SHERRILL LINGEL, JEFF HAGEN, ERIC HASTINGS, MARY LEE, MATTHEW SARGENT, MATTHEW WALSH, LI ANG ZHANG, DAVID BLANCETT

Joint All-Domain Command and Control for Modern Warfare

An Analytic Framework for Identifying and Developing Artificial Intelligence Applications

Support RAND
Make a tax-deductible charitable contribution at
www.rand.org/giving/contribute

www.rand.org

Preface

In 2019, Air Combat Command Directorate of Plans, Programs and Requirements (A5/8/9) asked RAND Project AIR FORCE (PAF) to examine and recommend opportunities for applying artificial intelligence (AI) and, more broadly, automation to deliberate planning for joint all-domain command and control (JADC2). JADC2 as envisioned integrates the planning, tasking, and assessments of operations across the domains of space, information, cyber, air, land, and sea.

The Air Operations Center (AOC) is the primary operational-level central command and control (C2) node for the U.S. Air Force today. Notwithstanding its historical effectiveness, the AOC construct has been recently challenged for myriad reasons. First, AOC systems and personnel are typically located at a forward-deployed, centralized facility. This constitutes a significant vulnerability and single point of failure. Second, many AOC information systems date back to the inception of the AOC in about 2000. The cancellation of the AOC 10.2 modernization effort has delayed the delivery of critical hardware and software upgrades to the AOC. Third, the growing emphasis on improved cyber and space integration has placed new functional and technical demands on the AOC and increased interest in multidomain operations. Fourth and finally, numerous breakthroughs have occurred in the fields of AI and machine learning (ML). Emerging technologies in these areas can enable new capabilities or, alternatively, can constitute new threats. Thus, this research on AI applications for JADC2 was conducted in PAF's Force Modernization program in order to address the question of how the Air Force can incorporate AI/ML and automation to achieve JADC2. A companion volume,[1] not available to the general public, provides technical details, information about the assessment process, and descriptions of the vignettes.

This report may be of interest to U.S. Department of Defense leaders and stakeholders in AI and/or C2 and to congressional audiences with an interest in AI and enabling the U.S. great power competition.

The research reported here was commissioned by Air Combat Command A5/8/9 and conducted within the Force Modernization and Employment Program of RAND Project AIR FORCE as part of a fiscal year 2019 project, AI Applications for JADC2.

---

[1] Sherrill Lingel, Jeff Hagen, Eric Hastings, Mary Lee, Matthew Sargent, Matthew Walsh, Li Ang Zhang, Dave Blancett, Edward Geist, and Liam Regan, *Joint All-Domain Command and Control for Modern Warfare*: *Technical Analysis and Supporting Material*, Santa Monica, Calif.: RAND Corporation, 2020, Not available to the general public.

RAND Project AIR FORCE

Contents

Figures

Tables

Summary

Issues

- A key concern for the Air Force is the air component commander's ability to integrate capabilities from domains other than air into multidomain operations (MDO).
- The air component commander and Air Operations Center (AOC) staff's ability to integrate MDO is limited by processes, systems, training, and planning and execution experience.
- Given the increased complexity of MDO planning and the greater data requirements, the Air Force will require new tools, including those based on artificial intelligence (AI), to enable joint all-domain command and control (JADC2).
- Focusing investments requires understanding what AI applications offer the greatest increase in operational effectiveness for an identified multidomain concept of operations (CONOPS) across the forces.
- Introducing AI tools requires that the appropriate supportive technological ecosystem be in place.

Approach

- Our research involved a literature review, site visits, and semistructured interviews with operators at geographic AOCs; analysis of technical documents describing the AOC baseline; leveraging emerging MDO CONOPS from recent wargames; and project team workshops on the command and control (C2) processes to enable MDO, AI approaches to address them, and necessary data sources.

Conclusions

- The Air Force AOC 72-hour air-tasking cycle is incongruent with the current digital world. The future will flip the balance between today's emphasis on deliberate planning to dynamic planning, and the JADC2 tools and processes need to support this change.
- Migrating the AOC structure to a modern digital environment poses many challenges: reliance on human-centric subject-matter expert meetings and boards, multiple classification levels of data on air-gapped systems, and heavy reliance on Microsoft Office products.
- Additional factors limit the speed and scope of MDO: authorities and command relationships, synchronizing battle rhythms across domains, different procedures for

different domains, different C2 structures in different theaters and regions, and robust and resilient communications systems and procedures.

- Three enabling categories should be aligned to support MDO: determining a C2 construct for JADC2, the data sources and computing infrastructure available for MDO, and algorithm development to support machine-to-machine processes with multidomain (MD) decisionmakers "on-the-loop" (i.e., the process is overseen by a human who can step in an interrupt it as necessary).
- There are multiple future MDO CONOPS plans, and the needs will vary by campaign. The future C2 structure should be flexible to accommodate the variations.

Recommendations

- The Air Force Warfighting Integration Center (AFWIC), working with Pacific Air Forces and U.S. Air Forces in Europe and Air Forces Africa, should postulate and continue to explore MDO concepts in support of the National Defense Strategy through wargames and tabletop exercises and then inform the broader community about agreed-upon MDO concepts. This would facilitate necessary engagements with other services and U.S. Department of Defense authorities.
- The Air Force Chief Data Office should establish a data-management policy across operations centers to ensure that data are saved and properly tagged for accessibility (including tagged security streams) and that data-storage capacity is sufficient. Experimentation with standardization would help validate the path forward.
- AFWIC should, with Air Combat Command (ACC), evaluate alternative C2 constructs to enable MDO. Additional wargames and workshops to compare and contrast alternatives are needed. ACC follow-on work would develop, organize, train, and equip plans.
- JADC2 progress should occur in a cohesive, progressive, interactive way. The entire enterprise—C2 construct, data management, and the tools, apps, and algorithm development—should adhere to an overarching strategy as illustrated in Figure S.1. As the warfighting integration center, AFWIC should ensure adherence to the strategy, reporting to the Chief of Staff of the Air Force. As new MD concepts emerge from wargames and tabletop exercises, the analytic community evaluates them and the warfighters design command post and other exercises to refine and operationalize the concepts. Moving down the MDO column, the concepts are further developed through live-fly exercises and weapons and tactics efforts. CONOPS move from conceptual to part of the force. These efforts should also inform progress across the three enablers. At the top of each column, guidance from leadership frames the operational and tactical level below it in order to move out on efforts, for example, to develop data infrastructure, data access, and algorithms and to train and equip the commander's staff. As the enablers evolve, changes help inform the MDO CONOPS refinement.

Figure S.1. MDO Informs C2 Construct, Data, and Algorithm Progress
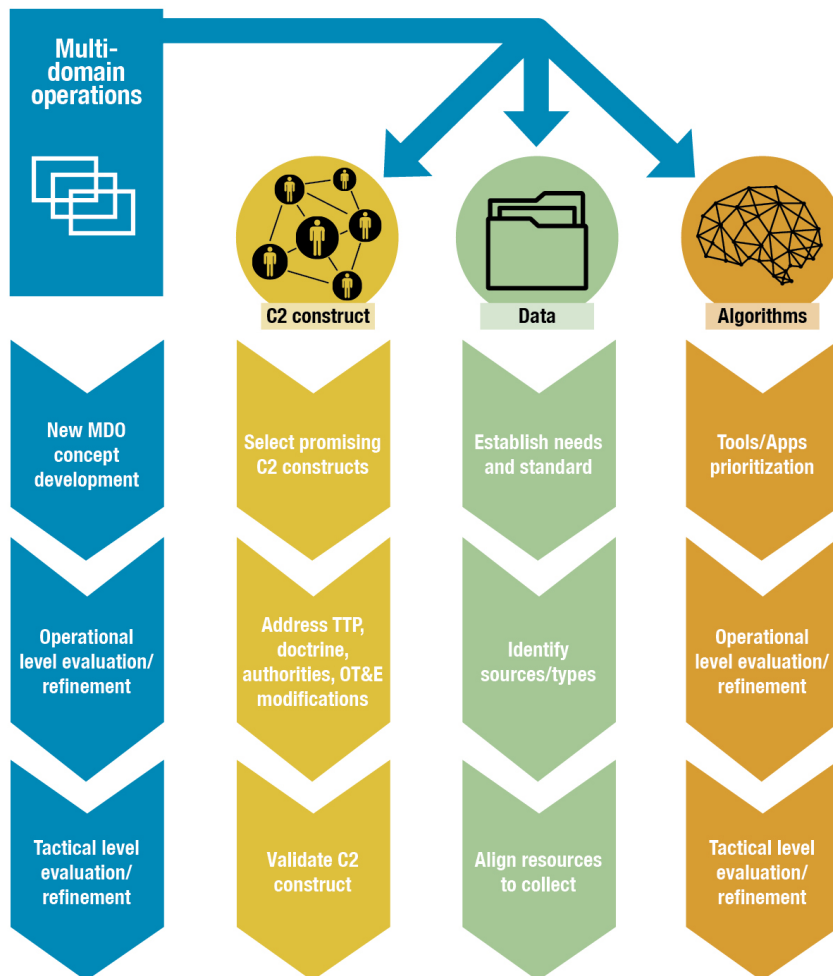
NOTE: TTP = tactics, techniques, and procedures.

Acknowledgments

Abbreviations

ACC	Air Combat Command
AFRL	Air Force Research Laboratory
AFWIC	Air Force Warfighting Integration Center
AI	artificial intelligence
AOC	Air Operations Center
AP	automated planner
APAN	All Partners Access Network
API	application programming interface
ATC	air-tasking cycle
ATO	air tasking order
C2	command and control
C-ATO	continuous authority to operate
CDO	Chief Data Office
CIO	Chief Information Officer
CMCC	Common Mission Control Center
COA	courses of action
COCOM	combatant commander
COD	Combat Operations Division
CONOPS	concept of operations
COP	common operating picture
CYBERCOM	Cyber Command
DaaS	data as a service
DARPA	Defense Advanced Research Projects Agency
DoD	U.S. Department of Defense
DOTMLPF-P	doctrine, organization, training, materiel, leadership and education, personnel, facilities and policy
DT	dynamic targeting
FY	fiscal year
GAN	generative adversarial neural networks
GPU	graphics processing unit
HADR	humanitarian assistance and disaster relief
IaaS	infrastructure as a service
ISR	intelligence, surveillance, and reconnaissance
IT	information technology
JADC2	joint all-domain command and control

JAIC	Joint Artificial Intelligence Center
JEDI	Joint Enterprise Defense Infrastructure
JFACC	joint force air component commander
JFC	joint force commander
KMCWG	Knowledge Management Capabilities Working Group
KR	Kessel Run
MAAP	master air attack plan
MD	multidomain
MDC2	multidomain command and control
MDO	multidomain operations
ML	machine learning
MLS	multilevel security
NLP	natural language processing
OA	operational assessment
OC	operations center
OODA	observe, orient, decide, and act
OT&E	Operational Test and Evaluation
PaaS	platform as a service
PACAF	Pacific Air Forces
PE	program element
RDT&E	research, development, test, and evaluation
RSPACE	Resilient Synchronized Planning and Assessment for the Contested Environment
SaaS	software as a service
SAF	Secretary of the Air Force
SAF/AQ	Secretary of the Air Force/Acquisition, Technology and Logistics
SAM	surface-to-air missile
SEAD	Suppression of Enemy Air Defenses
ShOC	shadow operations center
SRD	Strategy Division
sUAS	small unmanned aircraft system
TTP	tactics, techniques, and procedures

1. Challenges of Implementing Joint All-Domain Command and Control Within the U.S. Air Force's Current Operational Level Construct

Multidomain operations (MDO) arguably represent a means of waging warfare increasingly employed by the United States and others for decades. Land, sea, and air forces are brought to bear against adversary air, land, and sea forces that often rely on capabilities from space-based systems and digital information from computer networks. Given these additional domains, what is driving today's modern warfare vision for MDO? In modern warfare, space is no longer a sanctuary in which offensive and defensive actions are rare, and cyberattacks on adversary networks and network self-defense are expected and planned for likely before kinetic hostilities even begin. In fact, actions in these newer warfighting domains are happening in the current gray zone, further emphasizing the urgency for close synchronization of efforts among all domains available in competitions and conflicts.

The Air Operations Center (AOC) is the primary operational-level central command and control (C2) node for the U.S. Air Forces. These physical centers with large staffs are where planning, execution, and assessment of air operations occur. There are currently several region AOCs and six functional AOCs (Table 1.1) around the world. Functional operations centers (OCs) support functional combatant commanders (COCOMs) in the areas of global strike, space, mobility, special operations, cyber, and intelligence, surveillance, and reconnaissance (ISR). Regional AOCs support geographic COCOMs and are used for planning and executing theater operations in support of a joint force commander (JFC).

Current Operational-Level Command and Control Challenges

The technological architecture underlying the baseline AOC system is made up of a patchwork of C2 systems that enable all phases of the air-tasking cycle (ATC). The AOC is normally employed by the joint force air component commander (JFACC) to exercise control of air forces in support of combined and joint force objectives. The effectiveness of the AOC construct, along with the associated doctrinal concepts of a functional air component commander and a centralized air planning process with decentralized execution, was demonstrated to great effect in Operation Desert Storm, galvanizing the role of the AOC in operational-level C2. The AOC has been tested in numerous conflicts since Desert Storm and, although methods for employing it have evolved, the AOC has remained the accepted paradigm for C2 of air forces.

Nevertheless, the AOC construct has recently been challenged for a number of reasons, not the least of which is the increasing attention on MDO: The growing emphasis on improved cyber and space integration has placed new functional and technical demands on the AOC. Its ability to

handle these challenges is limited by processes, systems, training, and planning and execution experience. In addition, emerging technologies in the fields of artificial intelligence (AI) and machine learning (ML) are enabling new capabilities and new threats.

RAND Project AIR FORCE was asked by Air Combat Command (ACC), Directorate of Plans, Programs and Requirements (A5/8/9) to examine and recommend opportunities for applying AI and, more broadly, automation, to deliberate planning and products for multidomain command and control (MDC2). The companion report, not available to the general public, provides technical details for this research and an extensive review of previous research.[1]

Table 1.1. Air Force Air Operations Centers

Type	AOC	Major Command	Station
Regional	601 AOC	ACC	Tyndall Air Force Base, Florida
Regional	603 AOC	U.S. Air Forces Europe	Ramstein Air Base, Germany
Regional	607 AOC	Pacific Air Forces (PACAF)	Osan Air Base, South Korea
Regional	609 AOC; 609 AOC Detachment 1	U.S. Air Forces Central Command	Al Udeid Air Base, Qatar Shaw Air Force Base, South Carolina
Regional	611 AOC	PACAF	Joint Base Elmendorf-Richardson, Alaska
Regional	612 AOC	U.S. Air Forces Southern Command	Davis-Monthan Air Force Base, Arizona
Regional	613 AOC	PACAF	Joint Base Pearl Harbor-Hickam, Hawaii
Functional	608 AOC	Air Force Strategic Strike Command	Barksdale Air Force Base, Louisiana
Functional	614 AOC	Air Force Space Command	Vandenberg Air Force Base, California
Functional	618 AOC	Air Mobility Command	Scott Air Force Base, Illinois
Functional	623 AOC	Air Force Special Operations Command	Hurlburt Field, Florida
Functional	624 OC	Air Force Space Command	Lackland Air Force Base, Texas
Functional	625 OC	ACC	Lackland Air Force Base, Texas

The AOC commander is responsible for managing operations, establishing battle rhythms, and planning, coordinating, executing, and assessing air operations based on JFC or commander Air Force guidance.[2] The AOC battle rhythm follows a nominal 72-hour ATC (Figure 1.1). The

---

[1] Sherrill Lingel, Jeff Hagen, Eric Hastings, Mary Lee, Matthew Sargent, Matthew Walsh, Li Ang Zhang, Dave Blancett, Edward Geist, and Liam Regan, *Joint All-Domain Command and Control for Modern Warfare: Technical Analysis and Supporting Material*, Santa Monica, Calif.: RAND Corporation, 2020, Not available to the general public.

[2] Air Force Instruction 13-1AOC, *Operational Procedures—Air Operations Center (AOC)*, Vol. 3, Washington, D.C.: Department of the Air Force, November 2, 2011, incorporating change 1, May 18, 2012.

cycle begins with the Strategy Division (SRD) defining objectives, effects, and guidance for the air tasking order (ATO) period. The continuous nature of the 72-hour ATC means that as one day's ATO is being executed and monitored by the Combat Operations Division, the ATOs for the next two days are being planned by the SRD and Combat Plans Division.

Figure 1.1. Air-Tasking Cycle
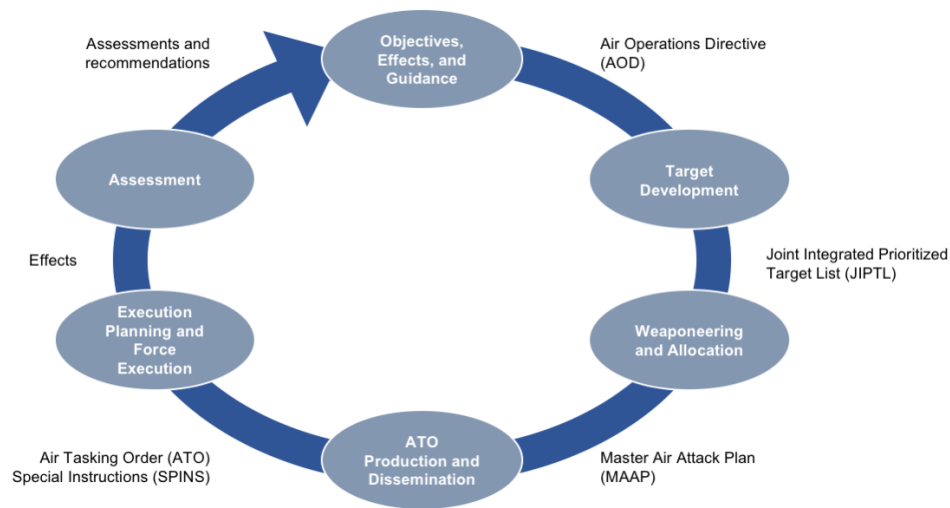
SOURCE: U.S. Air Force, *Operational Employment: Air Operations Center*, AFTTP 3-3.AOC, March 31, 2016a, Not available to the general public.

The Air Force currently operates under the paradigm of centralized control and decentralized execution, in which the C2 of air operations is concentrated in the AOC and execution of the ATO is expected to be done across subordinate nodes. AOCs work in coordination with higher headquarters; lateral organizations, such as other joint force centers and components; subordinate units; and other OCs as necessary. AOCs are connected to each other and to forward nodes through classified and Non-Classified Internet Protocol Router Network networks and telecommunications systems that enable voice, phone, chat, and email capabilities. Examination of alternative C2 constructs, such as distributed control or multidomain (MD) C2, are underway with the Doolittle Wargames and activities at geographic commands.[3]

Two challenges of the current AOC construct of great concern and worth special mention are the speed of the current planning process and the focus on deliberate planning.

Planning Speed

AOC systems and processes have evolved to best enable a commander's objectives to be efficiently transformed into targets and sorties to find and attack them. Conflicts in the past 20

---

[3] Miranda Priebe et al., unpublished RAND Corporation research, 2019.

years conducted with AOCs have allowed a reasonable amount of planning and reaction time. Support elements, such as electronic warfare, cyber, and space, have been integrated to the greatest extent possible and in a largely manual fashion. This would be challenging in a high-end conflict and will become increasingly infeasible as the scope of MDO expands. The canonical 72-hour ATC works well in such fairly steady-state operations as a strategic bombing campaign; strikes against static, dug-in forces; support to preplanned ground operations; and planning ISR support against an insurgency. Much of this planning cycle is taken up by coordination meetings among components, commanders, and various echelon levels, which begin up to 72 hours before execution. Targets are carefully developed using many sources of intelligence and munitions, and platforms are deliberately selected one to three days in advance to maximize probability of kill and to minimize risk and collateral damage. Commander and legal reviews occur at multiple levels. Coalition partners must be integrated, synchronized, or at least deconflicted. Plans also are released to wing- and squadron-level units more than ten hours before execution to allow detailed tactical-level planning to minimize risk and maximize effectiveness.

All of this care allows for the management of risk and the efficient employment of airpower, but these plans take time. This comprehensive planning process is exacerbated by the fact that the United States is likely to be on the defense in most near-peer conflicts, attempting to delay and stop a plan that an adversary has put into motion. Unless the preplanning is almost perfect, a 72-hour ATC will not work well if the adversary can achieve its objectives in 60 hours, for example.

Deliberate Versus Dynamic Planning

Most AOC processes are directed toward building an ATO to service preplanned targets. This is a direct and natural consequence of the types of conflicts in which the AOC has matured and can be exemplified by the number of personnel authorized to the deliberate targeting process and those assigned to dynamic targeting (DT) (the companion report provides details about manning).

The Combat Operations Division, where the DT cell is located, is much smaller than it appears. First, almost half of the division is dedicated to defensive—not offensive—operations. Second, the vast majority of the other cells, such as strategy, combat plans, ISR, and Air Mobility Division, are almost fully dedicated to *deliberate* planning and do not support DT.

If critical targets are not known in advance, most of the AOC's deliberate planning processes are less relevant and sorties must simply wait for targets to be found. This was the case in Operation Iraqi Freedom, for example, in which 79 percent of the desired mean points of impacts were against targets planned outside the ATO cycle.[4] Because much of the deliberate targeting process is omitted in these cases, these sorties and their support (such as tankers and ISR) are

---

[4] T. Michael Moseley, *Operation Iraqi Freedom: By the Numbers*, PDF slides, U.S. Air Forces Central Command, Prince Sultan Air Base, Kingdom of Saudi Arabia, April 30, 2003, p. 15.

inherently less efficient and riskier than those with preplanned targets and mission packages. To maximize efficiency and minimize risk, planning for DT involves determining the best C2 process (such as kill boxes), coordinating with the ISR or other component assets that might find and identify targets, determining the most-flexible platform types and munition loadouts, deconflicting airspace when routes to targets are determined, providing tanker support for sorties that might be loitering for long periods of time, and performing target engagement approval if necessary. All of these DT steps are primarily conducted manually using human-to-human communications and reasoning to navigate the large number of stove-piped databases and chains of command. As a result, simply using the current AOC to generate sorties for broad use in DT is not a realistic expectation, given the current AOC staffing bias toward deliberate planning; the manpower and tools are simply not adequate to handle that number of sorties, assuming that the number of dynamic targets could be found. In addition, timeliness, a critical element of a successful engagement against dynamic targets that can relocate, would be affected.

Differences Across Air Operations Centers

The AOC baseline architecture is made up of more than 40 applications connected by thousands of machine-to-machine interfaces spread across three to ten networks.[5] The two applications installed for the greatest number of duty positions are Internet Relay Chat and business services products. The former is a major point of integration between human planners, and the latter is an all-purpose tool that allows operators to bypass limitations of major command, control, communication, and intelligence systems. Problematically, Internet Relay Chat and business services products have no interfaces with other AOC applications, pointing to major seams between a primary communication channel, a primary planning tool, and all other AOC applications.

Each AOC weapon system has evolved to be its own distinct entity because of the unique needs and concerns of its functional demands or area of responsibility. For example, the remote location of the 613 AOC and the large distances between its subordinate and commanding nodes are a challenge to air operations, whereas the 603 AOC must manage challenges and threats within its areas of responsibility, which include Europe and most of Africa, as well as interface with the North Atlantic Treaty Organization. Funding streams to various AOCs also differ.

Functional AOCs also differ by mission. Although all AOCs should be capable of basic tasks, such as planning, execution, and operational-level assessments, some AOCs also have additional tasks, such as cyberspace defense or theater air and missile defense.

Although no *true* AOC baseline exists, and each has evolved to meet the needs dictated by its mission, many needs—and challenges—are also shared across AOCs.

---

[5] Air Force Life Cycle Management Center, Battle Management Directorate, *Descriptive List of Applicable Publications (DLOAP) for the Air Operations Center (AOC)*, Hanscom Air Force Base, Mass., April 1, 2019, Not available to the general public.

Challenges Unique to Space and Cyber Operations Centers

Air Force Chief of Staff General David L. Goldfein described effective MD integration as "using dominance in one domain or many, blending a few capabilities or many, to produce multiple dilemmas for our adversaries in a way that will overwhelm them."[6] The challenges with better integrating the air, space, and cyber domains are numerous and daunting, including such problems as multilevel security, lack of authorities at the regional AOC level, and lack of compatible communications. In the case of cyber, many different levels of security are involved, and authorities for planning and execution often reside far from the theater and component levels, and few of these authorities are delegated forward.

Challenges in Emergence of Near-Peer Threat Environments

The AOC and the JFACC's execution of the ATC face external threats. As the AOC has grown and evolved, it has been upgraded and modified many times to keep pace with advancements in technology, doctrine, and the forces under its control. Although the systems, processes, and personnel that constitute the weapon system have been tested in a variety of conflicts—from regional war in Europe to counterinsurgency operations in the Middle East— they have never faced a challenge similar to the conflicts envisioned by the National Defense Strategy involving China or Russia. Such high-end conflicts raise two additional concerns with current AOC operations:

- survivability of the AOC facility and personnel itself
- robustness to cyber and communications attacks.

These concerns will be exacerbated in MDOs, which involve a greater number of personnel and information systems and for which planning processes are even more prolonged and operations are more reliant on fragile connections that a sophisticated adversary has capabilities to deny, degrade, and disrupt.

Survivability

As the centerpiece of air operations planning, a JFC is wholly dependent on AOC processes for planning and executing an air campaign. However, for overseas AOCs with significant warfighting responsibilities, their ability to operate in the face of kinetic attack is not at all assured.[7] For example, the 613 AOC at Hickam Air Force Base on Oahu and the 603 AOC at Ramstein Air Base in Germany would support U.S. Indo-Pacific Command and U.S. European

---

[6] David Goldfein, "CSAF Focus Area: Enhancing Multi-Domain Command and Control . . . Tying It All Together," Washington, D.C.: Chief of Staff, U.S. Air Force, March 2017.

[7] Kristin F. Lynch et al., unpublished RAND Corporation research, 2017.

Command, respectively, in fights against China or Russia. However, both organizations place their systems and personnel in unhardened above-ground centralized facilities.

Although all AOCs are responsible for creating and maintaining a continuity of operations with backup facilities and personnel, it appears likely that a significant delay would occur as damage assessment and handover were conducted between sites.

The loss of personnel to an AOC attack is also likely to have far-reaching effects because AOCs contain the personnel most qualified to plan and conduct air operations in a particular theater, as well as liaisons with many different components and units, making an AOC a particularly high-value target for adversaries.

Robustness to Cyber and Communications Attacks

Similar to the physical attacks on AOC facilities and personnel, the information and communications systems of the AOC are also likely to be threatened. As numerous cases with civilian infrastructure illustrate, even supposedly air-gapped computer networks can be attacked with modern tools and sufficient preparation. The AOC operates primarily with commercial hardware and software, leading to concerns that many tools and techniques commercially developed could be used against it. Migrating AOC systems to higher classification networks and computer systems to increase the security posture would be challenging because the AOC is dominated by information at the collateral SECRET level.

Likewise, communications networks are classic high-payoff targets for attack, and, in the Pacific theater, there is already significant worry that plans created on Hawaii or in the continental United States will be difficult to push forward to units deployed to such locations as Guam, Japan, and the Philippines. Much of the military beyond-line-of-sight communication occurs over leased channels on commercial satellites, which are vulnerable to kinetic, electronic warfare, and possibly cyberattacks. Underground and undersea cables are likely more robust against many of these threats, but spying and interruption are clear threats here as well.

Central Multidomain Operational Challenges for AOCs

One of the key concerns for the AOC moving forward is the ability of the air component commander to integrate capabilities from domains other than air into so-called MD operations. MDO has received much attention recently, particularly from the U.S. Army and the Air Force, and has been the subject of numerous speeches, publications, wargames, and exercises. Brigadier General Chance Saltzman, until recently the head of an Air Force Enterprise Capability Collaboration Team on MDC2, now evolved into joint all-domain command and control (JADC2), lists four key elements:

- AF Multi-domain Ops must include generation of offensive and defensive effects from air, space & cyber with capability to independently and directly support JFC objectives

- MDO is more than systems in one domain supporting operations in another domain (necessary but not sufficient)
- MDO are *high velocity*, *operationally agile operations* that present multiple dilemmas for an adversary at an operational tempo they cannot match.
- MDO requires seamless, dynamic and continuous integration of capabilities generating effects in and from all domains . . . *MDC2*[.][8]

The emphasis on improved space and cyber integration (as opposed to sea or land, for example) by the Air Force appears to be an outgrowth of three trends:

- Near-peer threats are assessed to be on a path to match the Air Force in force-on-force capabilities.
- The Air Force possesses significant and growing organic capabilities in space and cyber.
- Training, exercises, and real-world operations all indicate that, for many reasons, attempts to *integrate* air operations with space or cyber are rare, slow, cumbersome, and fraught with error when they are attempted.

Although naval and ground forces have their own component commands under a JFC and mechanisms for coordinating with the air component, space and cyber forces do not provide capabilities to the JFC in the same way and typically do not provide the JFC as much control over those forces. Some capabilities require the Secretary of Defense or the President to approve. Some (such as gaining cyber accesses) require much longer planning times than the ATC or can be executed only at specific times (because of orbital mechanics in the case of non-geostationary satellites). In addition, these capabilities are often handled at classification levels above collateral SECRET, where the AOC operates;[9] thus, discussions and planning must be performed using specialized facilities and communications channels. The air, space, and cyber communities, even within the same service, have unique cultures and languages and typically have few opportunities to operate from the same location and share knowledge. Most training activities are single domain, and opportunities for integration are usually prescribed. Even Air Force space and cyber capabilities, which, on the surface, would appear to be easily available to the JFACC, are largely controlled at much different levels of authority in very different chains of command by organizations, such as U.S. Space Command and Joint Force Space Component Commander or U.S. Cyber Command (CYBERCOM).[10] These supporting commands may also face legal

---

[8] Italics original, Chance Saltzman, "MDC2 Overview," briefing for 2018 C2 Summit, June 2018.

[9] These challenges are not unique to space and cyber. For example, they also exist with special access programs (e.g. F-35), as well as coalition environments.

[10] Both U.S. European Command and U.S. Indo-Pacific Command have space-integrated project teams at the COCOM level (like cyber-integrated project teams).

restrictions, such as the review and approval process for cyber operations, and may require permission from other organizations, such as the National Security Agency, to use infrastructure.

Each domain has evolved with its own resources, limitations, and culture—and, as a result, its own procedures. For example, space operations are not planned and executed in the same way as air or cyber operations. Functional command operating centers—such as the 624th for cyber, the Combined Space Operations Center for space, the 618th for mobility, and the 608th for global strike—each have distinct planning cycles.[11] As these operating centers plan their own unique missions, human liaisons at each OC translate and integrate capabilities and effects into a *normal* air-targeting cycle. The targeting process is another example. The air component focuses on selecting aimpoints to achieve desired effects. Although the cyber-targeting process has the same overall goal (achieving effects), the cyber warfighter may need to "attack" an entire series of targets (routers on different networks, firewalls, servers, user accounts, etc.) to reach the ones that make that effect possible. Although potential opportunities exist for kinetic effects and cyberattacks to complement one another, synchronization across domains is again mediated by human liaisons in a time-consuming and limited manner.[12] Although timelines could possibly be synchronized with automated systems, it is more difficult to imagine differences in process being resolved automatically in an intelligent and mutually acceptable way. For instance, cyber tools may be designed for one-time use, and their effects may be nonconventional, complicating comparison with and selection against traditional kinetic effects.

Approach and Organization of This Report

Approach

We conducted a multimethod research project that involved literature reviews; site visits and semistructured interviews with operators at geographic AOCs, the science and technology community, and other Air Force stakeholders; and analysis of technical documents describing the AOC baseline. This research had a two-pronged approach toward C2 for multiple domains.

First, the team examined the AOC mission threads and other technical documents, collected input from AOC personnel on their use of business products (e.g., Microsoft Office) and the security level of AOC products, and used network mapping with R programming to explore opportunities for automation and AI/ML within the existing AOC processes. Key findings from that analysis determined that AI/ML opportunities are sprinkled throughout the AOC processes, space and cyber are largely disconnected from the overall AOC process, and multilevel

---

[11] Although the 618th and 608th are both air components, the specialization and scale of air mobility and global strike planning necessitate dedicated AOCs.

[12] A secondary concern with these human-centric negotiations is that they are rarely captured in any kind of structured way to allow algorithms to eventually learn from them. Chat room archives may be the most likely source of this type of data, if they can be processed to learn content and meaning.

classification provides a significant complicating factor for AI/ML. Details of the assessment process and results are in the companion report. This approach proved problematic because of its adherence to existing AOC processes, which are linked to the challenges discussed earlier.

The second line of effort started with examining C2 processes needed to enable three example MDO concepts of operations (CONOPSs). Recent wargames within the Air Force and between the services have suggested new MD CONOPSs that will purportedly provide cost-effective means of meeting military campaign objectives. We started with these MD operational concepts and examined the C2 capabilities needed to enable them (the companion report contains detailed descriptions of example vignettes). AI/ML categories are then mapped to each C2 process to identify promising AI capabilities to address MDO CONOPS needs. These inputs were synthesized to propose applications of AI/ML that can address future MDO concepts and their associated challenges and requirements. The team also examined the ecosystem necessary to leverage AI/ML for JADC2.

Terminology

For the purpose of this report, we define *AI* as an academic discipline concerned with machines demonstrating intelligence. Definitions vary along two dimensions. The first concerns how machines reason, and the second concerns how they act.

ML is a subfield of AI that concerns machines performing tasks without first receiving explicit instructions. Instead, the machine learns to perform the task from training data or through interactions with a simulation environment.

Although the terms *MDC2* and *MDO* are pervasive, they are poorly defined and mean different things to different communities. With the change to JADC2, the confusion persists. To provide a consensus view on the meaning and scope of MDC2 (and JADC2), we reviewed several recent briefings, articles, and technical reports on MDC2 (the companion report provides details) to address three questions: (1) What are common elements of MDC2 (currently JADC2) definitions? (2) What are the challenges to achieving JADC2? and (3) What technologies and capabilities are needed to enable JADC2?

Our review of these documents revealed three recurring themes in MDC2 definitions: (1) the use of cross-domain effects in conjunction with or in lieu of one another, (2) the ability to share information across domains, and (3) a framework for joint and combined mission command to enable cross-domain effects and information sharing.

Organization of This Report

This report is divided into two volumes, of which this is the first. Chapter 2 briefly describes current Air Force efforts and funding levels for C2 for air, space, and cyber and the science and technology community's pursuit of C2 and AI capabilities and the limitations of these efforts; the programs are described in detail in the companion report. Chapter 3 introduces an analytic approach that uses vignettes to identify AI/ML opportunities for future MDOs and highlights

themes in types of algorithms needed; the three vignettes are outlined in Chapter 3, with details provided in the companion report. Chapter 4 discusses the ecosystem needed to support an AI/ML-enabled JADC2. Chapter 5 concludes by presenting a cohesive strategy to move forward toward a modern JADC2 capability.

2. Command and Control Modernization

This chapter focuses on AOC modernization, although a cursory look at space and cyber C2 within the Air Force and the Navy, and Army C2 funding streams are included for perspective. AOC funding and modernization efforts demonstrate three trends: (1) single domain to MD; (2) manual to machine; and (3) local to cloud-based.

Modernization of the Air Operations Center in the Direction of Multidomain

In 2016, cybersecurity vulnerabilities and third-party application integration challenges led to the termination of the AOC modernization (version 10.2) architecture.[1] The AOC Program Management Office shifted direction, implementing the in-house modernization effort called the AOC Pathfinder, which has since evolved into Kessel Run (KR). KR is applying Agile software development to AOC modernization.[2] Given KR's early success in developing and delivering applications to the AOC using Agile and given Agile's widespread adoption for commercial software development, Agile development and operations has been identified as one component of the strategy to evolve C2 and nascent MD capabilities for the AOC. Other components include virtualized (i.e., cloud-based) data structures and a common C2 platform. Most recently, in response to the MDC2 Enterprise Capability Collaboration Team's recommendation to leverage advanced technologies for C2, the Air Force created ShadowNet, a multinode network of development and operations environments linked together to experiment for the purpose of addressing enterprisewide C2 challenges (including security, latency, bandwidth, data access, network resiliency operations, and information technology [IT] challenges).[3] In parallel to pursuing the KR development efforts, the Air Force has initiated a series of experimental campaigns under ShadowNet to evolve MD concepts and to identify and develop technologies to enable them.

Since 2000, AOC modernization has increasingly moved in the direction of MD operations, as evidenced by budget line items. Table 2.1 shows the fiscal year (FY) 2020 budget for the

---

[1] Thomas Light, Brian K. Dougherty, Caroline Baxter, Frank A. Camm, Mark A. Lorell, Michael Simpson, and David Wooddell, *Improving Modernization and Sustainment Outcomes for the 618th Air Operations Center*, Santa Monica, Calif.: RAND Corporation, 2019, Not available to the general public.

[2] Agile software development involves iteratively developing, testing, and gathering user feedback on software before adjusting goals for the next iteration.

[3] David L. Goldfein, *Multi-Domain Command and Control (MDC2): Enterprise Capability Collaboration Team (ECCT), Campaign Plan, Strategy Document*, Washington, D.C.: U.S. Air Force, January 2, 2018, Not available to the general public; Saltzman, 2018.

MDC2 program element (PE) by cost category item.[4] These values underscore the nascent emphasis on MDC2 (now JADC2). Just one year prior, a single cost category item valued at $14.9 million explicitly identified MDC2. The FY 2020 MDC2 cost categories are informative with respect to the Air Force's MDC2 research, development, test, and evaluation (RDT&E) strategy and emphasize (1) a common platform for accelerating the development and delivery of C2 applications, (2) an enterprise data lake to fuel C2 applications and enable future uses of AI/ML,[5] and (3) an experimentation campaign to identify new technological needs arising from MDC2. The FY 2020 President's Budget identifies KR and ShadowNet as leading these MDC2 thrusts, although the detailed activities they are intended to support and the timeline for delivering new capabilities are poorly defined.[6] Future funding lines may identify MDC2 emphasis; however, whether the funding will continue to be stove-piped by domain and service is unclear.

The companion report, not available to the general public, describes the roles and contributions of KR and ShadowNet, as well as those of the Defense Advanced Research Projects Agency (DARPA), Air Force Research Laboratory (AFRL), AFWERX, and others, to the evolution toward JADC2.

Table 2.1. Product Development for Multidomain C2 Project

Cost Category Item	Performing Activity	Cost (in Millions)
MDC2 C2 common platform	KR	$60
MDC2 enterprise data lake	Various	$40
MDC2 ShadowNet experimentation	Various	$36
MDC2 development	Various	$7
Support, test and evaluation, and management	Various	$8
	Total	$151

SOURCE: Under Secretary of Defense (Comptroller), undated.

---

[4] Headquarters U.S. Air Force, Financial Management and Comptroller, "*Air Force President's Budget FY21*," 2019.

[5] An *enterprise data lake* is a centralized repository to store vast amounts of structured, semistructured, and unstructured data as they are generated.

[6] Under Secretary of Defense (Comptroller), "DoD Budget Request: Defense Budget Materials—FY2021," webpage, undated.

C2 efforts extend beyond the MDC2 PEs. Space and cyber within the Air Force also have C2 PEs, as do the Army and Navy. For FYs 2018 and 2019, cumulative spending during each year was highest for Air Force space, followed by Air Force cyber, and this pattern is projected to continue for FY 2020. Of particular note were RDT&E investments in the Joint Space Operations Center, the primary operational-level C2 node for space, and the investments in capabilities to better integrate the cyber-tasking process with the AOC and to increase visibility of cyber operations to combatant, service, and joint commanders. Army and Navy expenditures for C2 were of the same order of magnitude as for the Air Force. The C2 PEs do not fully capture all the efforts for JADC2 or AI within the U.S. Department of Defense (DoD). Other stakeholders are part of associated efforts, as listed in Table 2.2.

Table 2.2. Additional Stakeholders and Contributors for Air Force C2

Effort	Emphasis	Description
Common Mission Control Center (CMCC)	Capability	Central hub for battle management and C2 across a range of next-generation aircraft and in a multilevel, secure environment
Joint Enterprise Defense Infrastructure (JEDI)	Capability	Amazon Web Services–provided service in charge of hosting and distributing mission-critical workloads and information to warfighters across the globe
AFWERX	Concepts and requirements	Catalyst for agile Air Force engagement across industry, academia, and nontraditional contributors to drive innovation
Combined Air Operations Center–Nellis	Concepts and requirements	Provides advanced training for operational C2 through MD and joint integration
Secretary of the Air Force/Chief Information Office (SAF/CIO)	Information	Advises the SAF and the Chief of Staff of the Air Force on IT, cyberspace, and national security systems
Air Force Chief Data Office (AF CDO)	Information	Develop and execute an enterprise data management strategy for the Air Force
Air Force Knowledge Management Capability Working Group	Information	Manages the Air Force Knowledge Management (KM) community and KM-related requirements and interests
Cyber mission platform	Platform	Infrastructure and platform for the Air Force's Offensive Cyber Product Line. Uses Agile acquisition strategy and has implemented rolling authorization to operate
National Space Defense Center	Platform	Creates unity of effort and information in space operations among DoD, the Intelligence Community, and interagency, allied, and commercial space entities. Provides data as a service (DaaS) to facilitate movement of data across classification boundaries and provides common, standard interfaces that facilitate large-scale integration
Open Architecture Distributed Common Ground System	Platform	Technology component of Air Force Distributed Common Ground System modernization framework. Provides computing hardware, server virtualization, data cloud, and service-oriented architecture stack to support rapid development and transition of incremental capabilities

Several efforts are underway to work toward standardization, infrastructure for connecting databases, and the development and operations and ATO necessary to rapidly develop C2 capabilities; however, there are no AI or ML uses to date. Several of these endeavors hold promise for the basis for AI applications. OneChat is one example of an app developed in such a way as to be accessible for ML by providing DaaS functions in a multilevel security (MLS)

context.[7] KR produces front-end service apps (or software as a service) that could act like Trojan horses collecting user OC data for future data mining. DARPA and AFRL have early research in AI/ML for JADC2. Still, an overarching vision or strategy for automating or leveraging AI is elusive.

Summary

As explained in Chapter 1, no single standard baseline exists for today's AOC, and the challenges for enabling MD C2 are myriad. Furthermore, given the threat environment, the demise of the 10.2 modernization effort, and the envisioned MD concepts of employing forces, a disparate series of efforts to modernize the AOC weapon system have emerged, as summarized here (and further described in the companion report). Volume II also provides an examination of the existing ATC framework in order to look for opportunities for AI to enable JADC2. Opportunities discovered in this manner may help inform current efforts to address existing shortfalls.

In response to the vulnerability of centralized AOCs, the Air Force has many initiatives and ideas for modernizing C2 that use a distributed construct of C2 rather than the traditional physically centralized tenet of centralized control and decentralized execution. The design of distributed C2 should be based on an analysis of risk factors, such as feasibility, inefficiencies, costs, resources, and threats by peer adversaries. Whatever the implementation, various investments in new technologies and practices will be necessary to evolve from centralized C2 physical centers. AI/ML can help enable this shift to distributed control—for example, by providing predictive tools (e.g., for force readiness at a wing operations center), dynamic courses-of-action (COA) generation at a subordinate node, and decision tools for commanders at forward operating nodes.

---

[7] Jeffery Compoc, *Obtaining MDC2 Capabilities Inside Technocratic Boundaries: Satisfying Multi-Domain Operational Desires Through an Acquiescent Technical Strategy*, Nellis Air Force Base, Nev.: Air Combat Command, 505 Test and Training Squadron, Combined Air and Space Operations Center–Nellis/Shadow Operations Center–Nellis, March 1, 2019.

3. Artificial Intelligence Opportunities for Future Multidomain Operations

This chapter illustrates an approach to identify AI/ML capabilities that enable JADC2. The chapter begins with brief definitions and a summary of the various categories of AI/ML technologies with relevant examples. Then we outline three MD vignettes; the companion report, not available to the general public, provides additional details. We also show the mapping of C2 processes from the vignettes to salient AI/ML capabilities and conclude with a summary across the three vignettes.

The Enabling Role of Artificial Intelligence and Machine Learning

For the purposes of this report, we define *AI* and *ML* as follows: *AI* is an academic discipline concerned with machines demonstrating intelligence. Definitions vary along two dimensions. The first concerns how machines reason, and the second concerns how they act. For example, a machine may act rationally by reasoning exhaustively and logically or by using simple heuristics. Alternatively, a machine could reason and act like a human, which may be suboptimal but nonetheless adaptive.[1]

ML is a subfield of AI that concerns machines performing tasks without first receiving explicit instructions. Instead, the machine learns to perform the task from training data or through interactions with a simulation environment. Neural networks are one class of ML techniques, along with many other parametric and nonparametric statistical techniques. A neural network comprises a collection of units arranged in layers. The network learns a set of connection weights between units to produce the correct outputs given different inputs. Deep learning is a particular implementation of neural networks that involves transforming inputs across a large number of intermediate layers prior to producing an output. AI/ML systems using deep learning have now achieved superhuman-level performance in tasks involving image classification, speech recognition, and game play. Advances in deep learning have been enabled by parallel developments in computer engineering. The recent success of AI/ML systems in such games as Go, no-limit poker, and StarCraft hints at the possibility of using these systems for solving practical decision problems. Much of the attention stems from the ability of contemporary AI/ML systems to cope with the substantial size and complexity of these games, the need to make decisions based on imperfect information, and the incorporation of real-time continuous play. The progression of task characteristics, from perfect to imperfect information and from turn-based to real-time play, increasingly resembles operational C2 processes. The

---

[1] Kristin Lynch et al., unpublished RAND Corporation research, 2017.

successes of such AI/ML-based systems as AlphaZero, Libratus, and AlphaStar on these games suggests that AI/ML could similarly be applied to solving certain real-time decision problems with imperfect information.

AI technologies may be grouped into six categories:[2]

1. *Computer vision* involves detecting and classifying objects in the visual world. Project Maven contains examples of AI algorithms to detect and classify objects in natural scenes.
2. *Natural language processing* (NLP) performs speech and text recognition and translation and is currently in millions of U.S. homes that use Google Home and Amazon Echo.
3. *Expert systems* are rule-based systems created using large amounts of expert knowledge. They have been used for decades in medical and financial decisionmaking and other fields.
4. *Planners* solve scheduling and resource allocation problems and have been used, for example, to reduce Google's energy use and to achieve superhuman-level performance in real-time strategy games.
5. *ML* involves acquiring knowledge from curated examples in a training set (i.e., supervised learning) or through interactions with a real or simulated environment (i.e., reinforcement learning). ML has been applied to such practical problems as credit card–fraud detection and is a component of many of the AI/ML systems mentioned in the previous paragraph. The methods in this category are general and can be applied in conjunction with each of the first four categories.
6. *Robotics* use a combination of AI/ML capabilities for sensing, scene processing, planning, and action selection to allow an embodied system to interact with the environment. This category reflects the integration of methods from some or all of the preceding categories.

Of these six, the first five categories all have clear applications to JADC2. Computer vision could be used to process multisource intelligence and to perform data fusion. NLP could be used to extract intelligence from speech and text but also to monitor friendly chat to route relevant information to individuals and to alert them to potential conflicts or opportunities. Expert systems could be used to recommend effects to achieve operational and tactical objectives. Planning systems could be used to pair various air, space, and cyber assets against targets and to generate a time-phased scheme of maneuver. Finally, ML could be used in conjunction with other categories of AI to allow C2 systems to learn how to perform tasks when expert knowledge is not available or when optimal tactics, techniques, and procedures (TTPs) are unknown. These

---

[2] Stuart J. Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed., Upper Saddle River, N.J.: Prentice Hall, 2009.

technologies have considerable potential to enhance JADC2. But to transition them, the Air Force must first invest in an ecosystem to enable them.

Our initial research approach leveraged existing AOC processes, TTPs, manning numbers, and mission threads to explore opportunities for AI and ML (the companion report contains details). To take a more forward-thinking approach that is not tied to existing C2 processes and structures, the research team first identified the warfighting CONOPS that are needed for peer conflicts and then examined the C2 needed to enable those CONOPS. Recent wargames within the Air Force and among the services have suggested new MD CONOPS that will purportedly provide cost-effective means of meeting military campaign objectives. We started with these MD CONOPS and examined the C2 capabilities that are needed to enable them.

To assess the suitability of AI/ML to C2 forces across multiple domains, we developed three MD vignettes—Suppression of Enemy Air Defenses (SEAD), humanitarian assistance and disaster relief (HADR), and proliferated mesh ISR—focusing on the overarching functions and data needs for each scenario.

The vignettes are meant to serve as MDO case studies that help illustrate how AI/ML can be leveraged in a single mission. The companion report provides a reference of the current C2 processes based on the AOC ATC for the three vignettes. The modernized processes were developed for a time frame ten years from now, agnostic to the existence of an AOC building and with the assumption of a more distributed C2 structure. The purposes of designing this modernized C2 framework are to illustrate how the process can be enabled through AI/ML, determine the various data flows and AI approaches that would need to be captured and developed, and make recommendations on the various investments that would be necessary.[3]

In the remainder of this chapter, we briefly summarize the key points from each vignette about applications of AI/ML to JADC2. We also touch on how AI/ML might enable two larger C2 capabilities that are common to all operations, an MD operational picture and operational assessment (OA). The companion report contains a detailed discussion of both and the vignettes.

Suppression of Enemy Air Defenses Vignette

The first vignette—suppressing or destroying an air defense threat, such as a surface-to-air missile (SAM) system—is a very traditional Air Force and, indeed, joint mission.[4] However, we construct an operational concept, based on the Training and Doctrine Command and ACC MDO wargames mentioned in Chapter 1, that emphasizes the need for MD operations by incorporating capabilities from many domains: air-, land-, and sea-based; space; electromagnetic spectrum; and cyber. Theoretically, all of these capabilities are currently available; however, it is very

---

[3] In many instances, simple automation may be sufficient.

[4] Joint Publication 3-01, *Countering Air and Missile Threats*, Washington, D.C.: Joint Chiefs of Staff, April 21, 2017, validated May 2, 2018.

unlikely that they could dynamically be used together in the way described in this report, let alone in the timeline necessary to prosecute a very mobile system, such as a modern SAM.[5] Table 3.1 shows a list of exemplar domain contributors to the Air Force kill chain to suppress or destroy a modern SAM.[6]

Table 3.1. Exemplar Domain Contributors to SEAD Vignette

Domain	Detect (Find)	Track and Identify (Fix)	Strike (Engage)	Support
Air	Airborne ISR (e.g., rivet joint)	Nontraditional ISR (e.g., F-35)		• Decoys to stimulate radars • EA-18G standoff jamming
Space	Overhead ISR			• Navigation and timing • Beyond-line-of-sight communications (e.g., from ISR to AOC)
Cyber	Integrated air defense system network monitoring			• Delay integrated air defense system tracks sent to SAMs • Insert false targets
Land			Long-range fires (e.g., guided multiple launch rocket system)	

Modernized C2 Model for Suppression of Enemy Air Defenses

For this vignette, most of the early planning simply focuses on ensuring that the needed resources are available at the right time, whereas the DT process ensures that those resources are actually used. Thus, for our prospective future look, we focus on this DT process during mission execution.

Figure 3.1 maps out a possible highly automated engagement sequence for our SEAD vignette, starting in the upper left with initial ISR cues and ending in the lower right with assessment.[7] From the top, the first row focuses on target detection and COA selection, the middle row is the resource-selection process to perform the COA, and the bottom row details mission execution. Red rectangles and text indicate processes that appear suitable for the

---

[5] Timelines vary by system, operator skill, and mode of employment, but, typically, modern SAMs can activate their tracking radar, engage, tear down, and begin relocating after five to 30 minutes (Jane's, "S-300V," *Land Warfare Platforms: Artillery and Air Defence*, last updated October 7, 2019).

[6] Adam J. Hebert, "Compressing the Kill Chain," *Air Force Magazine*, June 18, 2008.

[7] Here, *highly automated* refers to the digital nature of the data and the machine-to-machine transference of information. We examine opportunities for using AI in these C2 processes in this section.

application of AI, while the black parallelograms highlight processes in which it seems likely that humans would need to remain involved. Black rectangles show external inputs to the process, such as ISR sources or other domain capabilities.

Figure 3.1. Modernized C2 for Multidomain Suppression of Enemy Air Defenses
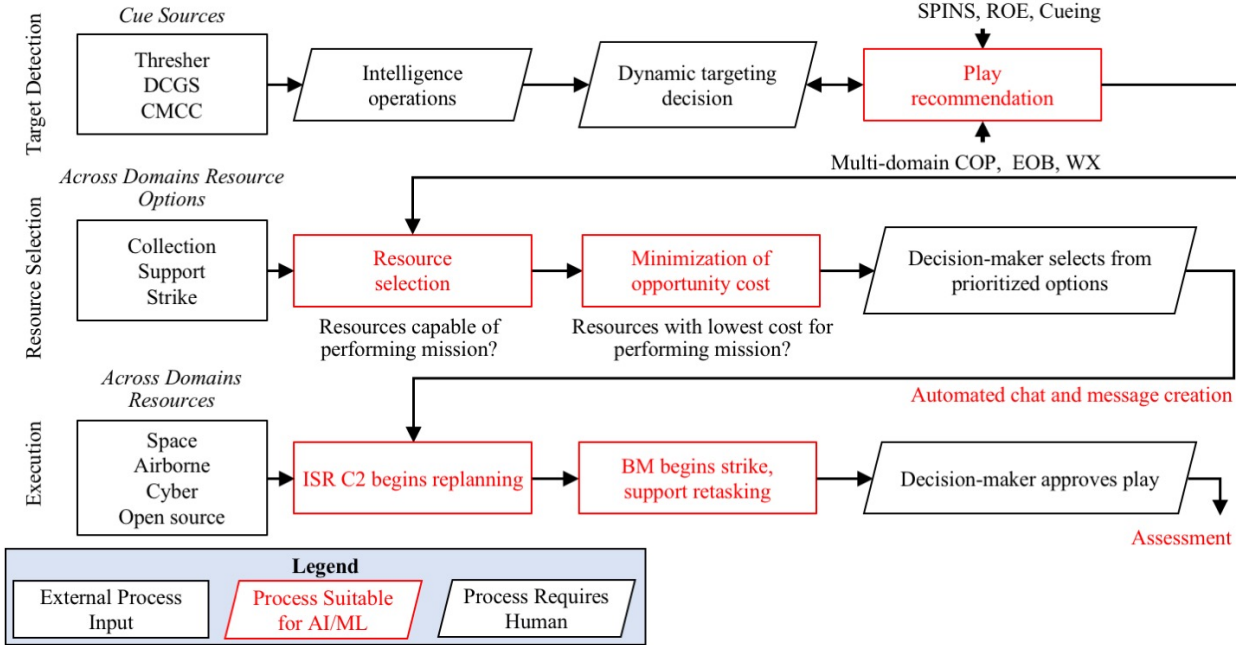
NOTE: BM = battle management, DCGS = Distributed Common Ground System, EOB = electronic order of battle, ROE = rules of engagement, WX = weather.

Beginning across the top row, many potential ISR sources can detect emitting targets, such as SAMs.[8] A decisionmaker must choose whether he or she wants to allocate resources based on this emergent target information. If action is taken, an automated play recommendation is provided. The middle row involves locating specific capabilities to resource the selected play. The final row employs the selected resources.

AI Approaches for Automated C2 Steps

Identifying the seven possibilities for automation (i.e., the red text or boxes in Figure 3.1) allowed us to assess the most-appropriate techniques to support them. We based our assessment on current or near-future capabilities, which, given the rapid advances in the AI field, may underestimate future capabilities. However, it also seems certain that further progress may be

---

[8] These ISR efforts may need to be coordinated with air-launched decoys that are stimulating the air defenses to engage.

more limited or might slow down in some areas (the companion report outlines each of the approaches):

- **Automated play**: Possible approaches include an expert system, modeling and simulation of a variety of combat scenarios, or the use of generative adversarial networks.

- **Automated resource selection from Blue common operating picture (COP)**: Techniques, such as NLP, will likely be necessary to incorporate information that is not available directly machine-to-machine or data that are not tagged appropriately (such as a PowerPoint briefing detailing a scheme of maneuver or chat message with an updated cyber effect timeline).

- **Automated minimization of opportunity cost**: Identifying the "best" resources for a task remains a challenge. Some researchers are examining the challenge of how best to compare the cost of using various resources by implementing *virtual liaisons* and a *capability marketplace*.[9] The marketplace matches consumers (those who need to accomplish a mission) with suppliers (those with such capabilities as a sensor or weapon) who use the virtual liaisons to offer their services at a price.[10]

- **ML**: Modeling and simulation and real-world exercises, such as Red Flag, could generate training data for algorithms that would likely resemble those being developed to play real-time strategy games, such as StarCraft.[11]

- **Automated chat and message creation**: NLP trained using data from current military chat rooms should be able to easily generate automatic tasking messages and even likely respond to simple clarification queries.[12]

- **Automated replanning of ISR and strike retasking by ISR C2/battle management**: Optimizing route planners that minimize exposure to threats and myriad other costs, such as fuel, are a well-developed area and could operate in conjunction across platforms to allow mutual support.

- **Assessment**: This may be one area most ripe for application of AI techniques, including image recognition algorithms, NLP, and expert systems or unsupervised learning.

---

[9] Craig Lawrence, "Adapting Cross-Domain Kill-Webs (ACK): A Framework for Decentralized Control of Multi-Domain Mosaic Warfare," Strategic Technology Office, Defense Advanced Research Projects Agency PDF briefing slides, July 27, 2018.

[10] The difficult part is calculating the price and the units to use for it. However, each supplier should theoretically know what tasks it is performing and when, and the priority of those tasks should be known from the planning process. So, theoretically, this price can be calculated as a function of time.

[11] AlphaStar team, "AlphaStar: Mastering the Real-Time Strategy Game StarCraft II," deepmind.com, January 24, 2019.

[12] Extended two-way conversations are a much more difficult problem, but those would likely not be necessary here.

Recommendations on Investments

Algorithms, curated data, and computing infrastructure are the three key capabilities that are needed to support introduction of AI techniques:

- **Algorithms.** Some key issues raised by applications for the JADC2 include the concern that almost all algorithm development is currently driven by commercial demands and academic interests, whereas the needs of the military will likely require more complex algorithms, larger amounts of data, and greater user expertise. A simple example would be an image-recognition algorithm with some layers of the network tuned to recognize features common to such objects as people, animals, and cars, as seen in photographs. If this network is then trained to recognize military vehicles in satellite imagery, performance can remain poor despite seemingly sufficient data.
- **Curated data** are likely to be one of the major investment requirements for AI approaches to this vignette and JADC2 as a whole; however real-world data are lacking, and simulation data risk being inapplicable.
- **Infrastructure** is unlikely to be a unique need for this vignette, although the need for MLS could be quite challenging here; security policy issues will need to be resolved with collaboration between the Air Force Chief Data Office (AF CDO), the Air Force Chief Information Office (SAF/CIO/A6), or possibly DoD's Joint Artificial Intelligence Center (JAIC) and various classification authorities.

Humanitarian Assistance and Disaster Relief Vignette

The HADR vignette was modeled after Operation Tomodachi, the U.S. Pacific Command–led operation to support Japan following the 2011 Great East Japan earthquake and tsunami and subsequent Fukushima Daiichi nuclear disaster, and Operation Pacific Passage, the U.S. Northern Command–led operation to evacuate more than 7,000 U.S. military families from Japan after the disaster.[13] Our study vignette is a noncombatant evacuation operation mission in a foreign country during a similar nuclear disaster. Table 3.2 shows example contributors for lift, ISR, and search and rescue by domain.

---

[13] U.S. Forces Japan, *Baseline Special Instructions (SPINs)*, Version 6, undated; Jennifer D. P. Moroney, Stephanie Pezard, Laurel E. Miller, Jeffrey Engstrom, and Abby Doll, Stephanie Pezard, Laurel E. Miller, Jeffrey Engstrom, and Abby Doll, *Lessons from Department of Defense Disaster Relief Efforts in the Asia-Pacific Region*, Santa Monica, Calif.: RAND Corporation, RR-146-OSD, 2013; Under Secretary of Defense for Acquisition, Technology and Logistics, *Lessons Learned/Strengths: Operation Tomodachi / Operation Pacific Passage*, undated.

Table 3.2. Exemplar Domain Contributors to HADR Vignette

Domain	Lift	ISR	Search and Rescue
Air	Inter- and intratheater	Global Hawk	Rotary wing
Space		commercial and national technical means	
Cyber		Open-source intelligence	
Maritime			Rotary wing
Land	Ground transport		

Modernized C2 Model for HADR

We developed a modernized C2 model for this vignette that is agnostic to the existence of an AOC.[14] The new process can generally be described in three phases, which are depicted as the three rows in Figure 3.2: target identification (in this case, evacuees) and COA selection, target and resource matching for the COA, and execution of the operation. Red text indicates products or processes that are not currently in existence or not currently automated. Black parallelograms indicate processes in which personnel will still need to make decisions, while black rectangles are inputs into the process. The companion report contains a detailed description of the modernized process.

Figure 3.2. Modernized C2 Process for a HADR Operation
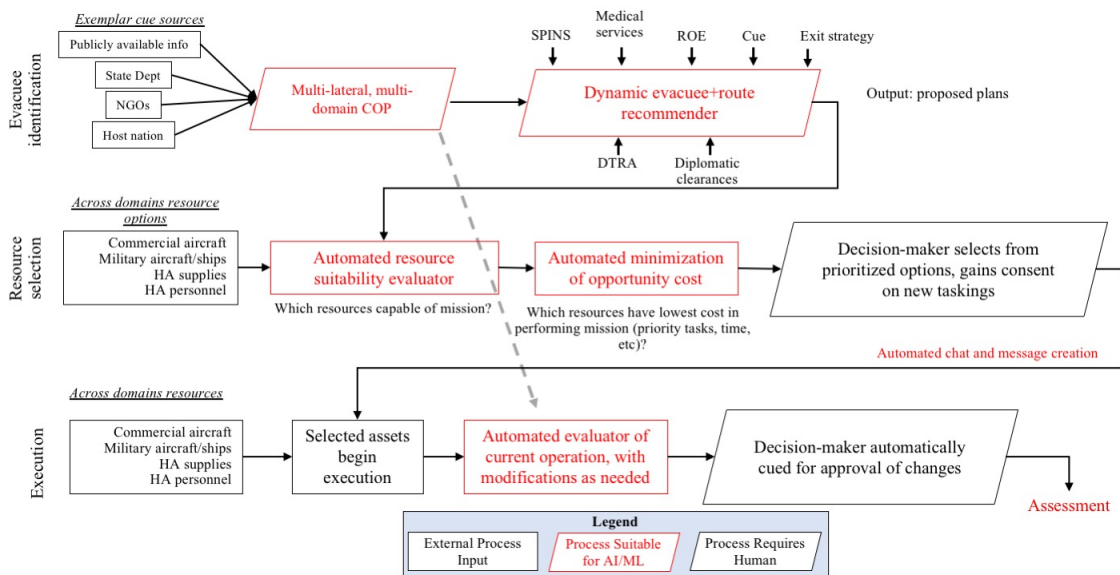
NOTE: DTRA = Defense Threat Reduction Agency, HA = humanitarian assistance, NGO = nongovernmental organization, ROE = rules of engagement.

---

[14] Lingel et al., 2020, contains information about the current C2 processes for this vignette.

AI Approaches for Automated C2 Steps

The crisis-management planning and execution of HADR touches on many processes that are similar to the SEAD mission—situational awareness of where different parties are in the area of interest and the movement of these parties, assigning resources to mission needs, and monitoring execution of operations. Automation opportunities identified for this vignette were multilateral, MD COP, dynamic evacuee and route planner, resource-suitability evaluator, automated minimization of opportunity cost, automated chat and message creation, and automated evaluator of current operation. The companion report contains details about each opportunity.

Recommendations on Investments

Some investments are currently underway for HADR operations. JAIC was established to help with national mission initiatives, such as humanitarian aid. Currently, the JAIC is assisting with Defense Support of Civil Authorities for the U.S. Department of Homeland Security. This involves support for such domestic events as natural disasters. In 2021, JAIC will expand to international efforts.[15]

There are three functional areas in Defense Support of Civil Authorities operations that could benefit from AI:

- search and discovery: searching for people in need of care during a disaster and determining whether infrastructure is functional, damaged, or denied
- resource allocation: determining where to send resources and relief supplies in a disaster zone
- execution of relief, rescue, and recovery operations: conducting the operation and identifying obstacles and opportunities along the way.

Two JAIC products currently in development—Fireline and Floodline—use computer vision models for sensor videos to automatically identify fire or flood boundaries on ArcGIS (Architecture Geographic Information System) maps.[16] There are also plans to use NLP to read data coming from various sources, such as hospitals, social media, and the Federal Emergency Management Agency during disasters. Data infrastructure is already in place for HADR missions through the All Partners Access Network (APAN), formerly the Asia-Pacific Area Network. APAN is a DoD information-sharing website that allows for communication and collaboration among various multinational organizations, partners, governments, and military services as a single touch point for those who might not have access to restricted networks.

Although APAN served an important purpose during Operation Tomodachi, it suffered from a number of problems (contained in the companion report). The Air Force can therefore help

---

[15] Based on discussions with JAIC staff. Also, see DoD, "JAIC and DSTA Forge Technology Collaboration," press release, Singapore, June 27, 2019.

[16] Based on discussions with JAIC personnel on June 12, 2019.

guide investment in making APAN more robust and accessible—for example, by requesting increased computing power and options for MLS filters. It will also need to incorporate a wide variety of data feeds and provide a less detailed version of the multilateral MD COP for all users. Such information as availability and readiness of personnel, supplies, and commercial and military aircraft and vessels needs to be input into any modernized HADR process. This includes personnel Air Force Specialty Codes, aircraft mission-capability status, and inventory of radiation protection and relief supplies.

Proliferated ISR Vignette

The next vignette is not so much a mission as a new type of force that will likely require a new type of C2 and battle management: the operation of a large number (on the order of hundreds) of fairly small, reusable, unmanned ISR platforms (referred to as a *small unmanned aircraft system* [sUAS]) operating in close proximity to one another to provide detailed surveillance of an area of about 10,000 square kilometers.[17] Figure 3.3 shows an example of such an operation over the Taiwan Strait. Because of the large number of platforms (blue triangular icons), vehicle control and sensor operation would be performed autonomously, while only fused target detections from multiple platforms would trigger close of the kill chain.[18] Fairly large numbers of these simple, low-cost platforms might be destroyed by air defenses, so replacement aircraft would need to be launched from distributed operating locations to replace them.

---

[17] Thomas Hamilton and Dave Ochmanek, unpublished RAND Corporation research, 2019.

[18] Such groups are often referred to as *swarms*, but we prefer to reserve this term to describe a type of behavior. Swarming behavior (as seen with some birds and fish, for instance) is not necessary to this concept. Here, we use the term *infestation*, since it better captures the type of behavior we are looking for. Another useful analogy to the desired environment is walking onto the floor of a Las Vegas casino. You know you are likely to be watched by a large number of surveillance cameras, but you do not know exactly which ones are looking at you or when.

Figure 3.3. Proliferated ISR Employment

Since this operational concept is not one currently employed by the AOC, there is no obvious "today's approach" to C2. The essence of this approach does not naturally fit into the air-targeting cycle (much as today's unmanned aircraft initially did not) and raises complicated issues with rules of engagement, airspace deconfliction, and the DT process.

Modernized C2 Model for Proliferated ISR

Planning and execution for this vignette touch on many similar processes already seen in the ISR division of today's AOC. Thus, the future AI-enabled mission flow in Figure 3.4 would apply to many aspects of the collection-planning process as a whole. Key differences are likely to be found in mission execution. For this vignette, most of the early planning focuses on determining whether a proliferated ISR system is appropriate, planning its operation, and then collecting and acting on its information and making sure it is performing well during mission execution.[19]

Figure 3.4 maps out a planning and execution process for proliferated ISR, starting in the upper left with initial ISR objectives and ending in the lower right with mission execution. From the top, the first row focuses on collector selection, the middle row is mission planning, and the

---

[19] The companion report provides a more detailed discussion of key steps that appear necessary in planning and then operating this concept.

bottom row details mission execution. Red rectangles and text indicate processes that appear promising to automation, while black parallelograms highlight processes in which it seems likely that humans would need to remain involved. Black rectangles show external inputs or outputs for the process, such as from higher headquarters or launch elements.

Figure 3.4. Modernized C2 for Proliferated ISR Employment

NOTE: BLOS = beyond line of sight, EEI = essential elements of information, exfil = exfiltration, ID = identification, INT = intelligence, NIIRS = n National Imagery Interpretability Rating Scale, OS = operating system, params = parameter, ROE = rules of engagement, reqts = requests, tgt = target, TNL = target nomination list.

AI Approaches for Automated C2 Steps

Automation opportunities identified for this vignette were automated collector recommendation; clearing airspace for ingress, area of interest, and egress; creating and disseminating a launch and recovery schedule; allocating beyond-line-of-sight bandwidth, alerting data receivers; creating and disseminating sUAS parameters during planning and execution; adjusting mission parameters; adjusting target parameters; and updating the operating system for airborne sUAS. Each opportunity is described in detail in the companion report. Creation and dissemination of sUAS parameters during planning and execution is the most unique task here, as there is no easily analogous situation today. In some ways, what is needed is a mission plan, such as that generated by the Joint Mission Planning System, but for hundreds of identical aircraft. The plans will not include exact flight paths in the mission area, because those will be determined autonomously in cooperation with other nearby sUAS, but they do need to include detailed technical parameters controlling the functioning of the automatic target-

recognition system. Automatically generating these data elements does not appear technically challenging; the key issues might revolve around policy, such as which C2 node should generate this plan (possibly a distributed node, such as a deployed-wing OC, or a centralized unit, such the 9th Reconnaissance Wing in the United States). How the plan is integrated with other C2 products, such as tasking orders, and who would have authority during mission execution—a Combat Operations Division [COD], a dedicated mission-control element as with an unmanned aircraft system today, or possibly a more ISR-focused OC, such as the CMCC—should be addressed.

Investment Recommendations

To enable proliferated sUAS operations, algorithms for cooperative flight and sensor tasking would have to be developed and tested. AI techniques are needed to teach such intermediate-level behaviors as threat avoidance, providing multiple sensor view angles, and best positioning for providing a missile guidance datalink.[20] Automatic target recognition would also be needed, such as what is already under development by Project Maven.

Most of the training-data needs are likely to be more receptive to synthetic generation than for our SEAD mission because there is less Blue-versus-Red move and countermove in our mission. The ability of the sUAS "infestation" to ignore attrition also enables less concern about minimizing risk with precise tactics. This only plays well with AI techniques that are likely to work fairly well most of the time. There is a significant amount of physical infrastructure needed to support this proliferated ISR architecture (such as launch and recovery elements), but the computational infrastructure is not likely to be stressing. The main challenge is likely to arise if individual air vehicles (or clusters of them) are unable to support target-recognition tasks with onboard computing and must be offloaded. However, sufficient resilient communication is likely to be the main concern because even today's facilities (e.g., the distributed common ground system) can provide significant computing resources to such an ISR task.

Common Themes Across the Three Vignettes

Although the vignettes are simplified abstractions of highly complex MD operations, they help illustrate various needs to enable JADC2. Moreover, although each vignette is unique, some common tools could be beneficial to multiple missions. Table 3.3 summarizes the various opportunities for automation or AI/ML across our three vignettes (i.e., the red text in Figures 3.2,

---

[20] Cooperative aircraft techniques have been studied for some time. For an early example, see John S. Bellingham, Michael Tillerson, Mehdi Alighanbari, and Jonathan P. How, "Cooperative Path Planning for Multiple UAVs in Dynamic and Uncertain Environments," *Proceedings of the 41st IEEE Conference on Decision and Control*, Las Vegas, Nev., December 2002. For use of more-modern techniques, see Nicholas Ernest, David Carroll, Corey Schumacher, Matthew Clark, Kelly Cohen, and Gene Lee, "Genetic Fuzzy Based Artificial Intelligence for Unmanned Combat Aerial Vehicle Control in Simulated Air Combat Missions," *Journal of Defense Management*, Vol. 6, No. 1, 2016.

3.3, and 3.5 as described in the modernized C2 processes for each vignette). Because this summary is limited to three vignettes, it does not cover the wide variety of MD missions in which the Air Force may participate, but some common themes were found. For example, there are several C2 tasks or tools that would benefit both the SEAD and HADR vignette, such as an all-domain COP or automated resource identification; likewise, there are some tools that would benefit both the SEAD and proliferated ISR missions. All three missions could benefit from an automated airspace-clearance tool. This exercise could be broadened to include other MD missions to encourage prioritization of new tools, data standardization and storage, and new algorithms.

Table 3.3. Modernized C2 Tasks Across Three MD Vignettes

C2 Task	SEAD Vignette	HADR Vignette	Proliferated ISR Vignette
Multilateral MD COP	x	x	
Automated play recommendation	x		x
Automated resource identification	x	x	
Automated resource selection (minimize opportunity cost)	x	x	
Automated airspace clearance	x	x	x
Automated chat and message creation	x	x	
Automated collection planning	x		x
Automated evacuee planning		x	
Automated strike and support planning	x		
Automated evaluator of current operation and cue for modification	x	x	

Many of the recent high-profile demonstrations of AI have involved ML. Another long-standing theme in AI research is the development of automated planners (APs) which, given a world model and set of decision rules, generate a plan to achieve the specified state. APs have been successfully applied to a large number of real-world problems. Indeed, many DARPA and AFRL programs that involve operational C2 (i.e., Joint Air/Ground Operations: Unified, Adaptive Replanning, Resilient Synchronized Planning and Assessment for the Contested Environment [RSPACE], and distributed C2 operations) involve some type of AP. The companion report contains more details on AP development.

To enable the JADC2 processes we just described to be successful, relevant data sources need to inform the tools. Table 3.4 lists the relevant data inputs that would need to be fed into the tools from Table 3.3 (the companion report includes more-detailed tables). Several sources are

duplicative for all three missions. This type of analysis about data sources can be done with additional mission vignettes to determine the most-critical data needs. In this effort, a similar approach was proposed for two C2 processes that cut across mission vignettes: joint all-domain COP and joint all domain OA. We briefly discuss them here (additional details are in the companion report).

Table 3.4. Modernized C2 Data Needs Across Three MD Vignettes

Data Sources and Input	SEAD Vignette	HADR Vignette	Proliferated ISR Vignette
Airspace	x	x	x
ATC tracking	x	x	
Blue COP	x	x	x
Blue force tracking	x		
Blue plan	x	x	x
Community of interest list	x	x	
Current airspace control order	x	x	x
Diplomatic clearances		x	
Down-selected cross-domain participants	x	x	
Electronic order of battle	x		x
Exit strategy		x	
Flight plans		x	
Host nation		x	
Intel cue	x	x	x
ISR collects	x	x	
Mapping		x	
Medical services		x	
Mission reports	x	x	
Nongovernmental organizations		x	
Objectives and desired effects	x	x	
Performance metrics	x	x	
Plume information from Defense Threat Reduction Agency		x	
Publicly available information	x	x	
Resources selected	x	x	x
Rules of engagement	x	x	x
Special instructions	x	x	x
U.S. Department of State		x	
Target location	x		x
Task prioritizations	x	x	x
Talon Thresher	x	x	
Weather	x	x	x

Joint All-Domain Common Operational Picture

An all-domain COP was a need identified in the clean-sheet vignette study to enable JADC2 and a recurring need expressed by AOC staff (the companion report contains more information). An all-domain battlefield picture could be empowered by ML for prediction, visualization of the operating environment across all domains, and leveraging algorithms to fuse data and provide confidence levels. Tools already exist for various operational pictures and situational awareness that may benefit from an infusion of AI/ML, advanced data analytics, and better visualization. Operational pictures and situational awareness tools are commonly used—for example, the joint COP at the geographic combatant command, the tactical air picture at the AOC, and the Talon Thresher tool—but, at the moment, they are all relatively simple amalgamations and visualizations of information across various databases. Details about the joint all-domain COP are in the companion report.

Joint All-Domain Operational Assessment

OA is a continuous process to measure progress toward the achievement of operational objectives and the desired end state.[21] OA begins with developing operational measures of effectiveness, quantifying their values, and then assessing the operation with support from other functional areas. Using the OA, the OA team provides "recommended adjustments to plans and guidance to achieve the desired end state conditions."[22] OA, in at least the regional AOCs, has mostly focused on the air domain. For MD operations, OA needs to expand its scope to include other domains. The major elements in a modernized OA process are automated information processing, analysis that includes forecasting (e.g., leveraging NLP), assessment (e.g., using expert systems), a joint all-domain operational picture, and automatic assessment reporting. OA processes would benefit from AI, with emphasis on reducing the time to collect data, giving more time for analysis and assessment. Analysis and assessment could also benefit from supporting AI/ML tools. The infrastructure for effective assessment is likely the largest challenge. A MLS solution should be incorporated to integrate data across security domains. A data lake and data pipelines, discussed in Chapter 4, would greatly benefit OA. Having standardized, structured, and curated data that can be stored—and through which analysis, ML, dashboarding, and assessment can be conducted—could potentially offer a more efficient assessment process and lead to better recommendations (the companion report contains a detailed discussion of the application of AI to OA).

---

[21] Joint Publication 5-0, *Joint Planning*, Washington, D.C.: Joint Chiefs of Staff, June 16, 2017.

[22] U.S. Air Force, 2016a.

Summary of Vignette-Driven Approach

Examining a series of MDO vignettes and then evaluating the C2 processes needed to enable them present opportunities to leverage AI/ML. This chapter presented three MDO vignettes: SEAD, HADR, and proliferated ISR. Additionally, there are C2 processes that cut across mission vignettes, which are important to capture for JADC2. This chapter highlights just two: COP and OA. Trends that emerge could help direct decisionmakers in prioritization across algorithm development for military JADC2. This approach further helps identify the data needed to enable these MDO CONOPS, as well as the necessary data architecture for machines to access them. The next chapter builds on this to address AI ecosystem needs for JADC2, commercial best practices, and the challenges with applying them to military operations.

4. Artificial Intelligence Ecosystem for Joint All-Domain Command and Control

One of the foundational drivers of JADC2 is to speed up processes in each domain such that observe, orient, decide, and act (OODA) loops are measured in minutes rather than days. Currently, the domains of air, space, and cyber operate at different battle rhythms (hours, weeks, months). Coordination among the three domains is performed via phone, emails, and meetings instead of through machine-to-machine communication, which further slows synchronicity. Machine-to-machine communication and AI-assisted decisionmaking are required to achieve the goal of battle rhythms of minutes. AI functions on large amounts of data: The more data that AI can access, the more successful its implementation will likely be, as long as corrupted data are not introduced.[1]

The cornerstone of an AI-enabled Air Force is data accessibility. The Air Force produces large amounts of data daily, but much of the data are siloed because of the lack of a unified data-management policy and insufficient IT. To foster AI, the Air Force will need to adopt a data-centric ecosystem and make relevant changes to IT, personnel, and policy, as we describe in this chapter.

This ecosystem begins with rethinking IT to store, share, and compute massive amounts of data. Specifically, such an ecosystem needs to handle data ingestion, provide advanced storage and computational resources, and host applications and tools that support decisionmakers.

A 2017 study revealed that companies leading in AI tend to be the most digitized (relative to their respective industries) and have invested in cloud and big-data infrastructure.[2] This form of data-driven ecosystem drives AI development and supports the machine-to-machine communication required by both AI and JADC2.

To begin to address how the Air Force can tackle these issues, we looked at the commercial sector—both the work conducted strictly there and what that sector does for DoD—to determine how it addresses these issues. We provide suggestions on how the Air Force could adapt these commercial best practices to its own needs. We then examine DoD's own in-house initiatives.

---

[1] Sydney J. Freedberg, "The Art of Command, the Science of AI," *Breaking Defense*, November 25, 2019.

[2] Jacques Bughin, Brian McCarthy, and Michael Chui, "A Survey of 3,000 Executives Reveals How Businesses Succeed with AI," *Harvard Business Review,* August 28, 2017.

Commercial Best Practices

Data Collection

Commercial applications commonly collect user data via a *Trojan horse* approach. These user applications typically fulfill a user need (shopping, searching the internet, checking email, photo sharing) at little to no cost. In return, companies collect data about users by tracking how they interact with the application and the content that they view, create, and share. Use tracking is ubiquitous, ranging from mobile apps to webpages to operating systems. For example, many websites use Google Analytics, which is a small piece of JavaScript code embedded in a website. This code is able to track where a user is coming from (before visiting the website), how many seconds the user spends on the website, and where the user goes after. Companies can use these data to continuously improve their product and improve marketing (e.g., identifying which ads are driving the most visitors) and website layout (e.g., identifying that most users follow an inefficient path to reach the checkout page).

KR could adopt this approach with its web apps. KR's first app, Chainsaw, assists airmen in the AOC with tanker planning. Although Chainsaw is currently the most-used KR app, the user data that pass through it are not currently saved. To move from an improved user experience to an environment that can leverage AI/ML, data should be captured, stored, and accessible for future algorithms. KR Trojan horses are a way to do this. Chainsaw provides a great test bed to prototype a data-collection system within the AOC. The saved user data might be used to develop a near-autonomous tanker planning tool that provides airmen with suggestions.

The commercial sector is also appearing to converge on the "collect everything" philosophy toward data. The premise is that some data streams may contain undiscovered correlations and that it is difficult to anticipate future data needs. The data requirements of JADC2 are ever-changing as JADC2 CONOPS are actively developed and tested. To be in a position to satisfy all future data requirements, the Air Force needs to adopt a "save everything" approach. For the AOC, saving all AOC products, including draft and intermediary products and chat-room logs, would represent such an approach.

Data Storage

The save everything approach is a relatively cheap policy to adopt because the cost of data storage has decreased significantly over the past decades, with the cost per gigabyte shrinking from millions of dollars to cents. The commercial sector is migrating toward data lake structures, which shift storage from expensive server rack storage solutions.[3] A data lake is a network of connected computers that provides storage and computational resources to form a central repository for data collection and processing. Data lakes provide readily available storage and are

---

[3] Wesley Leggette and Michael Factor, "The Future of Object Storage: From a Data Dump to a Data Lake," IBM webpage, December 3, 2019.

designed to handle large amounts of data, such as every ATO ever produced or every master air attack plan (MAAP) briefing an AOC has made. The distributed computing and storage that come with a data lake provide the scalability and flexibility to adapt to growing data needs. Distributing storage across multiple computers also provides faster access to data. The performance gains come from parallelization when large files are distributed into chunks across different computers.

We observed that most AOCs use Microsoft SharePoint to share data. The Air Force currently operates under a database-storage model in which data are stored and siloed on local computers. Siloed data eventually make their way onto SharePoint via network upload or even manual entry in some instances. This storage format is easy to use and implement, but it comes at the cost of scalability and flexibility, which are necessary for a data-driven AI ecosystem.

Data lakes would address many data-accessibility problems in the AOC because all of the data are stored in their raw format in a central location. Once the data are properly tagged and annotated, data lakes are easily searchable for specific pieces of information via search functionality or a data catalog. To that point, data lakes may eliminate the need for manual entry. On the COD floor, we observed that many airmen had to manually enter information into their computers because the information was digitally stored in a document of higher classification on Joint Worldwide Intelligence Communications Systems. For example, a set of coordinates at the SECRET level gets aggregated with TOP SECRET information during targeting effects team processes. After the aggregation occurs, the COD floor is unable to access the original coordinates, which would need to be hand-copied from the terminal (as illustrated in Figure 4.1). Therefore, the flow of information moves serially within an AOC, adopting the highest classification with which it becomes aggregated. Within a data lake, data are preserved in their raw form, and the original coordinates are directly accessible within the COD's SECRET-level network. The data classification tag ensures that data remain accessible at the appropriate classification level.

Figure 4.1. Data Pipelines Address Current Bottlenecks
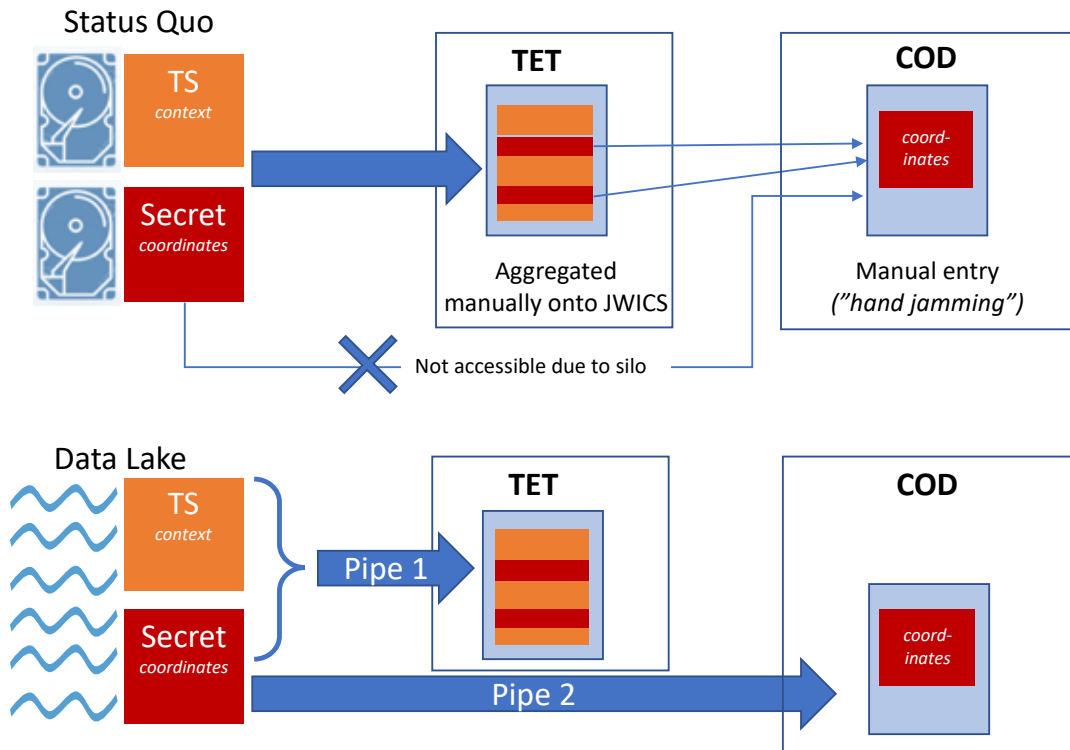
NOTE: JWICS = Joint Worldwide Intelligence Communications System, TET = targeting effects team.

Data Processing

Once data are collected, they must be cleaned: Data scientists spend the majority of their time cleaning data prior to analysis, a process with three common confounding issues. First, issues with a data set may not be readily apparent at first glance. When data scientists encounter additional errors during analysis, they must reclean the data. These feedback loops make it difficult to know what to clean unless one has some experience working with the data. Many companies retain data experts for their expertise with specific data sets. Second, data cleaning tends to require hundreds of small fixes, such as removing duplicate entries and typos, identifying unreliable entries, removing white space or extra zeros, imputing missing entries, converting timestamps, and making unit conversions. Finally, data cleaning is typically task specific, with no one-size-fits-all solution. Depending on the specific application, different types of data need to be joined (there are also different types of data-join algorithms). These three factors contribute to the disproportionate amount of time that data scientists spend on data cleaning.

AI in its current form is unlikely to replace humans in data cleaning. Automation is very helpful in this field, and data-cleaning tools can speed up the process by automating most of the processes. However, these tools are not autonomous and require column-by-column guidance by

a data scientist. More complex issues that arise during data cleaning as a result of unknown unknowns tend to be unique to each data set and therefore not well suited for such tools. Therefore, data scientists with experience in specific data sets continue to be invaluable in this field.

Another commercial best practice that we identified is the use of data-ingestion pipelines to automate the process of data cleaning.[4] Although there is no one-click AI solution to perform data cleaning, the commercial sector has greatly improved the process. This software-facilitated process is called *data ingestion*—pipelines that offer a way to encode a data scientist's knowledge of a data set by hardcoding the data's cleaning workflow. These pipelines represent a new framework for data cleaning and a shift away from the older approach of using Microsoft Excel. Although Microsoft Excel remains a popular data-cleaning tool, it offers neither transparency nor repeatability into the data-cleaning process. Furthermore, unless users are diligent in their use, raw data files can be overwritten by the clean versions. This may pose a problem if the raw data need to be processed differently for another application. The data pipeline offers a framework that addresses these issues by recording each data-cleaning step. Myriad free and commercial data-cleaning software programs exist and allow a user to use code, a point-and-click interface, or both. This framework allows data scientists to address all three issues described earlier by encoding all of the steps into a recipe that both preserves the raw data and enables transparency and peer-review of the data-cleaning process. This can be especially valuable during personnel turnover. For example, data scientists can work with subject-matter experts to create a recipe that identifies and removes duplicate entries for aircraft maintenance data. This recipe can be used to automatically process next month's data as they come in. Data-ingestion pipelines offer the ability to automate ingesting newer data of the same format (with little to no human input).

Data-ingestion pipelines offer a *do-once* approach to repetitive tasks. We observed that, during the Pacific Sentry 19-2 exercise, the MAAP cell spent two hours before each commander's briefing preparing presentations. Several commercial tools can automatically generate presentations from data streams using the data-pipeline approach. Joint integrated prioritized target list and MAAP briefings currently use PowerPoint slides displaying content in a standardized format. The data-pipeline approach can automatically generate slides for these briefings, and updating data streams will be automatically reflected at the time of briefing.

Within a data lake, data-ingestion pipelines require management. Pipeline management involves scheduling recurring jobs to process the data—identifying issues with incoming data—and to adapt to changing needs. Pipelines need to be scheduled to execute at certain times in accordance with a battle rhythm to generate the appropriate data products prior to a deadline or briefing time. These jobs provide inherent data quality checks because pipeline jobs typically fail

---

[4] Chiara Brocchi, Davide Grande, Kayvaun Rowshankish, Tamim Saleh, and Allen Weinberg, "Designing a Data Transformation That Delivers Value Right from the Start," McKinsey and Company, October 2018.

if unexpected issues occur with the data. Data scientists will need to continuously monitor data streams because they can change over time (e.g., switching vendors, change in data-collection policy, adding new data columns).

The ability to identify data issues quickly can help improve data-collection practices and serve as a driving force to improve data quality across the Air Force. Maintaining an organized data catalog and its associated data pipelines provides decisionmakers with the assurance that their data are free of corruption. Decisionmakers have differing needs and preferences for the data they consume. Duplicating and modifying an existing pipeline to tailor a briefing for a specific commander is faster and less error-prone than starting over.

Computation

The scalability and performance of computational resources needed to process data-pipeline jobs are similar to those for data storage. Fortunately, distributed computing also addresses this issue because parallelization improves both storage access times and processing speeds. Additional computers can be added, ad hoc, to a data lake network to provide additional computational resources. As previously mentioned, data are split into chunks across multiple computers in a data lake for both redundancy and performance benefits. A data lake leverages its network of computers to divide the workload into smaller pieces for parallelization.

A computational framework designed to process big data in a distributed computing environment exists today. The framework divides computation into smaller subtasks by grouping similar computations. These smaller computational tasks are distributed across the network for processing. This method has the benefit of having a smaller memory footprint, easier computation, and crash tolerance. The results of the computations are also stored in chunks across computers, preventing the need to load the entire large data set into the memory of a single computer. In contrast, trying to load and process a large Excel spreadsheet will be very slow on most AOC computers. The framework is based on the philosophy that storage is cheap and computational cycles are expensive. In exchange for relatively easy computational tasks in the reduce phase, large amounts of intermediate data are generated in the map phase (when filtering and sorting occur). The bottom line is that distributed computing and storage go hand-in-hand and transitioning to data lakes is the entry point to obtain these benefits.[5]

AI algorithm development, which uses deep (very large) neural networks, may require specialty computational resources in excess of the requirements for big data processing. In these instances, data lakes are used to process the data into a usable format that is stored on a separate database. This database typically resides on a platform with access to specialty hardware, which typically includes several graphics processing units (GPUs). Because of their design, GPUs have demonstrated speeds that are orders of magnitude faster when used to develop AI algorithms.

---

[5] Jure Leskovec, Anand Rajaraman, and Jeff Ullman, *Mining of Massive Datasets*, 2nd ed., Cambridge, UK: Cambridge University Press, 2014.

Briefly, GPUs sacrifice accuracy for speed when computing matrix operations by truncating numbers. Outside of small problems,[6] it is necessary in most cases to use GPUs to develop neural network AI algorithms.

Data Environments and Platforms

Although data lakes are inherently an abstract framework, most implementations involve the Apache Hadoop suite of software and systems. Hadoop is open source and allows anyone to create a data lake in-house. Alternatively, data lakes are accessible via cloud platforms. Most commercial tools in this field offer hardware services with Hadoop and proprietary management software installed to simplify the management of data, data lakes, and data pipelines. These services are referred to as *platform as a service* (PaaS). All hardware is owned, maintained, and upgraded by the commercial vendor. Users purchase cloud access to a preconfigured data lake on this hardware.

To implement a data lake, the Air Force has the option to develop one in-house or to purchase cloud access. The in-house approach would give the Air Force complete control of its data and data policies. Furthermore, the Air Force can leverage the benefits of globally redundant storage by having on-site hardware at each AOC. However, this option would require that the Air Force invest in high capital costs to purchase server racks and to build up all of the hardware and the associated electrical and networking infrastructure. Another necessary consideration is that the interconnected nature of data lakes creates potential cyber vulnerabilities. Furthermore, in-house approaches are typically less robust during surge demands: Storage or computational resources cannot be increased on-demand because that would require additional servers to be ordered, configured, and physically installed. In contrast, a cloud-based approach is beneficial because it incurs lower capital costs up-front and allows the Air Force to instantaneously scale up servers to meet demands. This approach also alleviates strain on IT personnel because all of the hardware, security, and networking issues are managed by the vendor. The data lake resides entirely on the cloud, but once data are processed and cleaned, they are sent to relevant local databases in each AOC.[7]

Once a data lake and its associated infrastructure are established, the Air Force would be optimally positioned to leverage AI on its data. The in-house solution would be to recruit AI talent to develop AI models or to select a company or organization to be a "trusted agent" that is independent from the commercial providers. Popular open-source AI software packages—called *deep learning frameworks*—include TensorFlow, PyTorch, Caffe, and Theano. All of these

---

[6] Outside most classification or regression problems (where the neural networks tend to be smaller), it is now standard practice to use deep neural networks to tackle such problems as image recognition; voice transcription, recognition, and translation; and advanced planning applications (such as playing chess or Go).

[7] When charting a path forward, the Air Force should carefully consider classification, cost, and security issues. In site visits, the team noted that the Air Force currently has many noninterconnected data lake systems in place. This shotgun approach achieves zero benefit of data lakes and is costing money.

frameworks allow users to develop, deploy, and maintain AI applications. Cutting-edge AI research and advancements are typically performed on these frameworks, and open-source implementations are readily available. These frameworks are powerful and allow fine-grained control, but this control comes at the cost of steep learning curves.

The Air Force has two options if chooses to purchase commercial data storage. AI software is available on many cloud platforms and is typically *accessible*—meaning that the AI software is usable with minimal training. Because of the GPU-heavy requirements for neural network training, vendors typically provide the hardware and allow users access to their AI software. These services are usually packaged as software as a service (SaaS). At the cost of being less feature-rich than deep-learning frameworks, this software provides easy-to-use point-and-click interfaces to develop AI models. The alternative commercial solution is to hire a consulting company that specializes in AI development: Many such companies exist, with AI subspecialties ranging from image analysis to text analysis to predictive modeling.

The Developing U.S. Department of Defense Artificial Intelligence Ecosystem

Aggregating data into a common cloud enables warfighters to process information, make predictions, and react at a much faster pace than would otherwise occur. The recent DoD request for proposals called JEDI is seeking a single vendor to provide DoD-wide cloud services over a ten-year period as an infrastructure as a service (IaaS) and a PaaS.[8] Through this contract, the Deputy Secretary of Defense will make an investment into a unified cloud infrastructure to provide data storage and computing solutions. The JEDI contract was awarded to Microsoft, but Amazon (Amazon Web Services) has protested the award.[9] In March 2020, the court signaled one flaw in the award process, and DoD asked for 120 days to "reconsider certain aspects" of the award.[10] The current DoD JEDI contract is an important consideration for the Air Force as it pursues enablers for MDO.

The Air Force will likely adopt the cloud capabilities from the winner of the JEDI contract. Any localized or cloud-computing solutions acquired in the next few years may likely be superseded or replaced by JEDI when it comes online. The JEDI request for proposals notes that managing communication among multiple clouds introduces higher security risks than does maintaining a single cloud. As a result, the Deputy Secretary of Defense is seeking a single cloud solution for all of DoD. Across all DoD organizations, there are currently 500 individual efforts

---

[8] Amanda Macias, "Pentagon Will Not Award JEDI Cloud Contract Until New Defense Secretary Completes Review," CNBC webpage, August 9, 2019.

[9] Jay Greene and Aaron Gregg, "Amazon's Formal Challenge to Huge Pentagon Award Uses Videos That Mark Potential Influence Exerted by Trump," *Washington Post*, November 23, 2019.

[10] Aaron Gregg, "Pentagon Asks to Reconsider Part of the JEDI Cloud Decision After Amazon Protest," *Washington Post*, March 12, 2020.

to acquire cloud capabilities.[11] These disparate efforts detract from the key benefits of a cloud-computing environment. Specifically, having separate clouds causes compatibility and data-accessibility issues, which dilute the effectiveness of AI/ML and automation solutions.

Having cloud capabilities is just one step toward an AI-enabled JADC2. As stated earlier, the JEDI cloud will be an IaaS and PaaS. IaaS means that the vendor will build and maintain the servers, data storage, and networking. Virtualization software enables computing and data resources to be disaggregated and split into smaller virtual servers tailored toward any requirement. PaaS provides an additional layer on top of IaaS: Operating systems and software are managed at this layer to facilitate data management, data security, and application development.

The Air Force is responsible for developing and maintaining the software layer. And, going forward, the Air Force will need to be mindful of the unintended ramifications of earlier decisions and make course corrections as needed. For example, KR is the Air Force's AOC modernization solution and has demonstrated the value of Agile development and microservices architecture, quickly transforming ideas from a whiteboard into working applications in the AOC. The platform that supports these apps, Pivotal technology stack, is currently accredited to operate in SECRET-level computing environments with a continuous authority to operate (C-ATO). This represents a major shift in software development because these ATOs were traditionally granted only to individual pieces of software. The Pivotal technology stack is a software factory created to develop, host, and maintain software. The C-ATO is helpful for agile development because all Pivotal software and updates are automatically authorized. However, an unintended consequence of KR's success is that their C-ATO is currently the easiest path toward deploying any application within the AOC. The process of obtaining an authority to operate software is typically time-consuming and difficult: Many stakeholders need to be involved, and incentives are not aligned to encourage this process. This means that non-KR-C2 apps now rely on the Pivotal platform. For example, the C2 tools that the DARPA RSPACE program researched and created for the AOC are currently in limbo because they cannot be fielded unless KR rewrites them using Pivotal's platform.

This example demonstrates that the app development base should not be limited to KR alone: A more expedient C-ATO process may help. In addition to RSPACE, we identified other parties, such as AFWERX and a small team within the 609th AOC, that are also experimenting with modern C2 apps. For example, a 609th AOC team has set up a data-management platform called Kibana that automatically ingests, processes (after defining the data once), and disseminates U.S. Air Forces Central Command data. This team assists other airmen with data-processing issues and then automates the process for incoming future data. Users are able to process, visualize, and

---

[11] U.S. Congress, "Combined Congressional Report: 45-Day Report to Congress on JEDI Cloud Computing Services Request for Proposal and 60-Day Report to Congress on a Framework for all Department Entities to Acquire Cloud Computing Services," undated.

present the data in customizable dashboards, which can replace the need for traditional Excel or PowerPoint products. AFWERX is developing proof-of-concept technologies to demonstrate that the Air Force can transition away from the siloed SharePoint model and into a cloud-based data lake model with multilevel classification.

Multilevel Security (and Security Concerns More Broadly)

As just discussed, data access is difficult within the military C2 enterprise because of historical storage of data locally or in SharePoint folders on shared drives. The reliance on distinct classified networks and having data on these networks at different classification levels present additional challenges. Security policies have focused on protection of information, often at the expense of integration. Transporting data between different classifications requires an approved cross-domain solution.

The Air Force has begun to address the duality of security concerns for data protection and integrity and the need to share information in various ways. For example, the CMCC uses application programming interfaces (APIs) to convert data into a universal C2 interface standard for sharing data across different mission-level systems.[12] In fact, an October 2018 memo cosigned by Lt Gen Robert McMurry (Commander, Air Force Materiel Command) and Will Roper (Assistant Secretary of the Air Force for Acquisition, Technology and Logistics) requires that the entire Air Force use a Modular Open Systems Approach by implementing open mission systems and a universal C2 interface to the maximum extent possible.[13] The Shadow Operations Center at Nellis Air Force Base developed OneChat, a chat application that supports chat with one interface but is supported by an MLS database providing storage across multiple classifications. Having an MLS IaaS backbone, such as the one envisioned with JEDI, creates the opportunity to continuously monitor all data coming into the cloud and all data at rest for adherence to Security Classification Guides.

Modeling and Simulation to Generate Data

Although AOC data are collected and stored, the quantity of data may be insufficient to train AI algorithms. The lack of data consistency in exercises also limits training data sources. Modeling and simulations can provide additional sources of data. The scenario cases for peer adversary conflicts would add data beyond current counterinsurgency and peacetime operations. Live, virtual, and constructive exercises could provide additional data by which to train algorithms. Building out a sufficient training set would likely take all three approaches for

---

[12] U.S. Air Force, Virtual Distributed Laboratory, "Common Mission Control Center (CMCC)," September 27, 2017.

[13] Robert D. McMurry, "Use of Open Mission Systems/Universal Command and Control Interface," memorandum to the Air Force Program Executive Officers, Washington, D.C., October 9, 2018.

gathering a suitable data set. To maintain a realistic optimism, see Chapter 3 for a discussion of the limitations of these data sets.

Personnel Considerations

Although personnel considerations were not a focus of this report, the topic came up in discussions at site visits and interviewing other JADC2 stakeholders. Thoughtful plans for the manning and training of C2 staff for a future joint all-domain AI-enabled environment where trust must be built for AI are also a part of DoD's AI ecosystem development. Roles and staffing levels will change from the existing AOC model as the Air Force matures its IT infrastructure and data policies. JADC2 will require the new 13O Air Force career field for C2 to prepare airmen for tasking capabilities from all domains, implying some working-level knowledge of capabilities from these domains.[14] Furthermore, these operators will likely need additional technical skills for a highly human-machine teaming environment. The increase in machine-to-machine communication will automate the time-consuming processes of many current roles, freeing up operators to spend more time to assess and plan. The location of these C2 staff may change to a more distributed C2 construct.

Effective management and navigation of an AI ecosystem will also involve new roles and new educational requirements. There will likely be an emphasis on data stewards and data scientist positions to keep the data flowing smoothly. Although the Air Force will need teams of airmen and contractors to develop AI algorithms, day-to-day operations rely on operators who understand how to use the algorithms and are experienced with data visualization and metrics, and the transition to a JADC2 will require senior leaders across the services who understand the MD resources available to them.

Another significant challenge is the lack of human trust of AI systems.[15] Decisionmakers are unlikely to make optimal MDO decisions if they cannot trust the machine-generated risk assessments, recommendations, or estimates. One of the biggest concerns involving AI is the lack of visibility and opaqueness of its classification or decisionmaking process. Traditional regression algorithms, such as generalized linear models or decision trees, are preferred in many fields today because of their interpretability. The rising popularity of neural networks has made explainable AI an active and hot topic in the field. This is a difficult field because neural networks can have millions of parameters. It is difficult to trust the judgment of such a large algorithm that cannot explain its outputs.[16]

---

[14] Amy McCullough, "USAF Looks to Create New Command and Control Structure," *Air Force Magazine*, June 6, 2018.

[15] A common theme heard in our AOC site visits.

[16] Freedberg, 2019.

Work regarding explainable AI is currently highly application specific.[17] For example, significant progress has been made in explaining image-recognition systems because of their visual nature. Generative adversarial neural networks (GANs) are a class of AI algorithms that can be used to generate new images based on training data. Combining an image-recognition system with a GAN could result in a new system that can identify content within an image. This new system could help explain a classification result by generating images that it believes also belong in the same category.[18] Although GANs are heavily used for visual applications, they may be less useful for nonvisual AI applications. Other work in this field is focused on visualizing the neural network layers. This feature allows us to see which areas of an image the algorithm is "paying attention to" and can give us a better sense of what is driving a decision.

Explainable AI outside of image applications is currently an ongoing field. One promising field of research is with text generation, which has been applied by AI toward writing scientific papers, generating captions for media, and real-time translation. Microsoft is researching how to apply this technology to generate sentence-level explanations for recommendation systems.[19] AI can be relegated to recommendation engines with a decisionmaker in the loop for now. As work toward explainable AI improves over time, algorithms can slowly be given control of more tasks and decisions.

---

[17] David Gunning, *Explainable Artificial Intelligence (XAI)*, Arlington, Va.: Defense Advanced Research Projects Agency, 2016.

[18] Chaofan Chen, Oscar Li, Chaofan Tao, Alina Jade Barnett, Jonathan Su, and Cynthia Rudin, "*This* Looks Like *That*: Deep Learning for Interpretable Image Recognition," 33rd Conference on Neural Information Processing Systems, Vancouver, Canada, 2019.

[19] Jingyue Gao, Xiting Wang, Yasha Wang, and Xing Die, "Explainable Recommendation Through Attentive Multi-View Learning," presentation at the Association for the Advancement of Artificial Intelligence Conference on Artificial Intelligence, Arlington, Va., March 2019.

5. Research Conclusions and Recommendations

> [T]he Air Force must build an integrated network of air, space, and cyberspace-based sensors, as well as leverage joint contributions from all domains. This integrated network and architecture will enable more rapid and *effective decisions from the tactical to the operational level."* —U.S. Air Force, *Air Superiority 2030 Flight Plan: Enterprise Capability Collaboration Team*, May 2016b, p. 6.

Issues

Although capabilities from the domains of air, space, maritime, ground, and cyber all contribute to military operations today, the ease and speed of synthesizing information across them into a coherent understanding of the environment for commanding and controlling MD forces is severely constrained. Legacy networks and security policies, limited data management policy and standards, and industrial-age ways of organizing, training, and equipping cannot keep pace and thereby create operational difficulties and missed opportunities. Cultural and language differences between domain communities and dispersed authorities to resource and employ domain capabilities further complicate integration.

We live in a digital age in which commercial capabilities promise to dramatically shorten the OODA loop by collecting, processing, storing, and mining data and rapidly recommending actions based on those data. Given the increased complexity of MDO planning and the greater data requirements, the Air Force will require automation and new algorithms that include AI and ML. The industrial-age 72-hour ATC within the Air Force's AOC is incongruent with today's digital world. To realize its potential, the Air Force must overcome technological, data access, and cultural challenges. In short, the Air Force needs an ecosystem to support AI/ML.

Conducting dynamic planning at speed and scale is problematic for the AOC given the current weight of emphasis on deliberate planning processes. One can imagine a future campaign in which the balance between deliberate planning and dynamic planning is reversed. AI/ML provides the potential to improve integration across domains and increase decision quality at time scales called for by dynamic situations.

Conclusions

Migrating the AOC structure to a modern digital environment poses many challenges: reliance on human subject-matter expert meetings and boards, basic IT infrastructure, multiple classification-level data on air-gapped systems, requirements to use unique stove-piped C2 systems, little-to-no data tagging or cleaning, and heavy reliance on Microsoft Office products. Although there are various JADC2 efforts underway, their reach is not far enough, and it appears

that there is little coordination between the organizations and their specific projects. In the wake of the cancellation of the AOC 10.2 weapon system upgrade in 2017, KR has emerged as the dominant modernization effort within the AOC program office. Currently, KR is tied to the Pivotal cloud foundry in which airmen work with Pivotal programmers to develop apps, but the Air Force could create opportunities for a larger and more-diverse set of contractors and airmen to contribute to JADC2. Current MDC2 RDT&E funding further includes an experimentation campaign to identify new technological needs and an enterprise data lake to enable C2 application development. The former includes the new shadow operations center (ShOC), where OneChat was developed to allow chat information to be stored in an MLS database and shared through an API, creating a DaaS capability. Such foundational elements as MLS and APIs are critical to MD capabilities and future applications of AI/ML. Yet ShOC has been underfunded relative to the scope and ambition of its experimentation campaign.[1]

Notwithstanding their successes, these near-term modernization efforts face challenges. KR has digitized only a small fraction of AOC workflows, and most still rely on Microsoft Office products. The unstructured data they create are not amenable to automation, let alone AI or ML. Even when structured or semistructured data exist, KR applications do not store them, a necessary condition for AI/ML. Additionally, KR's software factory supports delivery of applications to systems at the collateral SECRET level, whereas MD operations occur across systems at multiple classification levels. Moreover, most AOC users, who are a major source of requirements for KR, neither have been part of a high-end conflict or MD operations nor are long-term users of KR products. **As a consequence, requirements may reflect current day-to-day needs rather than needs for a high-end future MD conflict**. Finally, mechanisms and prospects for transitioning ShOC and other JADC2 contributors' developments using KR's software factory are unclear.

Although far-term science and technology modernization efforts from AFRL and DARPA show promise, they also face challenges. Despite decades of research, few C2 technologies have transitioned to the AOC, and, of those that have transitioned, none have involved AI/ML.[2] This is in part because integrating new software with the AOC baseline is hard, given the aforementioned AOC challenges. This is exacerbated in the case of AI/ML systems that need to share interfaces with many applications to receive data for planning and execution. Further complicating integration, data pedigree and data sources, which are taken for granted during RDT&E, vary widely during actual operations. Culturally, operators and AOC commanders are not routinely exposed to AI/ML technologies. Introducing such technologies will require training to ensure that human operators have appropriately calibrated trust in the AI/ML systems.

---

[1] The companion report, not available to the general public, provides more details on funding levels and ShOC scope.

[2] Based on information from teleconference with KR and on-site visits to AOCs.

Unfortunately, "simply" solving today's data-management and access, system-interoperability, and decisionmaking problems will not enable greatly expanded MD operations. As a result, application of automation and AI alone to the ATC is likely a necessary, but not sufficient, upgrade. Our interviews, site visits, and literature review pinpointed at least four additional factors that limit the speed and scope of current MDO, issues that should be resolved before additional progress can be made:

1. **Authorities and command relationships**: However much the JFC and the air component would like to employ MDO, and however good they are at it, they often lack the authority to employ some capabilities. Just as the maritime component puts some, but not all, aircraft sorties under control of the JFACC, CYBERCOM and Strategic Command do not intend to transfer (and may be prevented by law from transferring) operational control of their units to the COCOM or a component. The supporting commander "aids, protects, complements, or sustains another commander's force, and who is responsible for providing the assistance required by the supported commander."[3] Direct-support relationships can be powerful but can also lack key authorities, such as control over prioritization. In additional to control, COCOMs also lack classification authority, such as the caveats protecting information about satellite constellations.

2. **Synchronizing battle rhythms across domains**: Unsurprisingly, different domains work on different timelines and with different planning horizons. Although ground and air forces may plan for tomorrow's coordinated maneuver, cyber forces may work for weeks or even years to gain access to an adversary's networks. Even if access is established, it can take weeks to grant the requested authority to use a specific tool. Similarly, although the availability and capability of space forces are predictable, it can take hours or longer to reprioritize assets, and it can take days or weeks to reconstitute assets lost to attrition. There are also simple practical problems—the SRD may have a MAAP prebrief at the same time of day as the cyber OC has a shift handover. Although these conflicts can be resolved (possibly using commercial AI-enabled scheduling software), they require compromise from all players. Today, these compromises are reached through human-to-human interaction, which limits both capacity and speed.

3. **Different C2 structures in different theaters and regions**: Although AOCs share a common baseline, the distinct geographic, functional, and coalitional constraints placed on each have driven their divergent evolutions. These differences prevent a one-size-fits-all solution for improved MDO or C2. If CYBERCOM or Strategic Command changed processes to improve integration with the 613th AOC in U.S. Indo-Pacific Command, there is little guarantee that those changes would benefit the U.S. Central Command 609th AOC, and they may, in fact, make things worse. Furthermore, it is unclear whether

---

[3] Joint Publication 3-0, *Joint Operations*, Washington, D.C.: Joint Chiefs of Staff, January 17, 2017, incorporating change 1, October 22, 2018.

the AOC construct will be the operational-level C2 construct of future JADC2. The AI/ML tools used for JADC2 would need to be considered in the context of each C2 center and its unique workflows. Likewise, automated schedule synchronizations and process negotiations would have to be "tuned" for each combination of organizations. Multiplied by an exponential combination of possible organizational processes across domains, the task is quite daunting.

4. **Robust and resilient communication systems and procedures**: A key enabler to the C2 of forces is the communications backbone and associated procedures that convey the commander's guidance, data sharing, planning, and execution monitoring. Distributed control efforts in the geographic commands may help address risks to communication capabilities in the face of adversary threats, but the concepts may also create new vulnerabilities. The need for common resilient communication systems and processes among operation centers for the various domains may be obvious; however, whether the requirements are established and the funding allocated is another matter.

These four issues have not yet been addressed for myriad reasons: (1) Each touches on multiple aspects of doctrine, organization, training, materiel, leadership and education, personnel, facilities and policy (DOTMLPF-P), and hence no single organization can resolve them alone, (2) the timescales for addressing each issue exceed a single leader's tenure, (3) some stakeholder organizations have little incentive to change, and the change may require resourcing, and (4) each change comes with major trade-offs, leading to honest disagreement about the best way forward. These difficulties are characteristic of most of DoD's "hard" problems, and there is no obvious quick fix. In general, progress is seen only through sustained leadership investment at the highest levels, across multiple personnel transfer cycles, in a continual pursuit of a small set of well-defined goals. The processes across stakeholder communities should be in place to sustain progress in the face of leadership and staff turnovers. For MDO, these goals could be defined from the bottom up, such as operationalizing MDO CONOPs for specific missions, or from the top down, such as specifying new requirements for a joint all-domain OC and resourcing the acquisition of capabilities to achieve the goals. Regardless of the approach, the current modernization process does not seem suited to make progress on the four issues highlighted.

In fact, these challenges exceed the scope of the Air Force. A true JADC2 capability should be agnostic of the domain(s) employed. Service biases may incline Air Force C2 staff to prefer air solutions when Army capabilities may be more cost-effective (in the broadest sense) to meet the military objective, for example. A truly MD workforce may likely necessitate a joint training program for C2 operators to plan and monitor the execution of operations across all domains. It may be daunting to reimagine the DOTMLPF-P changes to truly embrace MDO. Yet there appears to be a systematic approach for the Air Force to move forward in the near term.

Recommendations

Fundamentally, C2 modernization efforts should be anchored in a clearly defined strategy. The basis of the efforts should be the military's ability to effectively conduct synchronized MDO. The goal is enabling the warfighting effort to employ capabilities from different domains to meet a military objective. MDO CONOPS are percolating from futures wargames within and among services.[4] Think tanks are evaluating CONOPS against current ways of performing the missions. Promising ideas from this process would then set the vision both for operational- and tactical-level training and exercises, as well as define the needs for technological infrastructure and applications.

Three primary enabling categories should be aligned (and improved) to support future MD operations: the C2 construct, data, and algorithms (Figure 5.1). Recommendations pertaining to each of these categories are presented in the remainder of this chapter.

---

[4] For example, the recent Training and Doctrine Command and ACC MDO wargame series in 2018.

Figure 5.1. Multidomain Operations Drive Three Enablers
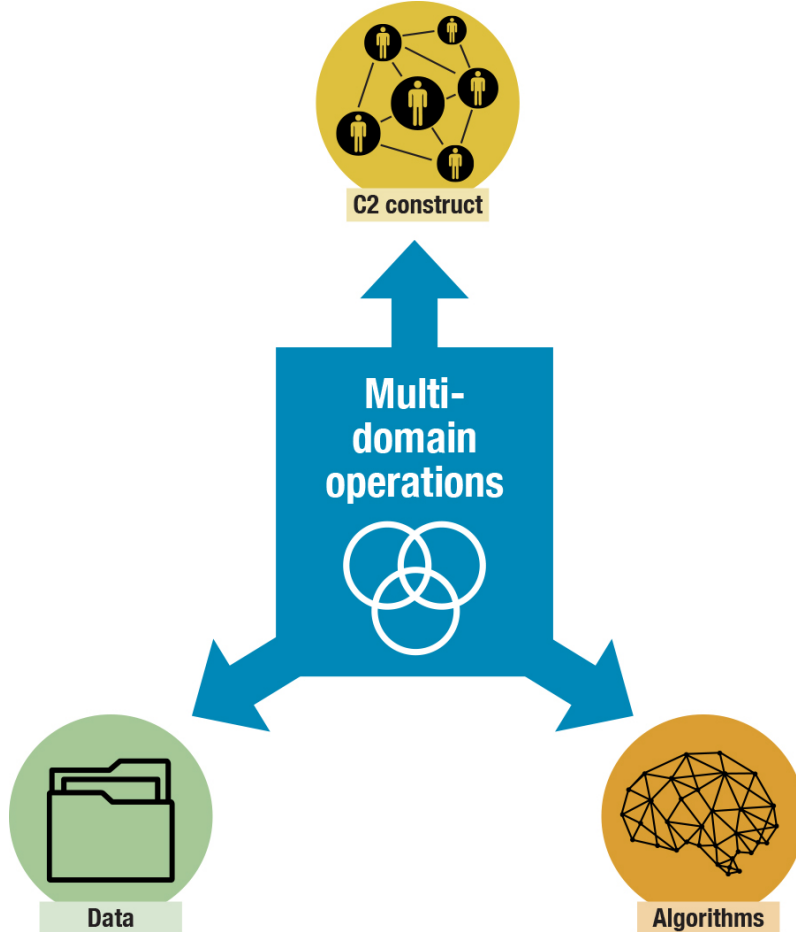
C2 Construct

The C2 construct includes how forces are organized, how they train to C2 missions for MDO, and how they are equipped to do so. The current AOC paradigm, characterized by a centralized facility staffed by subject-matter experts and convening boards and bureaus over the 72-hour ATC, is not likely to be the construct of the future. In both the European and the Pacific theaters, the Air Force is examining alternative ways of implementing centralized command and distributed control of forces. Key challenges to address include determining what authorities can be delegated to which echelon, how to transition to and from centralized C2 (i.e., defining what "distributed control" equates to), and training the people involved for MD force employment. The driving force is that the C2 construct of the future must support the operational needs of future MDO CONOPS. Since there are multiple concepts for a new distributed C2 architecture, the construct should be flexible enough to accommodate the variations as they mature.

The Air Force's *Air Superiority 2030 Flight Plan* calls for a campaign of experiments to address the most-challenging targets across multiple domains in a Defeat Agile Intelligent

Targets campaign of experiments.[5] The plan falls short of defining a campaign of experiments to address operational-level C2. Alternative C2 constructs should be tested and evaluated to understand the trade-offs in resources and risks across candidate constructs.

Data

The second supporting capability is the data necessary to C2 MD forces. The ongoing experimentation in AI by the Air Force is a promising start for developing an Air Force AI ecosystem to support air operations. Until JEDI deploys a DoD-wide cloud, the Air Force can ensure that it is well positioned to leverage cloud resources by pursuing the following steps.

Develop Data Standards and Processes for Storing and Sharing Data

The AOC generates many Microsoft Office files during an ATO cycle. Analyzing these data is possible only if they follow a standardized format, such as using templates or similar column names. A data-management policy should dictate how these products are archived to ensure their compatibility with future data pipelines. Existing SharePoint drives could be migrated to support this effort. Interviews also revealed that old files are deleted whenever SharePoint gets full. Clearly, the deletion of existing data is a problem for moving to an AI-enabled C2 that is hungry for data.

A critical part of the data-management policy includes tagging security streams. All incoming classified data should be properly tagged with their appropriate classification levels. Many data streams will contain information of various classifications. With the proper tagging, data pipelines can be established to separate data within these streams into separately classified data streams. Addressing these issues as soon as possible will make it easier to address MLS issues during implementation. This will also ensure that each AOC computer has access to the relevant data streams and enable machine-to-machine communication to eliminate manual entry of data between air-gapped computers.

We recognize that all AOCs differ in culture and in their workflow. Under a unified data-management policy, different AOCs can maintain their unique differences as long as the data they generate remain in a standardized and consistent format. When implementing a data lake, AOC-specific data pipelines can be generated to accommodate these differences.

The last aspect of an AI-focused data-management policy is to foster a "free-the-data" type of environment. There is a trade-off between operational efficiencies (and perhaps effectiveness) and information security. Information within the Air Force, other services, and intelligence agencies is siloed not only because of infrastructure issues but also because of policy challenges regarding the safe handling of classified information. Senior leadership can help "free the data" by emphasizing the need for integrated information to enable MDO and clearly articulating guidance on balancing this pressing need with security guidelines regarding need-to-know. Data

---

[5] U.S. Air Force, 2016b.

sharing is a slow process, especially when requests have to be approved by the originator, as in originator control authority.[6] There is heavy scrutiny to verify the need-to-know and what the data will be used for. Additionally, there is no incentive to undergo this process for data stakeholders. The AF CDO and Knowledge Management Capabilities Working Group (KMCWG) could develop overarching data-security policies in collaboration with other service and Intelligence Community authorities that encourage data sharing and collaboration in accordance with the initiative to digitize the Air Force. The KMCWG can then assist OCs in implementing the initiative.

In short, every AI/ML app will rely on data that may not currently be available because they are not currently collected, are imperfect, or are not accessible. Therefore, the Air Force needs to develop an Air Force–wide strategic plan for how the data will be collected, stored, secured, and shared. This effort includes setting in place data standards, authorities, integrity checks, and safeguards to protect against undetected intrusion. AF CDO would have primary responsibility for setting these standards and would work in concert with the SAF/CIO to balance safeguarding the data appropriately while ensuring they remain readily accessible as MDO CONOPS matures. An experimentation on standardization approach may allow refinement of guidance as capabilities evolve. A data-management policy is needed to ensure that all data are saved and properly tagged for machine accessibility. This includes tagging classification levels as described in Chapter 4.

Develop the Necessary Infrastructure

A major barrier toward creating and implementing a data-driven ecosystem lies in transitioning away from the current SharePoint framework. As mentioned in Chapter 1 and in the companion report, the AOC depends heavily on business services products for data processing and dissemination. These tools are not well suited for big-data analysis. They lack a data-processing workflow, which results in trapping a lot of the data within these files.

The Air Force should continue investments in cloud computing infrastructure and data lakes and ensure that they align with DoD JEDI moving forward. AF CDO and SAF/CIO—with support from KM—should be testing data practices and ensuring the compatibility of data standards with a cloud environment to facilitate the eventual migration into the cloud. Furthermore, the Air Force should consider granting C-ATOs additional software development platforms to support app building.

Update Training Requirements

Although many app builders promise easy-to-use interfaces, Air Force training will likely need to be modified to build workforce expertise in understanding data-analysis tools and

---

[6] *Originator control authority*, commonly referred to as *ORCON*, requires the release of the information to be approved by the originator of the information and can create a bottleneck for dissemination.

visualizations. New postmillennial ways, such as YouTube do-it-yourself videos, may be alternatives to traditional training approaches. The 13O C2 career field, if it in fact morphs into a JADC2 career field, will need to have a truly joint training and education plan to ensure all domains are understood in terms of force presentation and capabilities. The Air Force should work with sister services and the joint staff to realize the all-domain workforce.

Algorithms

The third key enabler for JADC2 comprises the tools, applications, and algorithms that leverage the data and IT infrastructure to greatly accelerate the necessary C2 processes. Science and technology organizations, such DARPA and AFRL, are actively developing new C2 approaches but without the guidance of well-defined MDO CONOPS to focus on. For this reason, the majority of efforts align with existing phases of the air-centric ATC. KR is addressing real-world operator "headaches" within the AOC with new applications. These applications produce real savings, but they do not address emerging C2 needs or provide new ways to employ MD forces. To truly realize MDO at scale and in a dramatically tightened OODA loop, AI algorithms will likely be needed.

Air Force Warfighting Integration Center (AFWIC) should prioritize AI algorithm development for JADC2, and SAF/Acquisition, Technology and Logistics (SAF/AQ) should select or establish an oversight organization to monitor progress: Currently, KR is the modernization branch of the AOC program office. Beyond KR, various other applications for JADC2 are being developed by a number of different groups (DARPA, small businesses, industry, academia, operational AOCs, AFRL, AFWERX, CMCC, and the like), but it is unclear whether any entity has oversight on all of the applications. A single organization to help manage all of these efforts can provide requirements for how the various groups can work together or be complementary to one another and would focus efforts toward common goals. AFWIC, as the JADC2 development lead for the Air Force, would be well positioned to set requirements with the help of ACC A5C support. This oversight organization would expand beyond KR to ensure that JADC2 requirements are pursued. As a necessary enabling technology, AI/ML efforts should be under the purview of the AF JADC2 development lead.

Convergence Across C2 Construct, Data, and Algorithms

The development of MDO and supporting capabilities is an iterative process, as captured by the framework in Figure 5.2. The first step for advancing new JADC2 processes is to recognize gaps in existing processes and identify new capabilities and CONOPS to address them. This could be done, for example, by AFWIC and regional COCOMs. This demand signal feeds MDO development and informs the three enablers: C2 construct, data, and algorithms. Developments across these three foundational areas enable experimentation and adoption of new CONOPS, which in turn fuel the next iteration of identifying functional and technical needs. This process aligns MDO capabilities with supporting enablers and delivers capabilities incrementally. As

new threats or U.S. capabilities emerge, the process is repeated with the new systems, CONOPS, or technologies. In a sense, this acquisition model resembles Agile development practices currently applied to the AOC, but it generalizes to MDO and pairs top-level requirements with end-user needs.

Figure 5.2. MDO Capabilities Flow Chart
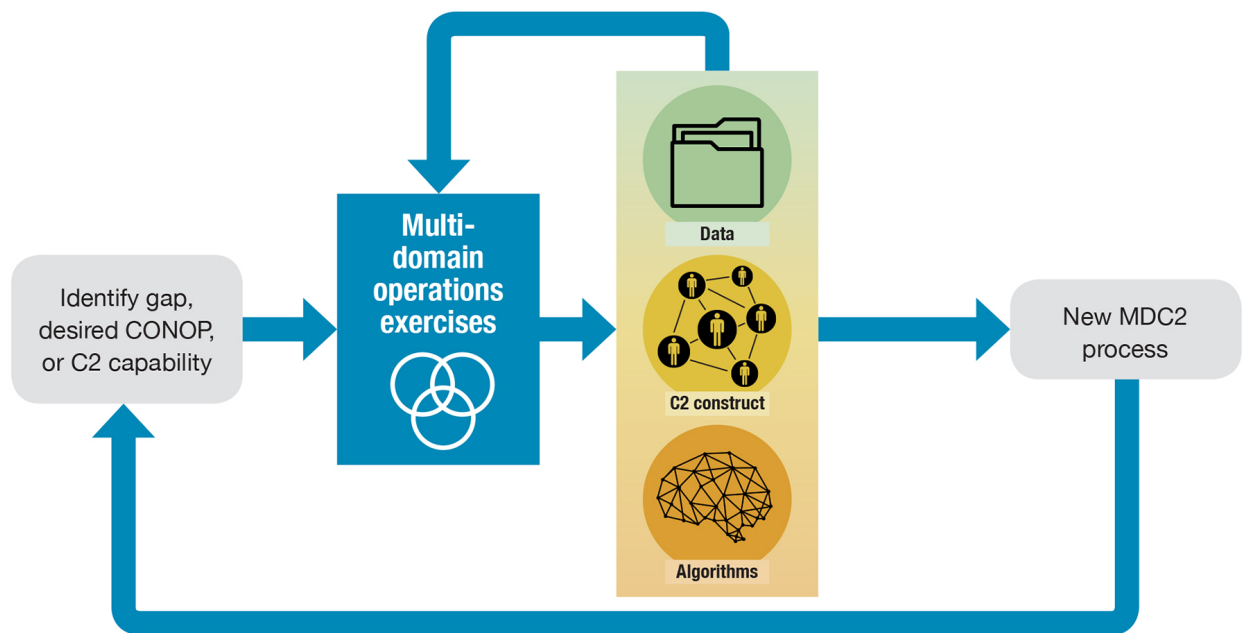
The details of the iterative steps under each line of effort are expanded in Figure 5.3. The identified MDO needs are provided as inputs into higher-level tabletop exercises for concept development. First, air, space, and cyber CONOPS can be explored in wargames by AFWIC with geographic and function command (space and cyber) representation, acknowledging that this approach is the incomplete picture. Next, sister service counterparts should participate in subsequent events to fully explore the best options across all domains. Finally, to refine the concepts, command-post and field-training exercises should be conducted. ShadowNet and the test and evaluation community for developmental and operational testing would be key participants in such events. Weapons and tactics conference topics would provide details about the tactical employment of these concepts.

Figure 5.3. Interactive Nature of JADC2 Progress
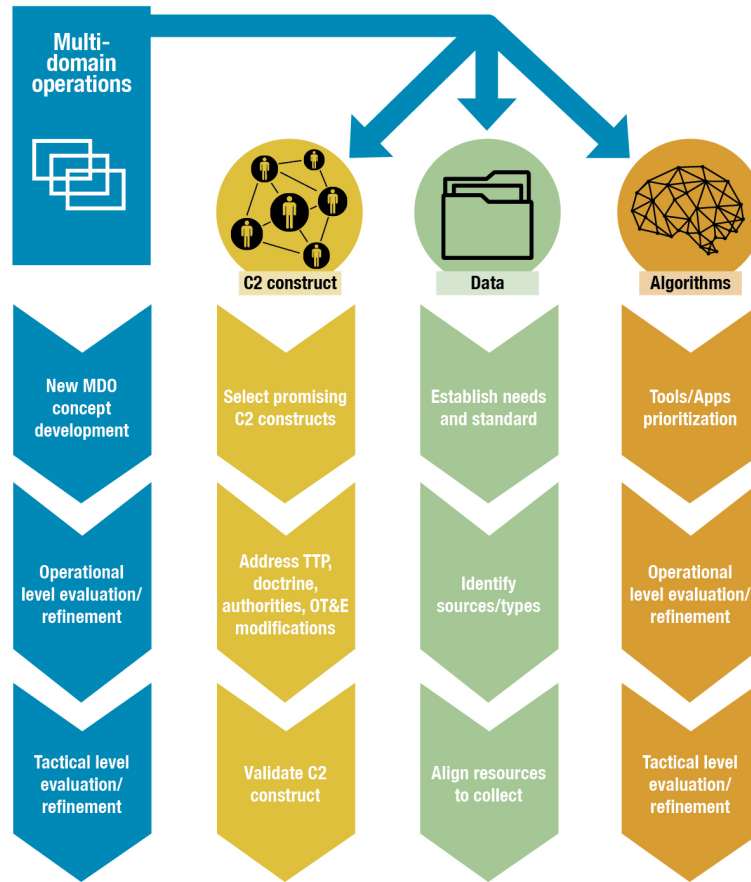
NOTE: OT&E = Operational Test and Evaluation.

The stream of MDO exercises, conducted at conceptual, operational, and tactical levels, would inform requirements for data, C2 construct, and the tools, apps, and algorithms needed (Figure 5.3). The framework allows feedback as progress is made on the three enablers, which should inform follow-on MDO exercises and modify MDO requirements. For example, a tabletop exercise could inform a preliminary set of requirements for an effective C2 construct, data (sources), data standards, and tools to process and display information. Once some of these requirements are met, even in part, the MDO concept is revisited to validate, replace, or refine the requirements. The outcome of the process is improved as additional gaps or capabilities are identified.

C2 Construct

As stated earlier in the report, the current centralized AOC structure with a 72-hour tasking cycle is problematic given modern threats; there are efforts to explore distributed control concepts to address these challenges. Alternative C2 constructs need to be evaluated to

understand the costs and benefits of each (see "C2 construct" in Figure 5.3). Alternative constructs should be considered as they support MDO. Starting within the Air Force, the 505 Command and Control Wing with AFWIC could evaluate alternative C2 constructs by holding wargames and command postexercises. Workshops would identify the resources, risks, and benefits of C2 alternatives in a structured manner. The results would be shared with the broader MDO community and would include engagement with the KM community. To address TTP, doctrine, authorities, and OT&E changes needed to enable MDO, the MDO CONOPS' impact across the DOTMLPF-P should be evaluated (middle arrow). The Air Force Doctrine Center, the Air Force Warfare Center, and the Air Force Weapons School are a few examples of stakeholders in this space. Findings here should inform future testing and experimentation by ShadowNet. The modeling and simulation community could evaluate candidate structures. At the same time, the emerging technologies needed to address gaps would inform the science and technology community. The 605th test and evaluation squadron would need to set up and execute a test plan for the selected C2 construct (bottom arrow).

Data

MDO will likely have a heavy data demand, and the AF CDO, working with the chief technology office, SAF/CIO, and SAF/AQ, will need to establish the data needs and standards for the MDO enterprise (top arrow, "Data" column). The KM community should help with setting, articulating, and implementing data standardization. Experimentation on data standardization will likely improve policy details. Data storage standards is another key element. To identify the sources and types of data (center arrow), the KMCWG should engage the AOCs. This line of effort will also need to work closely with the SAF/CIO office for approvals. Furthermore, resources need to be aligned to collect the data (bottom arrow). The OCs need to be informed of the data needs, standardization rules, and storage standards. The requirements should also be clearly articulated to the acquisition community to include the Program Management Office. The KMCWG and the AF CDO can advocate for resources (money and manning) to make this happen.

Algorithms

Algorithms are the last line of effort illustrated (far right column in Figure 5.3). The tools and apps for these capabilities should be prioritized to focus efforts to directly address the MDO CONOPS needs (top arrow, "Algorithms" column). Along this line, an AI algorithm-development plan for JADC2, a SaaS roadmap, and a ShadowNet testing schedule would all provide structure to enable moving forward in a coordinated fashion. Operational-level evaluation and refinement would occur with the ShOC or the CMCC for testing and operator workshops and events to identify automation and AI needs (center arrow). To tactically evaluate these tools, they should be incorporated in existing command post and live-fly exercises (bottom

arrow). User feedback collected by the 605th test and evaluation squadron would enable refinement.

To make future MDO a reality, there should be an enterprisewide strategy-to-tasks approach to get enabling capabilities down the same path. This requires multiple stakeholders at different echelons working to set policy, guidance, TTPs, training and exercising, infrastructure, and tools to operationalize these concepts. An enterprisewide approach is needed to move from the current situation in which, at the top, there are senior leader presentations on JADC2, and, at the bottom, science and technology and experimentation efforts develop narrow SaaS and DaaS examples, to a coherent and comprehensive enterprise push forward.

A truly MD process with supporting AI capabilities is still not fully captured here. Any Air Force capability will need to interface with capabilities of the other services and coalition and partner systems. Information-assurance concerns and policies for these other organizations will add further complexity to the enterprise. However, the strategy-to-tasks approach illustrated here can also aid stakeholders at different levels by framing the effort with whom, where, and when external stakeholders need to be engaged.

Imagine the MDO CONOPS and the enabling data, infrastructure, and tools, as well as how JADC2 is organizationally structured, developing and evolving in a cohesive, progressive way. Certainly, there will be roadblocks, failures, or even changes in emphasis, but the entire enterprise should adhere to an overarching strategy toward a common goal. Figure 5.3 illustrates how these lines of effort might interact and how this process is interactive in nature.

This report has provided one method for establishing an overarching strategy and method for aligning JADC2 key enablers under the goal of supporting MDO concepts for warfare against peer adversaries. The stakeholder base is broad and includes warfighters and others within the areas of evolving C2 construct, data, and algorithms. The challenge is significant, and such proposals will be needed to increase the chances of significant progress.

References

Air Force Instruction 13-1AOC, *Operational Procedures—Air Operations Center (AOC)*, Vol. 3, Washington, D.C.: Department of the Air Force, November 2, 2011, incorporating change 1, May 18, 2012.

Air Force Life Cycle Management Center, Battle Management Directorate, *Descriptive List of Applicable Publications (DLOAP) for the Air Operations Center (AOC)*, Hanscom Air Force Base, Mass., April 1, 2019, Not available to the general public.

AlphaStar Team, "AlphaStar: Mastering the Real-Time Strategy Game StarCraft II," deepmind.com, January 24, 2019. As of July 31, 2019:
https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/

Bellingham, John S., Michael Tillerson, Mehdi Alighanbari, and Jonathan P. How, "Cooperative Path Planning for Multiple UAVs in Dynamic and Uncertain Environments," *Proceedings of the 41st IEEE Conference on Decision and Control*, Las Vegas, Nev., December 2002. As of July 1, 2019:
https://www.semanticscholar.org/paper/Cooperative-path-planning-for-multiple-UAVs-in-and-Bellingham-Tillerson/8c9d12969607238d7071f2f45c99ab6bde88598d

Brocchi, Chiara, Davide Grande, Kayvaun Rowshankish, Tamim Saleh, and Allen Weinberg, "Designing a Data Transformation That Delivers Value Right from the Start," McKinsey and Company, October 2018. As of December 26, 2019:
https://www.mckinsey.com/industries/financial-services/our-insights/designing-a-data-transformation-that-delivers-value-right-from-the-start

Bughin, Jacques, Brian McCarthy, and Michael Chui, "A Survey of 3,000 Executives Reveals How Businesses Succeed with AI," *Harvard Business Review*, August 28, 2017. As of September 1, 2019:
https://hbr.org/2017/08/a-survey-of-3000-executives-reveals-how-businesses-succeed-with-ai?hootPostID=4247d7900bda767e984816fe65f00444

Chen, Chaofan, Oscar Li, Chaofan Tao, Alina Jade Barnett, Jonathan Su, and Cynthia Rudin, "*This* Looks Like *That*: Deep Learning for Interpretable Image Recognition," 33rd Conference on Neural Information Processing Systems, Vancouver, Canada, 2019. As of April 15, 2020:
https://arxiv.org/pdf/1806.10574.pdf

Compoc, Jeffery, *Obtaining MDC2 Capabilities Inside Technocratic Boundaries: Satisfying Multi-Domain Operational Desires Through an Acquiescent Technical Strategy*, Nellis Air Force Base, Nev.: Air Combat Command, 505 Test and Training Squadron, Combined Air and Space Operations Center–Nellis/Shadow Operations Center–Nellis, March 1, 2019.

DoD—*See* U.S. Department of Defense.

Ernest, Nicholas, David Carroll, Corey Schumacher, Matthew Clark, Kelly Cohen, and Gene Lee, "Genetic Fuzzy Based Artificial Intelligence for Unmanned Combat Aerial Vehicle Control in Simulated Air Combat Missions," *Journal of Defense Management*, Vol. 6, No. 1, 2016. As of August 1, 2019:
https://www.longdom.org/open-access/genetic-fuzzy-based-artificial-intelligence-for-unmanned-combat-aerialvehicle-control-in-simulated-air-combat-missions-2167-0374-1000144.pdf

Freedberg, Sydney J., "The Art of Command, the Science of AI," *Breaking Defense*, November 25, 2019. As of November 27, 2019:
https://breakingdefense.com/2019/11/the-art-of-command-the-science-of-ai/

Gao, Jingyue, Xiting Wang, Yasha Wang, and Xing Die, "Explainable Recommendation Through Attentive Multi-View Learning," presentation at the Association for the Advancement of Artificial Intelligence Conference on Artificial Intelligence, Arlington, Va., March 2019. As of April 15, 2020:
https://www.microsoft.com/en-us/research/publication/explainable-recommendation-through-attentive-multi-view-learning/

Goldfein, David, "CSAF Focus Area: Enhancing Multi-Domain Command and Control . . . Tying It All Together," Washington, D.C.: Chief of Staff, U.S. Air Force, March 2017. As of August 21, 2019:
https://www.af.mil/Portals/1/documents/csaf/letter3/CSAF_Focus_Area_CoverPage.pdf

Goldfein, David L., *Multi-Domain Command and Control (MDC2): Enterprise Capability Collaboration Team (ECCT), Campaign Plan, Strategy Document*, Washington, D.C.: U.S. Air Force, January 2, 2018, Not available to the general public.

Greene, Jay, and Aaron Gregg, "Amazon's Formal Challenge to Huge Pentagon Award Uses Videos That Mark Potential Influence Exerted by Trump," *Washington Post*, November 23, 2019. As of December 26, 2019:
https://www.washingtonpost.com/technology/2019/11/22/amazon-files-protest-pentagons-billion-jedi-award-under-seal/

Gregg, Aaron, "Pentagon Asks to Reconsider Part of the JEDI Cloud Decision After Amazon Protest," *Washington Post*, March 12, 2020. As of April 23, 2020:

https://www.washingtonpost.com/business/2020/03/12/pentagon-asks-reconsider-part-jedi-cloud-decision-after-amazon-protest/

David Gunning, *Explainable Artificial Intelligence (XAI)*, Arlington, Va.: Defense Advanced Research Projects Agency, 2016. As of May 13, 2020:
https://www.cc.gatech.edu/~alanwags/DLAI2016/(Gunning)%20IJCAI-16%20DLAI%20WS.pdf

Headquarters U.S. Air Force, Financial Management and Comptroller, "Air Force President's Budget FY21," 2019. As of September 3, 2019:
https://www.saffm.hq.af.mil/FM-Resources/Budget/

Hebert, Adam J., "Compressing the Kill Chain," *Air Force Magazine*, June 18, 2008. As of December 26, 2019:
https://www.airforcemag.com/article/0303killchain/

Jane's, "S-300V," *Land Warfare Platforms: Artillery and Air Defence*, last updated October 7, 2019. As of April 9, 2020:
https://janes.ihs.com/Janes/DisplayFile/JLAD0110

Joint Publication 3-0, *Joint Operations*, Washington, D.C.: Joint Chiefs of Staff, January 17, 2017, incorporating change 1, October 22, 2018. As of May 14, 2020:
https://www.jcs.mil/Portals/36/Documents/Doctrine/pubs/jp3_0ch1.pdf?ver=2018-11-27-160457-910

Joint Publication 3-01, *Countering Air and Missile Threats*, Washington, D.C.: Joint Chiefs of Staff, April 21, 2017, validated May 2, 2018. As of December 26, 2019:
https://www.jcs.mil/Portals/36/Documents/Doctrine/pubs/jp3_01_pa.pdf

Joint Publication 5-0, *Joint Planning*, Washington, D.C.: Joint Chiefs of Staff, June 16, 2017. As of August 19, 2019:
https://www.jcs.mil/Portals/36/Documents/Doctrine/pubs/jp5_0_20171606.pdf

Lawrence, Craig, "Adapting Cross-Domain Kill-Webs (ACK): A Framework for Decentralized Control of Multi-Domain Mosaic Warfare," Strategic Technology Office, Defense Advanced Research Projects Agency PDF briefing slides, July 27, 2018. As of July 23, 2019:
https://www.darpa.mil/attachments/ACK-Industry-Day-Briefing.pdf

Leggette, Wesley, and Michael Factor, "The Future of Object Storage: From a Data Dump to a Data Lake," IBM webpage, December 3, 2019. As of December 30, 2019:
https://www.ibm.com/cloud/blog/the-future-of-object-storage-from-a-data-dump-to-a-data-lake

Leskovec, Jure, Anand Rajaraman, and Jeff Ullman, *Mining of Massive Datasets*, 2nd ed., Cambridge, UK: Cambridge University Press, 2014.

Light, Thomas, Brian K. Dougherty, Caroline Baxter, Frank A. Camm, Mark A. Lorell, Michael Simpson, and David Wooddell, *Improving Modernization and Sustainment Outcomes for the 618th Air Operations Center*, Santa Monica, Calif.: RAND Corporation, 2019, Not available to the general public.

Lingel, Sherrill, Jeff Hagen, Eric Hastings, Mary Lee, Matthew Sargent, Matthew Walsh, Li Ang Zhang, Dave Blancett, Edward Geist, and Liam Regan, *Joint All-Domain Command and Control for Modern Warfare*: *Technical Analysis and Supporting Material*, Santa Monica, Calif.: RAND Corporation, 2020, Not available to the general public.

Macias, Amanda, "Pentagon Will Not Award JEDI Cloud Contract Until New Defense Secretary Completes Review," CNBC webpage, August 9, 2019.

McCullough, Amy, "USAF Looks to Create New Command and Control Structure," *Air Force Magazine*, June 6, 2018. As of May 13, 2020:
https://www.airforcemag.com/usaf-looks-to-create-new-command-and-control-structure/

McMurry, Robert D., "Use of Open Mission Systems/Universal Command and Control Interface," memorandum to the Air Force Program Executive Officers, Washington, D.C., October 9, 2018.

Moroney, Jennifer D. P., Stephanie Pezard, Laurel E. Miller, Jeffrey Engstrom, and Abby Doll, *Lessons from Department of Defense Disaster Relief Efforts in the Asia-Pacific Region*, Santa Monica, Calif.: RAND Corporation, RR-146-OSD, 2013. As of December 10, 2019:
https://www.rand.org/pubs/research_reports/RR146.html

Moseley, T. Michael, *Operation Iraqi Freedom: By the Numbers*, PDF slides, U.S. Air Forces Central Command, Prince Sultan Air Base, Kingdom of Saudi Arabia, April 30, 2003.

Russell, Stuart J., and Peter Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed., Upper Saddle River, N.J.: Prentice Hall, 2009.

Saltzman, Chance, "MDC2 Overview," briefing for 2018 C2 Summit, June 2018. As of July 23, 2019:
https://www.mitre.org/sites/default/files/publications/Special-Presentation-Gen%20Chance-Saltzman%20MDC2%20Overview%20for%20MITRE-June-2018.pdf

Under Secretary of Defense for Acquisition, Technology and Logistics, *Lessons Learned/Strengths: Operation Tomodachi / Operation Pacific Passage*, undated. As of April 9, 2020:
https://www.acq.osd.mil/dpap/ccap/cc/jcchb/Files/Topical/After_Action_Report/resources/Lessons_Learned_Operation_TOMODACHI.pdf

Under Secretary of Defense (Comptroller), "DoD Budget Request: Defense Budget Materials—FY2021," webpage, undated. As of April 23, 2020:
https://comptroller.defense.gov/Budget-Materials/

U.S. Air Force, *Operational Employment: Air Operations Center,* AFTTP 3-3.AOC, March 31, 2016a, Not available to the general public.

U.S. Air Force, *Air Superiority 2030 Flight Plan: Enterprise Capability Collaboration Team*, May 2016b.

U.S. Air Force, Virtual Distributed Laboratory, "Common Mission Control Center (CMCC)," September 27, 2017. As of September 3, 2019:
https://www.vdl.afrl.af.mil/programs/uci/cmcc.php

U.S. Congress, "Combined Congressional Report: 45-Day Report to Congress on JEDI Cloud Computing Services Request for Proposal and 60-Day Report to Congress on a Framework for all Department Entities to Acquire Cloud Computing Services," undated. As of December 26, 2019:
https://imlive.s3.amazonaws.com/Federal%20Government/ID15183034696552921558719522610265670631/180510_Final_Cloud_Combined_Congressional_Report__vg7_PDF_Redacted.pdf

U.S. Department of Defense, "JAIC and DSTA Forge Technology Collaboration," press release, Singapore, June 27, 2019.

U.S. Forces Japan, *Baseline Special Instructions (SPINs)*, Version 6, undated.

RAND PROJECT AIR FORCE

The authors examine and recommend opportunities for applying artificial intelligence (AI) and, more broadly, automation to deliberate planning for joint all-domain command and control (JADC2) for the U.S. Air Force.

The authors found that three primary enabling categories must be aligned to support future multidomain operations: (1) the command and control (C2) construct or how the forces are organized, where the authorities reside, and how they are trained and manned, (2) the data and data infrastructure needed to leverage data for C2, and (3) the tools, applications, and algorithms that leverage the data to C2 all-domain forces to include AI algorithms.

Moving to a modernized JADC2 requires various stakeholders to collaborate to set policy, guidance, tactics, techniques, procedures, training and exercising, infrastructure, and tools, likely leveraging AI, to operationalize concepts.

$21.00

www.rand.org