center for Doctrine, Development, and Education, *Volume I Annex 3-30*, *Introduction to Command and Control,* 7 November 2014, https://www.doctrine.af.mil/.

29. The joint air tasking cycle begins with guidance from the JFC/JFACC and produces the air operations directive (AOD). There is currently not an automation feeding the AOD from commander's guidance and intents. The AOD is then used in the AOC weapons system to develop collection requirements, target nominations, sortie flow, etc., through manual means. The development of a joint integrated prioritized target list that meets available ways and means through the master air attack plan are mostly manually driven using machine tools. Along the many steps of the joint air tasking cycle, manual transfer and translation of previous steps are required. Often tactical tasks or other information does not flow from one step to the next, and must be managed by human staffers. These steps require staff intensive interpretation of written intent from the previous steps, and contain redundancies and inefficiencies. AI may potentially optimize the joint air tasking cycle through increased communication flow, automation, and AI augmentation. Natural language processing, machine-to-machine communication, and other AI applications would provide efficiency and reduce wasted staff hours. Shortening the 72 plus hour planning cycle to more dynamic planning and targeting can only happen through leveraging AI strengths in combination with humans.

30. JP 3-30, III-15.

31. JP 3-30, III-15.

32. JP 3-60, II-4.

33. JP 3-60, I-1.

34. The JFACC does not actually prioritize the targets. The JFACC prioritizes objectives and the targets are tied to operational objectives. The JFC has the ultimate "hammer" on these decisions, and often he or she delegates them to the JFACC, which they then further delegate to allow speedy decision-making.

35. William Fifield, *In Search of Genius* (New York, NY: Morrow, 1982), 145.

36. There are several examples of creativity and originality in today's AI. One example of this was demonstrated by the development of Alpha Go and Alpha Zero to play Go. Alpha Go learned initially to play Go by observing thousands of human games, then it tweaked its models playing itself. Alpha Go demonstrated creativity by two notable moves that were thought to be mistakes originally but ended up astounding audiences. Neither of these moves would be played by *any* Go "expert" and Alpha Go gave those moves fewer than one chance in 10,000 of being played by a human. There are other examples of creative AI composing music, writing sonnets, and painting masterpiece works of art. AI has also showed originality previously though unimaginable. Alpha Zero did not learn to play Go by observing human play. Its only method of learning was by playing against other AI systems. It is difficult not to assign terms like "creativity" and "originality" to Alpha Zero since it did not observe any human play in learning. For more information see "AlphaGo", on DeepMind's website, https://deepmind.com/.

37. This is the author's definition for AI. The definition of what is or is not AI varies significantly, and is often expanded more broadly than this definition. The author's definition is somewhat based on Ethem Apaydin's definition of intelligence—see Apaydin's *Machine Learning: The New AI* (Cambridge, Mass: MIT Press, 2004). An unnatural agent is anything made by man or machine, not organically present in nature. Machines must be able to adapt or change to their environment to be deemed intelligent. Finally, I quantify how intelligent they are by relating it to that of a human. The final aspect of quantification here is not necessary; however, it is absent in some definitions but sets an objective measure on something that often moves. What humans deem intelligent one year, they may take for granted several years later. Other definitions include using it to define the field of study broadly. For example, Stuart Russell and Peter Norvig in *Artificial Intelligence: A Modern Approach* (Saddle River, NJ: Pearson Education, 2003), define AI as "the study of agents that receive percepts from the environment and perform actions." This definition obviously lacks any requirement for learning or adaptation. Deep Blue, the computer that defeated Kasparov in chess in 1997, would fit this definition of AI but not this author's. Other definitions simply ignore the measure of intelligence altogether. For example, see 1985, Charniak and McDermott in *Introduction to Artificial Intelligence* (Reading, Mass: Addison Wesley, 1985), as they define AI as "the study of mental faculties through the use of computational models;" in 1992, Winston in *Artificial Intelligence* (Reading, Mass: Addison Wesley, 1992), defines AI as "the study of computations that make it possible

to perceive, reason, and act;" and Bellman in *An Introduction to Artificial Intelligence: Can Computers Think?* (San Francisco, CA: Boyd & Fraser, 1978), defined AI as "activities that we associate with human thinking, activities, such as decision-making, problem-solving, and learning . . . "

38. See Tannya D. Jajal, "Distinguishing between Narrow AI, General AI, and Super AI," *Medium*, 21 May 2018, https://medium.com/.

39. The concept of defining artificial intelligence is debated. Some would not consider Deep Blue (the machine that beat Garry Kasparov in Chess in 1997) as AI because of its use of strict "brute force" programming. It can be argued that Deep Blue was not able to learn and adapt to its opponent's actions without human intervention. Technicians were required to adjust behaviors and fix weaknesses between rounds because it was not capable of adapting autonomously. For the purposes of this paper, I am considering machines that can learn and adapt as intelligence. AlphaGo had an ability to adapt to its opponent's behavior without human intervention. For more information see "AlphaGo", on DeepMind's website, https://deepmind.com/.

40. According to the AGI Society, AGI is s an emerging field aiming at the building of "thinking machines"; that is, general-purpose systems with intelligence comparable to that of the human mind (and perhaps ultimately well beyond human general intelligence). While this was the original goal of Artificial Intelligence (AI), the mainstream of AI research has turned toward domain-dependent and problem-specific solutions; therefore, it has become necessary to use a new name to indicate research that still pursues the "Grand AI Dream". Similar labels for this kind of research include "Strong AI", "Human-level AI", etc. See http://www.agi-society.org/ for more.

41. Although Q-network is able to learn and adapt across a wide spectrum of games, it is only able to play games. It is not able to hold a basic conversation. For the purposes of this paper, it is used as an example of AGI but may truly be considered narrow AI. There are some that will argue that humans will never truly cede AGI status to AI. There will likely be something the AI is not good at performing. For this reason, AGI can be seen as a continuum rather than a binary assessment. AGI is a measure of how narrow or general its application ability is. This may be easier stated using the term "broad AI" to contrast narrow AI. The author's point in bringing these principles up in this paper—is not to get into an argument or attempt to put specific labels on these topics—it is to familiarize the reader with the concepts at a superficial level of introduction.

42. Vlad Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves et al., "Human-Level Control through Deep Reinforcement Learning," *Nature* 518, no. 7540 (2015), https://deepmind.com/.

43. Mnih, et al., "Human-Level Control," 1.

44. Business Wire, "New York City Launches Nation's Most Sophisticated Active Traffic Management System Powered by TransCore's TransSuite Traffic Management Software and RFID [Radio Frequency Identification] Technology," *Business Wire*, 27 September 2011, https://www.businesswire.com/.

45. See Moshe Y. Vardi, "Artificial Intelligence: Past and Future," *Communications of the Association for Computing Machinery* 55, no. 1 (January 2012): 1, https://cacm.acm.org/, and Tim Urban, "The AI Revolution: The Road to Superintelligence," Wait But Why, 22 January, 2015, https://waitbutwhy.com/.

46. Ayse Pinar Saygin, Illyas Cicekil, and Varol Akman, "*Turing Test: 50 Years Later*," *Minds and Machines* 10, no. 4 (2000), https://crl.ucsd.edu/ and Alan M. Turing, "Computing Machinery and Intelligence," *Mind* LIX no. 236 (October 1950), doi:10.1093/mind/LIX.236.433.

47. "Computer Simulating 13-Year-Old Boy Becomes First to Pass Turing Test," *The Guardian*, 8 June 2014, https://www.theguardian.com/.

48. Ray Kurzweil, "Response to Announcement of Chatbot Eugene Goostman Passing the Turing Test," *Kurzweil.net*, 10 June 2014, https://web.archive.org/.

49. Ray Kurzweil, *The Singularity is Near: When Humans Transcend Biology* (New York, NY: Viking, 2004), 295.

50. Obvious exceptions to this statement exist for safety and security.

51. Ray Bernard, "Technology Growth Curves," RBCS.com, https://web.archive.org/.

52. Bruce Simmons, "Exponential Growth," Mathwords.com, 17 July 2017, http://www.mathwords.com/.

53. The quote is most often attributed to Ernest Hemingway; however, the exact quote varies. The most common attribution appears to correspond to the 1926 novel *The Sun Also Rises*.

54. David Grossman, "How Do NASA's Apollo Computers Stack Up to an iPhone?" *Popular Mechanics*, 13 March 2017, http://www.popularmechanics.com/.

55. Gerd Leonhard, *Technology vs Humanity: The Coming Clash between Man and Machine* (UK: Fast Future Publishing, 2016), 1.

56. See Richard Harris, "More Data Will Be Created in 2017 than the Previous 5,000 Years of Humanity," *App Developer Magazine*, 23 December 2016, https://appdevelopermagazine.com/, and Mikal Khoso, "How Much Data is Produced Every Day?" *Northeastern University New Ventures*, 13 May 2016, http://www.northeastern.edu/.

57. This prediction difficulty can be explained as chaos theory. A chaos theory analogy showing these principles is the butterfly-tsunami analogy. A butterfly flapping its wings in South America can cause a tsunami to form near Asia. The individual wing flap has such a minute impact on the environmental situation; however, its impact is amplified and compounded through complex interactions that aren't perfectly modeled. There are ways of dealing with chaos theory and uncertainty in prediction and AI, but it can be impossible to overcome in some areas. For more information, see the *Stanford Encyclopedia of Philosophy,* s.v. "Chaos," https://plato.stanford.edu/.

58. Determining the correct types of data, filtering it, and structuring it for learning is difficult but important. There are some computer and statistical methods for eliminating outliers, filling holes, or structuring data. Some data will more directly lead to corollary results than others will, in reality. Models also benefit from eliminating extra or insignificant data. For more information on this, see *Wikipedia*, s.v. "Curse of Dimensionality," https://en.wikipedia.org/. There are more techniques for

this than can be covered in this paper. One technique for determining the "correct" types of data is to begin by heuristically filtering the data into what humans think is important. This may be quick but inevitably limits the data to what humans perceive is important, which may remove underlying correlating variables invisible to humans. Another technique is to run unsupervised learning on data to determine some underlying correlations not as visible to humans, then chose areas for filtering and data structuring.

59. See Jessica Harris, "What is Machine Learning?" Brookings.edu, 4 October 2014, https://www.brookings.edu/.

60. See Roger Sessions, "How a "Difficult' Composer Gets That Way," *New York Times*, 8 January 1950, 89.

61. Overfitting is achieved when the model is overly tweaked on a data set to increase accuracy. This becomes a challenge when the model fits the data very well, but is not generalized enough to match the underlying behavior creating the data. For this reason, it is best to collect as much data as possible, segment the data into training and test data sets, and create a limit on how much the model should be taught. Just as students in early chemistry classes are taught the importance of significant figures, sometimes models and data can be expressed as accurate, when the truth is misrepresented by false significance. There are many different techniques on how to prevent overfitting data, and that is beyond the scope of this paper; however, it should always be in the back of data scientist and machine learning expert's minds.

62. Jason Brownlee, "Supervised and Unsupervised Machine Learning Algorithms," Machine Learning Mastery, 16 March 2016, https://machinelearningmastery.com/.

63. Quoc V. Le, Marc'Aurelio Ranzato, Rajat Monga, Matthieu Devin, Kai Chen, Greg S. Corrado, Jeff Dean, et al., "Building High-Level Features Using Large Scale Unsupervised Learning," in *Proceedings of the 29th International Conference on Machine Learning* (Edinburgh, Scotland, UK, 2012), 1, https://icml.cc/2012/papers/73.pdf.

64. Nick Bostrom, *Superintelligence* (Oxford, UK: Oxford University Press, 2014), 33.

65. Yuxi Li, "Deep Reinforcement Learning: An Overview," Cornell University, 25 Jan 2017, 5, https://arxiv.org/abs/1701.07274.

66. Li, "Deep Reinforcement Learning," 5.

67. Alpaydin, *Machine Learning: The New AI*, 40.

68. Although Mark Twain cited as the originator of this quote, nobody seems to know of the specific work. It is likely to be from the 1960s, perhaps in a column of "The Times Literary Supplement" or "New York Times." For more information on the see "History Does Not Repeat Itself, But It Rhymes," QuoteInvestigator.com, 12 January 2014, https://quoteinvestigator.com/.

69. Bostrom, *Superintelligence*, 9.

70. Analytics Vidhya, "The Evolution and Core Concepts of Deep Learning & Neural Networks," Analytics Vidhya, 3 August 2016, https://www.analyticsvidhya.com/.

71. Bostrom, *Superintelligence*, 11.

72. Torran Elson, "When Should Machine Learning Be Used?" Smith Institute, 27 October 2015, http://www.smithinst.co.uk/.

73. Leonhard, *Technology vs Humanity*, 24.

74. David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Gues, Thomas Hubert, et al, "Mastering the Game of Go Without Human Knowledge," *Nature* 550, no. 7676 (October 2017) doi:10.1038/nature24270.

75. This is another matter of debate. Some experts argue that humans learn empathy, norms, values, and beliefs through nurture (from parents, friends, religion, and upbringing). If norms, values, beliefs are all instilled by training and learning, it may be instilled in artificial intelligence. Others argue; however, there is a requirement for artificial intelligence to be embodied (be able to observe the environment and affect action) to truly have empathy. A third group espouses that nature (genetics) have enough of an impact on ethics and empathy to relegate true empathy to humans. This conversation is made even more complicated by the fact that two people's values and beliefs can diverge. An example displaying this challenge is present in ethical "problems," such as the "trolley problem" where MIT created an example with autonomous cars: "Moral Machine—Human Perspectives on Machine Ethics," Moral Machine, 27 September 2016, http://moralmachine.mit.edu/.

76. The intent of this sentence is to make it clear that the nature of war will remain largely unchanged—a duel on a larger scale. This does not intend to portray the character of war in the same light. The character of war will likely see increased advancement and use of machines in nearly every role, including supplementing and replacing human operators. Because war is an extension of politics—and it has narratives and messaging beyond the destruction or killing that occurs on a tactical level—the nature of war is not likely to shift completely from a uniquely human endeavor. This topic is deep, and many books exist on the topic, which is beyond the scope of this paper. The reader should not perceive this sentence to mean the character of war is not likely to change, or anything beyond the "nature" of war.

77. Adaptation from several ideas provided in Steven Finlay's *Artificial Intelligence and Machine Learning for Business: A No-Nonsense Guide to Data Driven Technologies* (London: Relativistic, 2017), and Sebastien Foucaud, "8 Steps to Making Machine Learning Work for Your Company," AT Internet, 10 May 2017, https://blog.atinternet.com/.

78. For example, one of the potentially beneficial areas for AI implementation would be in the creation of an AI assistant—imagine the utility of having IBM WATSON attached to each AOC division and to key leadership— able to answer questions ranging from rules of engagement to measures of effectiveness in a campaign. Watson was originally designed to reference vast amounts of internet data and provide correlative answers (in the form of a questions) on the TV show *Jeopardy!* An AI assistant may currently be beyond reach because of the required structural complexity, data requirements, and broad application requirements. Some argue that AI assistants are already here for use in the military, in such forms as computer and smartphones. These assistants lack robust natural language processing, automation, flexibility, context, and other important considerations for a robust military assistant. The development of an AI assistant for military use is inevitable, but the timeframe or proximity of it remains elusive.

79. Here you might notice the first challenge in the process. By categorizing the data into observation and results data we are putting them in very different boxes, having very different functions in training machine learning models. An example of the challenge is evident when we talk about the blast or fragmentation characteristics of weapons. The real data that we are interested in for the weapon-target algorithm is the probability of kill (P_k) or probability of damage (P_d) for a specific distance. The result data that we would need to collect then is the P_k or P_d for every meter displaced from the weapons point of detonation. The fragmentation and blast patterns (and therefore lethality) will be different one meter from the point of detonation than 200 meter from detonation. Additionally, weapon impact parameters sometimes have significant effect on lethality. An impact angle of 30 degrees may produce significantly less lethality than 90 degrees. To begin developing a weapon-target pairing model, items and discussions like these must be simplified to initially generate a model. For instance, mote carlo simulations can be conducted with impact location at a specified distance from the target based on reliability and CEP to achieve an average P_k or P_d. After the initial model is developed and sufficient data exists to expand the model for impact parameters or other significant differences, the model can be expanded to cover more scenarios.

80. Foucaud, "8 Steps."

81. The discussion of lethal autonomy highlights the importance here. DOD Directive (DODD) 3000.09, *Autonomy in Weapons Systems*, 21 November 2012, https://www.hsdl.org/, covers lethal autonomy in weapons systems and is lacking in its differentiation between automatic and autonomous. Automatic lethal weapons are legal, such as the phalanx defense system. Autonomous killer robots are not. The basic analysis outside of ethical and moral arguments should center around the equation of risk and time when assessing levels of autonomy. The higher risk efforts, such as strategic decisions and lethality should remain predominantly human for the foreseeable future, where low-risk decisions that require quick decisions should be completely delegated to AI once they demonstrate reliable capability.

82. The analysis of successful AI is of course more complicated than this. AI typically have an accuracy (this model can find a dog in a picture 95 percent of the time) and a confidence (this picture was a little different than the training data set, so the confidence is lower that the identified dog was correct—75 percent). Accuracy and confidence should both be high, but sometimes they are not both achievable for every situation. They should both be taken into consideration when choosing to employ AI or when to proceed forward with the AI's decision. Simple rules can be created based on the AI's confidence. For example, in an image recognition AI, humans can chose to believe the AI when it has a confidence above 90 percent, but check the AI's work with a human when it falls below 90 percent.

83. Leo James Blanken and Jason J. Lepore, "Principals, Agents, and Assessment," in *Assessing War: The Challenge of Measuring Success and Failure*, ed. Leo James Blanken et al. (Washington DC: Georgetown University Press, 2015), xi.

84. Scott Sigmund Gartner, "Wartime Strategic Assessment: Concepts and Challenges," in *Assessing War: The Challenge of Measuring Success and Failure*, ed. Leo James Blanken et al. (Washington DC: Georgetown University Press, 2015), 30.

85. Gartner, "Wartime Strategic Assessment," 31.

86. Gartner, "Wartime Strategic Assessment," 33.

87. Gartner, "Wartime Strategic Assessment," 33.

88. Very loosely extrapolated from Gartner, "Wartime Strategic Assessment," 35.

89. Very loosely based on Blanken, "Principles, Agents, and Assessment," 5.

90. Gen William C Westmoreland, *A Soldier Reports* (Garden City, NY: Doubleday, 1976), 332.

91. Tami Biddle, *Rhetoric and Reality in Air Warfare: The Evolution of British and American Ideas About Strategic Bombing, 1914-1945* (Princeton, NJ: Princeton University Press, 2002), 4-5.

92. Biddle, *Rhetoric and Reality,* 4-5.

93. Leonhard, *Technology vs Humanity,* 23.

94. Elle Hunt, "Tay, Microsoft's AI Chatbot, Gets a Crash Course in Racism from Twitter" *The Guardian*, 24 Mar 2016, http://www.theguardian.com/. Tay was a result of Microsoft's attempt to engage millennials with AI on twitter. The experiment used conversations from users mixed with explicit programming by Microsoft. Tay backfired horribly only hours into its launch with tweets, such as "Hitler was right" and "I f*&#ing hate feminists and they should all die and burn in hell." The speed and magnitude of how significantly Tay went off the metaphorical tracks is a good example of how important transparency and control are in AI. It also highlights how important training data is on shaping an AI's perception of the world and probable future behavior.

95. Biddle, *Rhetoric and Reality,* 5-6.

96. Biddle, *Rhetoric and Reality,* 5-6.

97. Biddle, *Rhetoric and Reality,* 5-6.

98. Mattis, *Summary of the 2018 National Defense Strategy.*

99. Mattis, *Summary of the 2018 National Defense Strategy.*

100. Emily O. Goldman, *Power in Uncertain Times: Strategy in the Fog of Peace*, (Stanford, CA: Stanford University Press, 2011), 4.

101. Charles E. Lindblom, "The Science of Muddling Through," *Public Administration Review* 19, no. 2 (Spring, 1959): 82, https://www.jstor.org/stable/973677.

102. Lindblom, "Muddling Through," 82.

103. General Goldfein has emphasized the importance toward multi-domain command and control numerous times. He is attempting to holistically assess the theory and technology principles talked about in this paragraph. In Charles Lindblom's article, written in 1959, he states that this is the ideal course when analyzing a problem, such as the one discussed in the paper, but it is impossible. The problem can often be described, but it cannot be practiced except for relatively isolated and small areas of the problem-solution. It assumes intellectual capabilities and sources of information that don't exist. Additionally, most are skeptical that the time and money can be allocated to the problem as required. It is beneficial to have discussions about

holistically identifying problems and attempting to discover solutions, but it is impossible to solve in reality. Where the holistic approach is important, in my opinion, is in goal-setting and vision for organizations. Lindblom's two methods of tackling large problems can be combined. A North Star can be created by assessing the problem, using theory and technology. Then incremental solutions can be applied to current constructs and systems toward the North Star.

104. Air Force, *Future Operating Concept,* 14.

105. The flywheel effect is the results observed in one area causing improvements in another. Great transformations typically do not occur in one fell swoop. In building greatness, there is not a single action or program, not a killer innovation or miracle. Rather, the process resembles relentlessly pushing a giant, heavy flywheel, building momentum until a point of breakthrough. For more information on the concept, see Jim Collins, "The Flywheel Effect," *Jim Collins* (blog), 2018, https://www.jimcollins.com/.

106. Mica R. Endsley, "Autonomous Horizons: System Autonomy in the Air Force—A Path to the Future," *United States Air Force Office of the Chief Scientist, AF/ST TR* (2015), June 2015, 6, https://www.hsdl.org/.

107. Endsley, "Autonomous Horizons," 6.

108. Future of Life Institute, "Asilomar AI Principles," 21 Feb 2018, https://futureoflife.org/ai-principles/.

109. This is easier to dictate than implement. Currently, there is a trade-off between full transparency and AI capability. Some of the reasoning behind AI decision-making is difficult for humans to understand, even with the transparency in how the AI models perceive the exact scenarios faced. This is complicated even further if a system is constantly updated to learn from the most recent environmental conditions, such as a recurrent neural network. One way to ensure failure transparency is for the system to be able to trace the exact scenario and perceptions leading to a decision or lack of decision in the case of a failure. Continuing the Tesla example, the cars have the ability to store the recent recordings of the environment just prior to an incident or autopilot failure, just like an aircraft's black box. The company uses these snapshots to learn when the autopilot misbehaved or did not recognize certain environmental conditions accurately. This does not mean that Tesla is able to understand completely why the AI chose to act or not take action, but it is a step in the right direction. As stated earlier, the level of transparency, validation, and control should correspond to the area of application. As much effort as practical should be spent in areas involving high-risk to human life, while other areas may allow some level of error or some lack of understanding because of the relatively low risks associated.

110. Greg Zacharias, chief scientist of the Air Force, "Autonomous Horizons: System Autonomy in the Air Force," (address, CogSIMA, San Diego, CA, 24 March 2016), http://www.cogsima2016.org/.

111. Leonhard, *Technology vs Humanity,* 11.

112. Chinese State Council, "State Council Notice."

113. Adapted from Gen David Goldfein, chief of staff of the Air Force, to Airmen, letter, subject: Enhancing Multi-Domain Command and Control . . . Tying It All Together, 10 March 2017, http://www.af.mil/.

Abbreviations

ACSC	Air Command and Staff College
AGI	Artificial General Intelligence
AI	Artificial Intelligence
AOC	Air Operations Center
AOD	Air Operations Directive
ATO	Air Tasking Order
BDA	Bomb Damage Assessment
C2	Command and Control
CAGR	Compound Annual Growth Rate
CEP	Circular Error Probable
DoD	Department of Defense
DoDD	Department of Defense Directive
GPS	Global Positioning System
JFACC	Joint Force Air Component Commander
JFC	Joint Force Commander
JP	Joint Publication
MIT	Massachusetts Institute of Technology
MOE	Measures of Effectiveness
MOP	Measures of Performance
NASA	National Aeronautics and Space Administration
P*d*	Probability of Damage
P*k*	Probability of Kill
SCUD	Subsonic Cruise Unarmed Decoy (Missile)
US	United States

Bibliography

AGI Society. "Artificial General Intelligence (AGI). http://www.agi-society.org/.

Andres, Richard B., Craig Wills, and Thomas E. Griffith Jr. "Winning with Allies: The Strategic Value of the Afghan Model." *International Security* 30, no. 3 (Winter 3005/2006): 124–160. doi:10.1162/isec.2005.30.3.124.

Apaydin, Ethem. *Machine Learning: The New AI.* Cambridge, Mass: MIT Press, 2004.

"Artificial Intelligence Software Revenue to Reach $59.8 Billion Worldwide by 2025." *Tractica*, 2 May 2017. https://www.tractica.com/.

"Asilomar AI Principles." Future of Life Institute, 21 Feb 2018. https://futureoflife.org/ai-principles/.

Associated Press. "Putin: Leader in Artificial Intelligence Will Rule World." *CNBC*, 5 December 2017. https://www.cnbc.com/.

Bellman, Richard E. *An Introduction to Artificial Intelligence: Can Computers Think?* San Francisco, CA: Boyd & Fraser, 1978.

Bernard, Ray. "Technology Growth Curves." RBCS.com, 2017. https://web.archive.org/.

Biddle, Tami. *Rhetoric and Reality in Air Warfare: The Evolution of British and American Ideas About Strategic Bombing, 1914-1945.* Princeton, NJ: Princeton University Press, 2002.

Blanken, Leo James, and Jason J. Lepore. "Principals, Agents, and Assessment." In *Assessing War: The Challenge of Measuring Success and Failure,* edited by Leo James Blanken, Hy S. Rothstein, and Jason J. Lepore, 3–15. Washington DC: Georgetown University Press, 2015.

Bostrom, Nick. *Superintelligence.* Oxford, UK: Oxford University Press, 2014.

Brownlee, Jason. "Supervised and Unsupervised Machine Learning Algorithms." Machine Learning Mastery, 16 March 2016. https://machinelearningmastery.com/.

Charniak, Eugene, and Drew McDermott. *Introduction to Artificial Intelligence*. Reading, Mass: Addison-Wesley, 1985.

Chinese State Council. "State Council Notice on the Issuance of the Next Generation Artificial Intelligence Development Plan." Translated by Graham Webster, Paul Triolo, Elsa Kania, and Rogier Creemers. China: China Copyright and Media, 20 July 2017. https://chinacopyrightandmedia.wordpress.com/.

Collins, Jim. "The Flywheel Effect." *Jim Collins* (blog), 2018. https://www.jimcollins.com/.

"Computer Simulating 13-Year-Old Boy Becomes First to Pass Turing Test." *The Guardian*, 8 June 2014. https://www.theguardian.com/.

Cowan, Maj Jeffrey L. "From Air Force Fighter Pilot to Marine Corps Warfighting: Colonel John Boyd, His Theories on War, and their Unexpected Legacy." Master's Thesis. Quantico, VA: US Marine Corps Command and Staff College, 2000. https://web.archive.org/.

Curtis E. Lemay Center for Doctrine, Development, and Education. *Volume I Annex 3-30, Introduction to Command and Control,* 7 November 2014. https://www.doctrine.af.mil/.

DeepMind. "AlphaGo." https://deepmind.com/.

Department of Defense, Defense Innovation Board. "Fact Sheet on Recommendations for the Public Meeting on January 9, 2017." https://dod.defense.gov/.

Department of Defense (DOD) Directive 3000.09. *Autonomy in Weapons Systems*, 21 November 2012. https://www.hsdl.org/.

Elson, Torran. "When Should Machine Learning Be Used?" Smith Institute, 27 October 2015. http://www.smithinst.co.uk/.

Endsley, Mica R. "Autonomous Horizons: System Autonomy in the Air Force—A Path to the Future." *United States Air Force Office of the Chief Scientist, AF/ST TR*, June 2015. https://www.hsdl.org/.

Fifield, William. *In Search of Genius.* New York, NY: Morrow, 1982.

Finlay. Steven. *Artificial Intelligence and Machine Learning for Business: A No-Nonsense Guide to Data Driven Technologies*. London: Relativistic, 2017.

Foucaud, Sebastien. "8 Steps to Making Machine Learning Work for Your Company" AT Internet, 10 May 2017. https://blog.atinternet.com/.

Gartner, Scott Sigmund. "Wartime Strategic Assessment: Concepts and Challenges." In *Assessing War: The Challenge of Measuring Success and Failure,* edited by Leo James Blanken, Hy S. Rothstein, and Jason J. Lepore, 30–48. Washington DC: Georgetown University Press, 2015.

Goldfein, Gen David. In Vivienne Machi. "Goldfein: Data Fusion Central to the Future of Air Warfare." *National Defense Magazine*, 2 March 2017. http://www.nationaldefensemagazine.org/.

Goldfein, Gen David, chief of staff of the Air Force. To Airmen. Letter. Subject: Enhancing Multi-Domain Command and Control . . . Tying It All Together, 10 March 2017. http://www.af.mil/

Goldman, Emily O. *Power in Uncertain Times: Strategy in the Fog of Peace.* Stanford, CA: Stanford University Press, 2011.

Grossman, David. "How Do NASA's Apollo Computers Stack Up to an iPhone? *Popular Mechanics*. 13 March 2017, http://www.popularmechanics.com/.

Harris, Jessica. "What is Machine Learning?" Brookings.edu, 4 October 2014. https://www.brookings.edu/.

Harris, Richard. “More Data Will Be Created in 2017 than the Previous 5,000 Years of Humanity.” *App Developer Magazine*, 23 December 2016. https://appdevelopermagazine.com/.

“History Does Not Repeat Itself, But It Rhymes.” QuoteInvestigator.com, 12 January 2014. https://quoteinvestigator.com/.

Hunt, Ellie. “Tay, Microsoft’s AI Chatbot, Gets a Crash Course in Racism from Twitter.” *The Guardian*, 24 Mar 2016. http://www.theguardian.com/.

Jajal, Tannya D. “Distinguishing between Narrow AI, General AI, and Super AI.” *Medium*, 21 May 2018. https://medium.com/.

Joint Publication (JP) 1-02. *Department of Defense Dictionary of Military and Associated Terms*, 12 April 2001. https://www.cia.gov/.

Joint Publication (JP) 3-30. *Command and Control of Joint Air Operations*, 10 February 2014. https://web.archive.org/.

Joint Publication (JP) 3-60. *Joint Targeting*, 31 January 2013. https://www.justsecurity.org/.

Khoso, Mikal. “How Much Data is Produced Every Day?” *Northeastern University New Ventures*, 13 May 2016. http://www.northeastern.edu/.

Kurzweil, Ray “Response to Announcement of Chatbot Eugene Goostman Passing the Turing Test.” Kurzweil.net, 10 June 2014. https://web.archive.org/.

_____. *The Singularity is Near When Humans Transcend Biology.* New York, NY: Viking, 2004.

Le, Quoc V., Marc’Aurelio Ranzato, Rajat Monga, Matthieu Devin, Kai Chen, Greg S. Corrado, Jeff Dean, et al. “Building High-Level Features Using Large Scale Unsupervised Learning.” In *Proceedings of the 29th International Conference on Machine Learning.* Edinburgh, Scotland, UK, 2012. https://icml.cc/2012/papers/73.pdf.

Lambeth, Benjamin S. *Air Power against Terror: America’s Conduct of Operation Enduring Freedom.* Santa Monica, CA: RAND Corporation, 2005. https://www.rand.org/.

_____, *NATO’s Air War for Kosovo: A Strategic and Operational Assessment.* Santa Monica, CA: RAND Corporation, 2001. https://www.rand.org/.

Leonhard, Gerd. *Technology vs Humanity: The Coming Clash between Man and Machine.* UK: Fast Future Publishing, 2016.

Li, Yuxi. “Deep Reinforcement Learning: An Overview.” Cornell University, 25 Jan 2017. https://arxiv.org/abs/1701.07274.

Lindblom, Charles E. “The Science of Muddling Through.” *Public Administration Review* 19, no. 2 (Spring, 1959): 79–88. https://www.jstor.org/stable/973677.

Mattis, Jim. *Summary of the 2018 National Defense Strategy of the United States of America*. Washington, DC: Department of Defense, 2018. https://www.defense.gov/.

Mnih, Vlad, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves et al. "Human-Level Control through Deep Reinforcement Learning." *Nature* 518, no. 7540 (2015). https://deepmind.com/.

"Moral Machine—Human Perspectives on Machine Ethics." Moral Machine, 27 September 2016. http://moralmachine.mit.edu/.

Neiger, Chris. "7 Artificial Intelligence Stats That Will Blow You Away." *The Motley Fool*, 4 October 2017. https://www.fool.com/.

"New York City Launches Nation's Most Sophisticated Active Traffic Management System Powered by TransCore's TransSuite Traffic Management Software and RFID [Radio Frequency Identification] Technology." *Business Wire*, 27 September 2011. https://www.businesswire.com/.

Ng, Andrew. "Artificial Intelligence is the New Electricity." Address. Stanford MSx Future Forum, Stanford, CA, 25 January 2017. https://www.youtube.com/.

Russell, Stuart and Peter Norvig. *Artificial Intelligence: A Modern Approach.* Saddle River, NJ: Pearson Education 2003.

Saltzman, Brig Gen Chance, deputy commander, US Air Forces Central Command. Subject: Multi-Domain Command and Control, PowerPoint, 30 Aug 17.

Saygin, Ayse Pinar, Illyas Cicekil, and Varol Akman, "Turing Test: 50 Years Later." *Minds and Machines* 10, no. 4 (2000): 463-518. https://crl.ucsd.edu/

Sessions, Roger. "How a 'Difficult' Composer Gets That Way." *New York Times*, 8 January 1950.

Silver, David, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Gues, Thomas Hubert, et al. "Mastering the Game of Go Without Human Knowledge." *Nature* 550, no. 7676 (October 2017): 354–359. doi:10.1038/nature24270.

Simmons, Bruce. "Exponential Growth." Mathwords.com, 17 July 2017. http://www.mathwords.com/.

Turing, Alan M. "Computing Machinery and Intelligence," *Mind* LIX, no. 236 (October 1950); 433–460. doi:10.1093/mind/LIX.236.433.

United States Air Force. *Air Force Future Operating Concept: A View of the Air Force in 2035.* September 2015. https://www.af.mil/.

United States Army, *Robotic and Autonomous Systems Strategy.* Fort Eustis, VA: US Training and Doctrine Command, March 2017. https://www.tradoc.army.mil/.

Urban, Tim. “The AI Revolution: The Road to Superintelligence.” Wait But Why, 22 January, 2015. https://waitbutwhy.com/.

Vardi, Moshe Y. “Artificial Intelligence: Past and Future.” *Communications of the Association for Computing Machinery* 55, no. 1 (January 2012):5-5. https://cacm.acm.org/.

Analytics Vidhya. “The Evolution and Core Concepts of Deep Learning & Neural Networks.” Analytics Vidhya, 3 August 2016. https://www.analyticsvidhya.com/.

Westmoreland, Gen William C. *A Soldier Reports.* Garden City, NY: Doubleday, 1976.

Winston, Patrick Henry. *Artificial Intelligence.* Reading, Mass: Addison Wesley, 1992.

Work, Robert O. “Artificial Intelligence, Autonomous Systems and the Third Offset.” *In Artificial Intelligence, Big Data, and Cloud Taxonomy*. Department of Defense, 2–5. Arlington, VA, Govini, 2017. https://en.calameo.com/.

______, “Remarks by Defense Deputy Secretary Robert Work at the CNAS [Center for a New American Security] Inaugural National Security Forum.” Address. CNAS Inaugural National Security Forum, Washington, DC, 14 December 2015. https://www.cnas.org/.

Zacharias, Greg, chief scientist of the Air Force. “Autonomous Horizons: System Autonomy in the Air Force.” Address. CogSIMA, San Diego, CA, 24 March 2016. http://www.cogsima2016.org/.